

An Introduction to Philosophy

W. Russ Payne PhD
Bellevue College

Copyright (cc by nc 4.0)
2023 W. Russ Payne

Permission is granted to copy, distribute and/or modify this document with attribution under the terms of Creative Commons: Attribution Noncommercial 4.0 International or any later version of this license. A copy of the license is found at <http://creativecommons.org/licenses/by-nc/4.0/>

Contents

Introduction	3
Chapter 1: What Philosophy Is	5
Chapter 2: Critical Thinking I, Being Reasonable.....	11
Chapter 3: Critical Thinking II: Logic.....	21
Chapter 4: Ancient Philosophy	41
Chapter 5: Rationalism	56
Chapter 6: Empiricism	68
Chapter 7: Philosophy of Science	79
Chapter 8: Philosophy of Mind	94
Chapter 9: Love and Happiness	101
Chapter 10: Meta Ethics	117
Chapter 11: Right Action	134
Chapter 12: Social Justice	146

Introduction

The goal of this text is to present philosophy to newcomers as a living discipline with historical roots. While a few early chapters are historically organized, my goal in the historical chapters is to trace a developmental progression of thought that introduces basic philosophical methods and frames issues that remain relevant today. Later chapters are topically organized. These include philosophy of science and philosophy of mind, areas where philosophy has shown dramatic recent progress.

This text concludes with four chapters on ethics, broadly construed. I cover traditional theories of right action in the third of these. Students are first invited first to think about what is good for themselves and their relationships in a chapter on love and happiness. Next a few meta-ethical issues are considered; namely, whether they are moral truths and if so what makes them so. The end of the ethics sequence addresses social justice, what it is for one's community to be good. Our sphere of concern expands progressively through these chapters. Our inquiry recapitulates the course of development into moral maturity.

Over the course of the text I've tried to outline the continuity of thought that leads from the historical roots of philosophy to a few of the diverse areas of inquiry that continue to make significant contributions to our understanding of ourselves and the world we live in.

As an undergraduate philosophy major, one of my favorite professors once told me that philosophers really do have an influence on how people think. I was pleased to hear that the kind of inquiry I found interesting and rewarding might also be relevant to people's lives and make a difference in the world. Then he completed his thought, "it only takes about 300 years." Over the course of my teaching career, it has struck me that the opinions many of my students come to class with have just about caught up with David Hume. So perhaps things are not quite as bad as my professor suggested. While Hume did publish young, he was still an infant 300 years ago.

My mission as a philosophy teacher has been to remedy this situation to some small degree. Most of the philosophy I read in graduate school was written by living philosophers, people I could meet and converse with at conferences. Every time I've done so I've come back with a new list of living philosophers I hoped to read. Philosophy has progressed as dramatically as the sciences over the last century or so. It is a great misfortune that the educated public by and large fails to recognize this.

Philosophers, no doubt, carry much of the blame for this. We have been better researchers than ambassadors. At no time in history have there been as many bright people doing philosophy as there are today. Clearly articulated fresh perspectives on important issues abound. But at the same time, philosophy's "market share" in the university curriculum has fallen to historic lows.

If the flourishing of philosophy over the past century or so is to continue, philosophy as a living discipline will have to gain a broader following among the general educated public. The front line for this campaign is the Philosophy 101 classroom.

This is an open source text. It is freely available in an editable, downloadable electronic format. Anyone is free to obtain, distribute, edit, or revise this document in accordance with the open source license. No one is free to claim proprietary rights to any part of this text. Academic publishing, both of research and textbooks, plays an unfortunate role in restricting access to information. This is quite against the spirit of free and open discourse that is the lifeblood of philosophy.

Introductory students should be exposed to as many philosophical voices as possible. To that end, links to primary source readings and supplemental material are imbedded in the text. I've restricted myself to primary source materials that are freely available on the Web. Students should require nothing more than a reliable Internet connection to access all of the required and recommended materials for this course. Limiting primary and supplemental sources in this way has presented some challenges. Classic sources are readily available online, though not always in the best translations. Many contemporary philosophers post papers online, but these are usually not intended for undergraduate readers. Most good philosophical writing for undergraduates is, unfortunately, proprietary, under copyright and hence unavailable for an open source course. This is my humble attempt to remedy that situation. The strength of an open source text is that it is continually open to revision by anyone who'd care to improve it. It can no doubt be improved upon and so I invite more capable thinkers and writers to do so.

Do feel free to email with questions or suggestions for improvements:
wpayne@bellevuecollege.edu

The latest version of this text is available at my blog:

[W. Russ Payne – A BC Commons Blog \(bellevuecollege.edu\)](http://w.russpayne.com)

Most recent update: 1-3-2023

W. Russ Payne

1. What Philosophy Is

What is philosophy?

Many answers have been offered in reply to this question and most are angling at something similar. My favorite answer is that philosophy is all of rational inquiry except for science. Perhaps you are tempted to think science is all of human inquiry. About a hundred years ago, many philosophers, notably the Logical Positivists, thought there was nothing we could intelligibly inquire into except for scientific matters. But this view is probably not right and here is a telling question: what branch of science addresses the question of whether or not science covers all of rational inquiry? If this question strikes you as puzzling, this might be because you already recognize that whether or not science can answer every question is not itself a scientific issue. Questions about the limits of human inquiry and knowledge are philosophical questions.

We can get a better understanding of philosophy by considering what sorts of things other than scientific issues humans might inquire into. Philosophical issues are as diverse and far ranging as those we find in the sciences, but a great many of them fall into (or across) one of three major branches of philosophy, metaphysics, epistemology, and ethics.

Metaphysics

Metaphysics is concerned with the nature of reality. Traditional metaphysical issues include the existence of God and the nature of human free will (assuming we have any). Here are a few metaphysical questions of interest to contemporary philosophers:

- What is a thing?
- How are space and time related?
- Does the past exist? How about the future?
- How many dimensions does the world have?
- Are there any entities beyond physical objects (like numbers, properties, and relations)? If so, how are they related to physical objects?

Historically, many philosophers have proposed and defended specific metaphysical positions, often as part of systematic and comprehensive metaphysical views. But attempts to establish systematic metaphysical world views have been notoriously inconclusive.

Since the 19th century many philosophers and scientists have been understandably suspicious of metaphysics, and it has frequently been dismissed as a waste of time, or worse, as meaningless. But in just the past few decades metaphysics has returned to vitality. As difficult as they are to resolve, metaphysical issues are also difficult to ignore for long. Contemporary analytic

metaphysics is typically taken to have more modest aims than definitively settling on the final and complete truth about the underlying nature of reality. A better way to understand metaphysics as it is currently practiced is as aiming at better understanding how various claims about the reality logically hang together or conflict. Metaphysicians analyze metaphysical puzzles and problems with the goal of better understanding how things could or could not be. Metaphysicians are in the business of exploring the realm of possibility and necessity. They are explorers of logical space.

Epistemology

Epistemology is concerned with the nature of knowledge and justified belief. Here are a few noteworthy epistemological questions:

- What is knowledge?
- Can we have any knowledge at all?
- Can we have knowledge about more specific matters, like the laws of nature, fundamental moral principles, or the existence of other minds?

The view that we can't have knowledge is called **skepticism**. An extreme form of skepticism denies that we can have any knowledge whatsoever. But we might grant that we can have knowledge about some things and remain skeptics concerning other issues. Many people, for instance, are not skeptics about scientific knowledge, but are skeptics when it comes to knowledge of morality. Later in this course we will entertain some skeptical worries about science and we will consider whether ethics is really in a more precarious position. Some critical attention reveals that scientific knowledge and moral knowledge face many of the same skeptical challenges and share some similar resources in addressing those challenges. Many of the popular reasons for being more skeptical about morality than science turn on philosophical confusions we will address and attempt to clear up.

Even if we lack absolute and certain knowledge of many things, our beliefs about those things might still be more or less reasonable or more or less likely to be true given the limited evidence we have. Epistemology is also concerned with what it is for a belief to be reasonable, or rationally justified. Even if we can't have certain knowledge of anything (or much), questions about what we ought to believe remain relevant.

Ethics

While epistemology is concerned with what we ought to believe and how we ought to reason, **Ethics** is concerned with what we ought to do, how we ought to live, and how we ought to organize our communities. It comes as a surprise to many new philosophy students that you can reason about such things. Religiously inspired views about morality often take right and wrong to be simply a matter of what is commanded by a divine being. Moral Relativism, perhaps the

most popular outlook among people who have rejected faith, simply substitutes the commands of society for the commands of God. Commands are simply to be obeyed, they are not to be inquired into, assessed for reasonableness, or tested against the evidence. Thinking of morality as so many commands based on the authority of God or society leaves no room for rational inquiry into how we ought to live, how we ought to treat each other, or how we ought to structure our communities. Philosophy, on the other hand, takes seriously the possibility of rational inquiry into these matters. If philosophy has not succeeded in coming up with absolutely certain and definitive answer in ethics, this is in part because philosophers take the answers to moral questions to be things we need to discover, not simply matters of somebody's say-so. The long and unfinished history of science should give us some humble recognition of how difficult and frustrating careful inquiry and investigation can be. So, we don't know for certain what the laws of morality are. We also don't have a unified field theory in physics. We are far more complicated than atoms, so why expect morality to be easier than physics?

So, we might think of **metaphysics** as concerned with "What is it?" questions, **epistemology** as concerned with "How do we know?" questions, and **ethics** as concerned with "What should we do?" questions. Many interesting lines of inquiry cut across these three kinds of questions. The philosophy of science, for instance, is concerned with metaphysical issues about what science is, but also with epistemological questions about how we can know scientific truths. The philosophy of love is similarly concerned with metaphysical questions about what love is. But it also concerned with questions about the value of love that are more ethical in character.

Assorted tangled vines of inquiry branch off from the three major trunks of philosophy, intermingle between them, and ultimately with scientific issues as well. The notion that some branches of human inquiry can proceed entirely independent of others ultimately becomes difficult to sustain. The scientist who neglects philosophy runs the same risk of ignorance as the philosopher who neglects science.

What is the value of philosophy?

Philosophy is a branch of human inquiry, and as such it aims at knowledge and understanding. We might expect the value of philosophy to be found in its results, the value of the ends that it seeks, the knowledge and understanding it reveals. But philosophy is rather notorious for failing to establish definitive knowledge on the matters it investigates. I'm not so sure this reputation is well deserved. We do learn much from doing philosophy. Philosophy often clearly reveals why some initially attractive answers to big philosophical questions are deeply problematic, for instance. But granted, philosophy often frustrates our craving for straightforward convictions. In our first reading, Bertrand Russell argues that there is great value in doing philosophy precisely because it frustrates our desire for quick easy answers. In denying us easy answers to big questions and undermining complacent convictions, philosophy liberates us from narrow minded conventional thinking and opens our minds to new possibilities. Philosophy often provides an

antidote to prejudice not by settling big questions, but by revealing just how hard it is to settle those questions. It can lead us to question our comfortably complacent conventional opinions.

The Value of Philosophy

Our first Reading is Chapter 15 of Bertrand Russell's *Problems of Philosophy*, "[The Value of Philosophy](http://www.ditext.com/russell/russell.html)." The whole book can be found here: <http://www.ditext.com/russell/russell.html>.

We humans are very prone to suffer from a psychological predicament we might call "the security blanket paradox." We know the world is full of hazards, and like passengers after a shipwreck, we tend to latch on to something for a sense of safety. We might cling to a possession, another person, our cherished beliefs, or any combination of these. The American pragmatist philosopher Charles Sanders Peirce speaks of doubt and uncertainty as uncomfortable anxiety-producing states. This would help explain why we tend to cling, even desperately, to beliefs we find comforting. This clinging strategy, however, leads us into a predicament that becomes clear once we notice that having a security blanket just gives us one more thing to worry about. In addition to worrying about our own safety, we now also have to worry about our security blanket getting lost or damaged. The asset becomes a liability. The clinging strategy for dealing with uncertainty and fear becomes counterproductive.

While not calling it by this name, Russell describes the intellectual consequences of the security blanket paradox vividly:

The man who has no tincture of philosophy goes through life imprisoned in the prejudices derived from common sense, from the habitual beliefs of his age or his nation, and from convictions which have grown up in his mind without the cooperation or consent of his deliberate reason. . . The life of the instinctive man is shut up within the circle of his private interests. . . In such a life there is something feverish and confined, in comparison with which the philosophic life is calm and free. The private world of instinctive interests is a small one, set in the midst of a great and powerful world which must, sooner or later, lay our private world in ruins.

The primary value of philosophy according to Russell is that it loosens the grip of uncritically held opinion and opens the mind to a liberating range of new possibilities to explore.

The value of philosophy is, in fact, to be sought largely in its very uncertainty. . . Philosophy, though unable to tell us with certainty what is the true answer to the doubts which it raises, is able to suggest many possibilities which enlarge our thoughts and free them from the tyranny of custom. Thus, while diminishing our feeling of certainty as to what things are, it greatly increases our knowledge as to what they may be; it removes the somewhat arrogant dogmatism of those who have never traveled into the region of liberating doubt, and it keeps alive our sense of wonder by showing familiar things in an unfamiliar aspect.

Here we are faced with a stark choice between the feeling of safety we might derive from clinging to opinions we are accustomed to and the liberation that comes with loosening our grip on these in order to explore new ideas. The paradox of the security blanket should make it clear what choice we should consider rational. Russell, of course, compellingly affirms choosing the liberty of free and open inquiry.

Must we remain forever uncertain about philosophical matters? Russell does suggest that some philosophical questions appear to be unanswerable (at least by us). But he doesn't say this about every philosophical issue. In fact, he gives credit to philosophical successes for the birth of various branches of the sciences. Many of the philosophical questions we care most deeply about, however - like whether our lives are significant, whether there is objective value that transcends our subjective interests - sometimes seem to be unsolvable and so remain perennial philosophical concerns. But we shouldn't be too certain about this either. Russell is hardly the final authority on what in philosophy is or isn't resolvable. Keep in mind that Russell was writing 100 years ago and a lot has happened in philosophy in the mean time (not in small part thanks to Russell's own definitive contributions). Problems that looked unsolvable to the best experts a hundred years ago often look quite solvable by current experts. The sciences are no different in this regard. The structure of DNA would not have been considered knowable fairly recently. That there was such a structure to discover could not even have been conceivable prior to Mendel and Darwin (and here we are only talking 150 years ago).

Further, it is often possible to make real progress in understanding issues even when they can't be definitively settled. We can often rule out many potential answers to philosophical questions even when we can't narrow things down to a single correct answer. And we can learn a great deal about the implications of and challenges for the possible answers that remain.

Suppose we can't settle some philosophical issue. Does that tell us that there is no right answer? No. That is not to say that every issue has a right answer. There is no answer to the issue of whether chocolate is better than vanilla, for instance. But when we can't settle an issue this often just tells us something about our own limitations. There may still be a right answer; we just can't tell conclusively what it is. It's easy to appreciate this point with a non-philosophical issue. Perhaps we can't know whether or not there is intelligent life on other planets. But surely there is or there isn't intelligent life on other planets. This question obviously has a right answer, we just haven't been able to figure out which it is. Similarly, we may never establish whether or not humans have free will, but, at least once we are clear about what we mean by "free will," there must be some fact of the matter. It would be intellectually arrogant of us to think that a question has no right answer just because we aren't able to figure out what that answer is.

Review and Discussion Questions

The first quiz covers this chapter and Bertrand Russell's essay "[The Value of Philosophy](#)." You will find a link to the quiz in the course module for this chapter. Watch the course calendar for

when to take the quiz. The following questions will help you prepare. Feel free to take these questions up on the discussion board.

On this Chapter:

- Why should we doubt that science covers all of human inquiry?
- What are some metaphysical issues? Some epistemological and ethical issues?
- What problem does the view that morality is simply a matter of the say-so of some authority lead to?

On Russell's "The Value of Philosophy":

- What is the aim of philosophy according to Russell?
- How is philosophy connected to the sciences?
- What value is there in the uncertainty that philosophical inquiry often produces?

On the commentary on Russell:

- Explain the "security blanket" paradox.
- How can understanding of issues be advanced even when definitive knowledge can't be had?
- What's the difference between saying we can't know the answer to some question and saying that there is no truth of the matter?

Finally, consider some of the definitions of philosophy offered by philosophers on the page linked at the opening of the lecture. A number of these would make for good discussion. Here's the link again: <http://www.brainpickings.org/index.php/2012/04/09/what-is-philosophy/>

Some Vocabulary from this Chapter

Metaphysics

Epistemology

Ethics

Skepticism

2. Critical Thinking I: Being Reasonable

The traditional goal in critical thinking courses is reasoning clearly and effectively to get at the truth and avoid falsehoods. Our conception of critical thinking can be expanded to include other goals and appropriate standards for achieving these. So, for instance, an architect isn't concerned with getting at the truth, but with good design. So, an architect will be concerned with standards and methods appropriate to achieving good design. But even in this context, it would be irrational for an architect to ignore truths about, say, the load bearing properties of various building materials. Whatever our goals, getting at the truth about many things will matter. So, truth oriented critical thinking in the context of inquiry is a vital and indispensable foundation for reasoning effectively.

Biases and confusions about the nature of truth and reason are widespread. So, we are going to start here by introducing the metaphysical and conceptual framework typically presupposed in any kind of inquiry. We will also discuss the personal traits and social conditions that are required for inquiry to proceed productively. In the next chapter we will get acquainted with the methods and skills employed in reasoning clearly and effectively. This will include an introduction to logical basics and some logical fallacies.

Subjects and Objects

Let's start with the modest metaphysical assumption that we all live on planet Earth. This means we have a shared reality. One that is populated with various and sundry objects (or better, containing stuff that can be divided into objects in any number of ways). This is the realm of objects, or the objective world. As embodied creatures, we are among the objects populating the objective realm. But in addition to being objects we are also subjects.

As subjects we have some experience of our shared reality, the objective realm. But, our experience of the world is limited by our perspectives. Further our impressions and beliefs are liable to be distorted by biases and assorted other intellectual bad habits. So, one thing we can all recognize about being subjects is that our impressions, beliefs and opinions are fallible. We are limited and imperfect in ways that make error quite possible. That's just life as a subject, having subjective impressions and beliefs means being prone to error. Fortunately, we can always expand the basis of evidence we reason from by sharing our impressions and beliefs with each other. And we can improve the reliability of our reasoning by cultivating the intellectual habits help us overcome biased and distorted ways of thinking. This is the point of critical thinking.

Philosophers typically use the word "subjective" to refer to what is mind dependent. This makes all of your thoughts, feelings and beliefs subjective. These are all aspect of you, a subject. But then your beliefs are often about the objective world. When I believe I have chocolate ice cream in the freezer, I'm representing a part of the objective world as being a certain way. So while it

exists in the subjective realm, my belief is also about the objective realm and it can accurately represent an aspect of the objective world. In which case my belief is true. Or it can fail to, say, if my son has finished the chocolate ice cream. In this case my belief is false.

So, your beliefs, are aspects of you, a subject, but they aim at representing things that are going on in our shared reality, the objective realm. For your belief to be true is just for your representation of how things are in the objective realm to fit, or correspond with, what is actually going on in the objective realm. That is, your belief is true when it represents some aspect of our shared reality the way it is.

Subjective realm includes whatever depends on your mind as a subject	Objective realm includes all aspects of our shared reality
All your thoughts, sense impressions, feelings, beliefs, fears and hopes are aspects of you as a subject.	The objective realm includes all the things, states of affairs and ways things are independent of you as a subject.

As subjects, it is generally good for us to have true beliefs and avoid false beliefs. When we have true beliefs, we are more capable of acting effectively, achieving goals, avoiding hazards, and generally having a good time. I suppose this is a value statement, but its not the sort of value statement anyone is likely to dispute. This much of the value of having true beliefs comes along with being subjects who have needs and goals in a world full of objects (and subjects) that can be helpful or harmful to us. Special cases aside, it's good to have true beliefs.

Truth

We just claimed that for your belief to be true is for it to represent things as they are. The basic idea here is that beliefs and claims are true when they correspond to how things are in our shared reality. Ordinary declarative sentences say something about how things are. So, take a few everyday examples:

- There is a spruce tree in Stuart's front yard.
- Lake Washington is east of downtown Seattle
- Your keys are on the kitchen counter.

Each of these sentences represents some aspect of our shared reality as being a certain way. The sentence is true if that part of the world is the way the sentence says it is. Since truth is about correspondence with our shared reality, it concerns the objective realm. Truth is objective.

Our society is currently riddled with some confusing ways of talking about truth. We have become accustomed to talking about "my truth" or "your truth." But if we stay focused on our ordinary understanding of truth as correspondence to reality, we can avoid confusions about truth being subjective or relative. Corresponding to our shared reality is obviously what we mean

when we count the belief that Lake Washington is east of Seattle as true. For my belief or claim to be true is just for it to represent some aspect of our shared reality as it is. What is true depends only on how things are objectively in our shared reality. Beyond shaping what is true about my own mind, I can't make things true merely by willing, wishing or believing them.

Given this ordinary everyday understanding of truth, it should be clear the truth doesn't belong to anyone. Nobody gets to dictate, define or decide what is the case, except in the very limited respect where a person decides what to do, how to think, or who to be. As a subject, I have this much power to shape our shared reality and no more. There is no "my truth" or "your truth." The only way to make sense out of truth being subjective would be to deny the existence of a shared reality. Truth could be subjective only if I live in my own little world and you live in your own separate reality. This would be to reject the very modest metaphysical assumptions we started this chapter with. I can't prove that we have a shared reality, an objective realm, but not having one sure sounds lonely.

When I believe something, I take it to be true. I suspect this is all most talk of "my truth" or "your truth" amounts to, a confusing way of talking about what we believe. But this kind of talk involves a rhetorical cheat in suggesting that my belief, which could well be just plain false, is still somehow to be associated with what is true. Talk of "my truth" and "your truth" blurs the difference between the subjective and the objective.

Before we leave the topic of truth, let's consider the difference between these two questions:

- What is it for a claim to be true?
- How do we determine that a claim is true?

It's important to keep these two questions separate. Questions about how we know whether something is true are epistemic questions. These questions are concerned with how our minds relate to the world. But the question of what it is for something to be true is not an epistemic issue. The truth of a claim is quite independent of how or whether we know it to be true. There are many truths we don't know and some of the things we think we know just aren't true. If you are not sure about this, consider these two claims:

- There is intelligent life on other planets.
- There is no intelligent life on other planets.

One of these claims is true. We can be sure of this on the basis of logic alone. Either claim being false would make the other true. We don't know which of these two claims is true and yet one of them is true. Whichever of these claims is true, its being true doesn't depend on whether we know it to be true. There are many truths that will never be known or believed by anyone, and appreciating this is enough to see that the truth of a claim is not relative to belief, knowledge, proof, or any other epistemic notion.

So, what it takes for a claim to be true doesn't depend on what we believe, or what we think we know, (except in the special case of claims about what we believe). What it takes for a claim to be true only depends on what on how things are. Once we get clear on subjects, objects and truth, the answer to our first question above is pretty clear. All it is for a claim to be true is for what it says to fit with how things are.

Naturally, most of us are more concerned with the second question: how we can determine when the claims people make and the things we believe are true? This question is more challenging. It's also what critical thinking is about. We will get to the issue how we can know in the next chapter when we turn do some logic and get acquainted with the basic methods for evaluating reasons and evidence. For now, we'll keep the focus on what it means to be reasonable and what value that brings.

Rationality

For your beliefs to be rational, or reasonable (we'll treat these terms as synonyms), is just for them to be held on the basis of the best available reasons. To be reasonable, in the literal sense of the word, is to be amenable to reason. That is, the reasonable person is the person you can reason with.

Good reasons are reasons that are truth-oriented. So, all it means for your belief to be rational, or reasonable, is for it to be held for the most truth-oriented reasons available to you. This much should make it clear why it is good to be rational. Being rational is more likely to get you true beliefs and true beliefs are good because they help you act effectively, achieve your goals, avoid hazards, and they give you a shared basis for understanding and communicating with others.

We should note that the words "rational" and "reasonable" can also refer to choosing or acting in ways that aim at maximizing some goal or value. So, for instance, the rational hedge fund manager is one who seeks a good return on investment. Words are often ambiguous. The way to be comfortable with ambiguity is to get clear on how words are being used and to track the various usages. Talk of rational or reasonable *belief* (as opposed to *choice* or *action*) can generally be understood as truth-oriented simply because to believe something is to take it to be true.

Rationality is not a kind of human imposed authority over what is true or what we should believe. The only thing that is authoritative concerning what we should believe is how things are in our shared reality. Again, to believe something is to take it to be true. To believe rationally is just to believe in a ways that target the truth well. To believe irrationally is to aim badly at the truth. Rational belief isn't guaranteed to hit the target of truth. But irrational belief involves a kind of unforced error.

Talk of rationality, objectivity and truth have some difficult connotations in the minds of many. These concepts often get associated with things like maleness, authority, power, and now even whiteness. These difficult associations appear to be based on antiquated stereotypes of one sort

or another. But thinking based on stereotypes is highly unreliable and perhaps we are in a position now to see how thinking in terms of stereotypes misleads many of us concerning these basic concepts. To be reasonable literally means to be amenable to good reasons (this especially includes the good reasons of others who think differently). To be a reasonable or rational believer involves a good measure of intellectual humility and a constant awareness of how easy it is to be misled in the pursuit of knowledge and understanding. So reasonable people are careful and cautious thinkers. Good critical thinkers get well acquainted with lots of logical fallacies, mistakes in reasoning, in order to avoid known pitfalls in thinking. The rational believer doesn't let ego and willfulness get in the way of her evaluation of claims and reasons. The reasonable or rational believer doesn't force things, but rather yields to best reasons and evidence. I won't delve into stock stereotypes here except to point out that being amenable to other perspectives, intellectually humble, cautious, and yielding to the better ideas don't stereotypically characterize maleness, authority or whiteness. But they do characterize reasonableness and rationality.

We started with two metaphysical assumptions here. Namely, that we have a shared reality, and that we each have limited and fallible experience of that reality. To this we've added a few definitional remarks about truth, rationality, reason, belief, subjectivity and objectivity. And we've reasoned a bit on the basis of these assumptions and definitions. In the definitional remarks I've tried to lay out standard philosophical usage clearly and straightforwardly.

The reason it is good to understand truth, rationality etc. in the manner I've laid out here is that it facilitates clearer communication and understanding of our diverse experiences and diverse ways of thinking. This allows us to cooperatively improve our ways of thinking and our limited grasp of what's true. And as a result of this, we are empowered to act more effectively, avoid hazards in our interactions, appreciate each other more significantly, and enjoy things.

Some will be tempted to object to what I've laid out here on the grounds that people are free to define things like truth and rationality as they please. In a sense, people are free to do so. Nobody has the power to prevent it. All the other concepts we might be tempted to attach to these words are out there, and as a linguistic community of a couple people or of a couple billion we could agree to name those concepts as we please. But to insist on defining things as we like amounts to the privatization of language, with the primary result of undermining our capacity to communicate with one another and understand each other in the limited ways that are open to us. I am no fan of such hyper-individualism. While we could quibble about how to define truth and rationality, the main result of this would be to talk about something else instead. Something other than how we stand as subjects to each other and our shared reality.

Philosophy as Inquiry

Philosophy is a branch of inquiry. Inquiry or investigation is the effort to figure something out, to get at the truth of some matter. Fruitful inquiry that produces lasting knowledge and understanding is typically a community affair. As individuals, all we have to work with is our own individual point of view based on our limited experience and whatever thought processes

we are used to or comfortable with. In this predicament, an individual has only new evidence to help check for errors. Worse, many of us are quite adept at only noticing the evidence that supports what we already think, which precisely misses all the evidence that might check for errors in our thinking. This is the problem known as confirmation bias. So, as individuals we are stuck with very limited evidence and little significant error checking. But as a community of inquirers, we can pool our diverse evidence and review our various thought processes. This is how inquiry has always proceeded. The current understanding in an area of science or philosophy is the result of a great many conversations, sometimes going back centuries or millennia, some in person, some in print.

Once a philosophical position is considered, we want to ask what arguments can be advanced in support of or against that position. In order to get at the best reasons, it is vital that our community of inquirers include people with diverse perspectives and diverse ways of thinking. Otherwise, we are liable to miss important evidence, overlook good reasons, or fail to catch our own mistakes. A community of like-minded people will tend to share the same blind spots, and as a result, ignore illuminating alternative perspectives, and neglect the error checking these can provide.

We then want to examine the quality of the arguments for and against a given position. Evaluating flawed arguments often points the way towards better arguments and the process of formulating, clarifying, and evaluating arguments continues.

This method of question and answer in which we recursively formulate, clarify, and evaluate arguments is known as **dialectic**. Dialectic looks a lot like debate, but a big difference lies in the respective goals of the two activities. The goal of a debate is to win by persuading an audience that your position is right and your opponent's is wrong. Dialectic, on the other hand, is aimed at inquiry. The goal is to learn something new about the issue under discussion. Unlike debate, in dialectic your sharpest critic is your best friend. Critical evaluation of your ideas and arguments brings new evidence and reasoning to light. The person you disagree with on a philosophical issue is often the person you stand to learn the most from (and this doesn't necessarily depend on which of you is closer to the truth of the matter).

Dialectic is sometimes referred to as the Socratic Method after the famous originator of this systematic style of inquiry. We will get introduced to some of Plato's dialogues chronicling the exploits of Socrates later. This will give you a good sense for how the Socratic Method works. Then watch for how the Socratic Method is deployed throughout the rest of the course.

The Fruits of Inquiry

We come to know what is true through inquiry. Sometimes this is as straightforward as making some observations. I know that it is sunrise by looking out the window. Sometimes inquiry is an involved process of formulating questions, identifying possible answers, formulating arguments that bear on these, and then critically evaluating the arguments in light of whatever evidence we

have to work with. The steps in this process may be repeated or elaborated as needed depending on the complexity of the issues raised.

Sometimes inquiry fails to yield definitive knowledge. Sometimes we don't have the evidence we need to settle an issue. And sometimes it is not so clear how to reason well from the evidence we have. As we heard from Bertrand Russell last week, philosophers often fail to uncover the truth of issues they inquire into. So, where inquiry yields no definitive right answer, what's the point of inquiry?

Inquiry bears many fruits even when it doesn't yield final answers. Inquiry can help us:

- clarify our questions
- distinguish different if closely related issues
- identify the plausible answers
- rule out some wrong answers
- appreciate the implications of some possible answers for other related issues
- increase our understanding of issues by doing some or all of the above

As we will see throughout the course, philosophical inquiry often follows a dialectical pattern where we propose a view, offer arguments for that view, critically evaluate the arguments, learn from our mistakes, rinse and repeat. Inquiry proceeds incrementally through a dialectical process of trial and error. As the Muslim philosopher, Alhazan, put it (around 1025):

The seeker after the truth is not one who studies the writings of the ancients and, following his natural disposition, puts his trust in them, but rather the one who suspects his faith in them and questions what he gathers from them, the one who submits to argument and demonstration, and not to the sayings of a human being whose nature is fraught with all kinds of imperfection and deficiency. Thus, the duty of the man who investigates the writings of scientists, if learning the truth is his goal, is to make himself an enemy of all that he reads, and, applying his mind to the core and margins of its content, attack it from every side. He should also suspect himself as he performs his critical examination of it, so that he may avoid falling into either prejudice or leniency

Notice in this rather militarized analogy that the discovery of truth happens when your attack fails and you “submit to argument and demonstration,” but not to human authority. The real action in this iterated process of dialectical inquiry happens in formulating and evaluating arguments. We'll get to this shortly, but first I want us to examine the personal traits and social conditions that lead to fruitful reasoning and inquiry.

Conditions for Critical Thinking

We've discussed how inquiry draws us towards knowledge and understanding through sharing and critically reviewing the ideas and arguments that emerge from diverse perspectives,

experiences, and ways of thinking. But this doesn't happen automatically. Diverse perspectives and ways of thinking can also drive polarization, conflict,

and hostility. In this section we want to explore how diversity can be harnessed to mutual benefit rather than mutual destruction. We've already alluded to a few of these conditions. Here I will list and characterize them more specifically. We'll start with the personal traits of effective critical thinkers, or just reasonable people. Then we'll consider the social conditions to be found in communities populated by reasonable people.

- **Fallibilism:** We mentioned earlier that as subjects we are fallible beings. Our evidence is limited and we are liable to make mistakes in reasoning. Given our nature as fallible subjects, we should never be entirely convinced that we have settled a matter once and for all. To be completely convinced we are right would lead us to neglect any further evidence and argument that might warrant some revision of our views.
- **Intellectual Humility:** Closely related to the idea of fallibilism is intellectual humility. Intellectual humility goes beyond merely recognizing our capacity for error. Fallibilism is not directly concerned with our attitudes towards others and their views and thought processes. Intellectual humility does concern these social factors. The intellectually humble person will keep ego out of their engagement with other people in inquiry. Pride and celebration of your excellence is fine in competitive contexts, but inquiry isn't a competition, it's a cooperative activity where respect for others is critical. Arrogance and pride are liable to drive others from the project of inquiry with the result of losing their insights and perspectives. Of course, there are times when someone understands more than others and it may be tempting to see arrogance in expertise. Dismissing expertise as arrogance, however, will be a failure of intellectual humility itself. Seeing arrogance in expertise is a self-protective way of propping up one's ego by judging another as flawed instead of trying to understand them and make good use of a learning opportunity. Bear in mind that genuine expertise is only acquired through the exercise of intellectual humility. This may be hard to see in people who have some hard-earned expertise, but even the smartest among us only move past ignorance by humbly yielding to the better argument.
- **Open Mindedness:** The open-minded person is open to fairly evaluating the reasons and evidence. Note that open mindedness focuses on our openness to reasons and evidence. A popular but misguided conception of open mindedness is that we should never have much confidence in our own beliefs but always grant that we are just as likely to be wrong as someone who disagrees with us. The problem with way of thinking about open mindedness is that the person who knows what they are talking about and holds a view with some conviction as a result of rigorous inquiry would not count as open minded. You might, for instance, encounter a climate change skeptic alleging that climate scientists are not open minded because they are unwilling to consider the possibility the warming of our climate is the result of sun spots. This is a fallacious attempt to undermine the science. We should hold

our beliefs with as much conviction as the best available reasons and evidence warrant. Often our reasons justify high levels of confidence, if not absolute confidence (see fallibilism).

- **Intellectual Courage:** Reasonable people, being open minded and intellectually humble, take the risk of discovering that they have things wrong once in a while. This can be hard. It's generally not pleasant to find to find that you are mistaken. It takes intellectual courage to bear this risk with grace. It helps to have a sense of humor here. It's best if curiosity and delight in discovery outweigh the dread we often feel about getting things wrong. But while critical thinking involves a degree of intellectual risk, it should not involve putting your personal safety on the line in any way. Critical thinkers attack ideas and arguments, not each other. If somebody attacks you, they are not being reasonable. It is possible for a person to feel attacked if they self-identify with an idea that comes under scrutiny. But feeling personally attacked when an idea you like faces criticism would be a failure of intellectual humility that results from investing ego into something that isn't you. You are not your ideas. You can change your mind about something without being personally harmed. When reasonable people do change their minds, it will not be because any other person is dominating or compelling them. Reasonable people change their own minds in response to compelling reasons, not domineering people.
- **Perseverance:** Even once we've acquired the traits discussed so far, clarifying and evaluating arguments can be challenging and frustrating work. You might feel this way about some things some things you've already read in this text. For an ounce of encouragement, bear in mind that confusion is often what it feels like to grow intellectually. Of course, sometimes things are confusing because they just don't make sense. But things that do make sense can feel confusing when they are novel, abstract or just complicated. Stick with it. That confusion is what it feels like to grow new neural pathways. You'll be smarter if you see it through. Even after 40 years of studying philosophy, I sometimes find myself feeling lost and confused in my first pass at reading the work of a philosopher I haven't studied before. Then in the second reading things will begin to make sense. Take some rest between passes. It also helps enormously to take notes on how terms are defined and how arguments are structured. Your brain will continue to sort things out even when you aren't actively reflecting on the material. By the third or fourth pass, maybe over the course of a few days, rich and clear understanding will emerge and you'll be wiser than you were before.

Now let's consider what a community of reasonable people who uphold these intellectual virtues will look like. I think it will be a community characterized by freedom from domination, tolerance and respect for diverse others, good humored civility, a healthy political capacity to deal with shared problems and challenges cooperatively and effectively, and intimacy in friendship. Let's consider each of these in more detail.

- **Freedom from domination:** Critical thinking provides a way of exploring, understanding and sometimes resolving differences between people. This is an alternative to bullying, manipulation, deceit and domination. Again, critical thinkers are responsive to good

reasoning and they cultivate intellectual defenses against rhetorical bullying and propaganda aimed at social control. Critical thinkers will resist dominating attempts to by-pass their own intellectual capacities through manipulation or deceit.

- **Tolerance and Respect for Diverse Others:** People who recognize their own fallibility and value intellectual humility will recognize that intolerance bars others from sharing their evidence and argument. This will introduce blind spots in inquiry and frustrate attempts to understand things and figure things out. Likewise disrespectful treatment of others is liable to drive them from participating in inquiry with the same result of ignoring potentially important evidence and argument. Intolerance and disrespectful treatment of others is literally a recipe for ignorance.
- **Politics:** I'm sure you have noticed how divisive politics is in America at the moment. Passionate conflict in politics often reflects a struggle for power aimed at sustaining or overcoming oppressive domination. But even here, conflict is driven and amplified by poor critical thinking. The political polarization in we current see in America is the result of people refusing to try to understand each other and evaluate each other's reasons and perspectives fairly. I'm afraid a great many Americans have become unreasonable people, disastrously poor critical thinkers. If we were better able to understand and evaluate each other's perspectives, we would be much more capable of finding common ground in addressing our shared problems. If we were better able to identify fallacies, mistakes in reasoning, we would be much less vulnerable to manipulation that divides us and undermines mutual understanding.
- **Friendship:** There may be no more basic human need than the need to be loved. As subjects, we are doomed to a sort of isolation. No other person, not matter how well they know you and care for you can share your subjectivity. We can only hope to understand each other to limited degrees. But I'd submit that the drive to charitably understand another people is itself a form of love.

This probably sounds idealistic to the point of being unrealistic. That is understandable given the current state of our world. We face multiple crises from political dysfunction to climate change and this engenders a great deal of fear and anxiety. In this state, critical thinking is not just intellectually challenging, but it is likely to feel emotionally remote as well. When people are fearful and anxious it natural to seek security in the familiar and defend that against all intrusions. The need for intellectual courage is all the more dire and may seem to carry with it a need for emotional courage just when this seems least available. What I want to suggest here, is that we can seek comfort and security not only in the familiar, but also in the project of building communities of critical thinkers. This obviously starts with cultivating our own critical thinking skills. And this may require loosening our grip on ideological security blankets. But as we saw last week in connection with Russell, clinging to opinions as a security blanket doesn't really provide security. A better strategy is to seek comfort and security in friends and loved ones. Critical thinking provides an avenue to expanding your community of friends and loved ones even across great differences of perspective.

3. Critical Thinking II: Logic

Philosophers, as we've previously mentioned, are mainly in the business of formulating, clarifying and evaluating arguments. This is how inquiry proceeds. In any realm of inquiry, this is how we determine what is true, when we can. An argument is a reason for thinking something is true. An argument consists of a set of premises which work together to provide a reason for accepting a conclusion as true. In this chapter we will get introduced to the basic standards and procedures for formulating, clarifying and evaluating arguments.

We've introduced the idea of an argument as a reason for believing something and most the chapter will focus on this primary function of arguments. But arguments are multifunction tools in inquiry and we will also want to discuss their various other uses along the way. Here are a few:

Arguments can be useful for

- Providing a reason for thinking their conclusions are true
- Clarifying our reasons
- Teasing out false premises
- Clarifying our own conceptual understanding
- Recognizing gaps on our own reasoning
- Understanding the views of others

We'll have to say more about how to clarify and evaluate arguments before explaining these points.

Arguments

The way to determine whether a claim is true or false, when this is possible, is to evaluate the evidence and argument for and against it. Sometimes good reasons take the form of simple observations. I have a good reason for thinking my bicycle has a flat tire when I see the tire sagging on the rim. But often the business of identifying and evaluating reasons is a bit more involved.

An argument is a reason for taking something to be true. Arguments consist of two or more claims, one of which is a conclusion. The conclusion is the claim the argument aims to establish as true. The other claims, there can be one or many, are the premises. The premises of an argument taken together are offered as a reason for believing its conclusion to be true.

Some arguments provide better reasons for believing their conclusions than others. In case you have any doubt about that, consider the following examples:

1. Sam is a line cook.

2. Line cooks generally have good of kitchen skills.
3. So, Sam can probably cook well.

1. Sam is a line cook.
2. Line cooks usually aren't paid very well.
3. So, Sam is probably a millionaire.

The premises in the first argument provide pretty good support for thinking Sam can cook well. That is, assuming the premises in the first argument are true, we have a good reason to think that its conclusion is true. The premises in the second argument constitute a pretty poor reason to think Sam is a millionaire. So, whether or not the premises of an argument support its conclusion is one important factor in evaluating an argument.

Now consider these examples:

1. Boston is in Massachusetts.
2. Massachusetts is east of the Rockies.
3. So, Boston is east of the Rockies.

1. Boston is in California.
2. California is west of the Rockies.
3. So, Boston is west of the Rockies.

Again, the first of these two arguments looks good, the second not so much. But the problem with the second of these arguments is different. The premises of both arguments provide good support for the conclusion. That is, in both arguments, if the premises were true, we'd have good reason for accepting the conclusion. In fact, for both arguments, if the premises were true, the conclusion would have to be true. So, in both of these arguments we have a good relation of logical support between the premises and the conclusion. But the first premise of the second argument just isn't true. Boston is not in California. So, the latter pair of arguments suggests another key issue for evaluating arguments. Good arguments have true premises.

That is pretty much it. A good argument is an argument that has true premises that support its conclusion. So, evaluating an argument involves these two steps:

- Determine whether or not the premises are true.
- Determine whether or not the premises support the conclusion (that is, whether we have grounds to think the conclusion is true if all of the premises are true).

Determining whether an argument's premises are true may involve evaluating further arguments in support of those premises. An argument might be the last link in a long chain of reasoning. In this case, the quality of the argument depends on the whole chain. And since arguments can have

multiple premises, each of which might be supported by further arguments, evaluating an argument might be more involved yet, since its conclusion is really supported by a rich network of reasoning, not just one link and then another. While the potential for complication should be clear, the basic idea should be pretty familiar. Think of the regress of “why” questions many of us tormented our parents with as children. Even at a young age we understood that the reasons for believing one thing can depend on the reasons for believing a great many other things.

However involved the network of reasons supporting a given conclusion might be, it seems that there must be some starting points. That is, it seems there must be some reasons for believing things that don’t themselves need to be justified in terms of further reasons. Otherwise, the network of supporting reasons would go on without end. The issue we are facing here is one of identifying the ultimate foundations of knowledge and justified belief. This is a big epistemological issue and we will return to it later in the course. For now, let’s consider one potential answer we are already familiar with. In the sciences our complex chains of reasoning seem to proceed from the evidence of the senses. We think that evidence provides the foundation for our edifice of scientific knowledge. Sounds great for science, but where does this leave philosophy? Does philosophy entirely lack evidence on which its reasoning can be based?

Philosophy does have a kind of evidence to work from and that evidence is provided by philosophical problems. When we encounter a problem in philosophy this often tells us that the principles and assumptions that generate that problem can’t all be correct. This might seem like just a subtle clue that leaves us far from solving the big mysteries. But clues are evidence just the same. As we will discuss in our chapter on the philosophy of science, science doesn’t really have it much easier. Sensory evidence by itself doesn’t tell us as much about the nature of the world as we often suppose. Scientific evidence provides clues, but there remains a good deal of problem solving to do in science as well as in philosophy.

So, we can assess the truth or falsity of the premises of an argument by examining evidence or by evaluating further argument in support of the premises. Now we will turn to the other step in evaluating arguments and consider the ways in which premises can support or fail to support their conclusions. The question of support is distinct from the question of whether the premises are true. The reason one of our arguments about Sam the line cook was good but not the other had nothing to do with false premises. We can grant that the premises in both arguments were true. The difference had to do with whether the premises provided good support of the conclusion. When we ask whether some premises support a conclusion, we are asking whether we would have good grounds for accepting the conclusion if we assume that the premises are true. It is important that we keep the two steps in evaluating arguments distinct in our minds. When we evaluate arguments holistically, as people often do, we wind up accepting or rejecting arguments based on how we feel about them overall without looking into whether the premises of the arguments really support the conclusions we draw. This is one of the ways we fall victim to confirmation bias, by endorsing just the arguments that point towards the conclusions we like without scrutinizing the logic of the argument.

Consider again the two good arguments in our examples above:

1. Sam is a line cook.
2. Line cooks generally have good of kitchen skills.
3. So, Sam can probably cook well.

In this example the premises do support the conclusion. We have pretty good reason to think Sam can cook well if he is a line cook. But these premises don't guarantee that Sam can cook well. It might be his first day on the job. He might be a really lousy line cook. Or he might be a breakfast cook and pretty useless in the kitchen beyond frying eggs and making hash browns. Still, the premises of this argument would give us good reason for trusting him with dinner. The premises being true would make it pretty likely he'd feed us well.

Now consider this one again:

1. Boston is in Massachusetts.
2. Massachusetts is east of the Rockies.
3. So, Boston is east of the Rockies.

In this argument the premises don't just make the conclusion likely. The premises being true would guarantee the truth of the conclusion. These two examples point us towards our two standards of support, deductive validity and inductive strength. A deductively valid argument is one where the premises, if they are true, would guarantee the truth conclusion. The support relation in the case of deductively valid arguments is logically necessary. Inductively strong arguments are arguments where the premises, if they are true, would provide good reasons for thinking the conclusion is true. But good reasons in inductively strong arguments are a matter of probability, not necessity. A strong inductive argument with true premises doesn't guarantee the truth of the conclusion.

Deductive Validity

The deductive standard of support is validity. An argument counts as deductive whenever its aiming at validity. Deductive validity is the strictest standard of support we can uphold. In a deductively valid argument, the truth of the premises guarantees the truth of the conclusion. Here are two equivalent definitions of deductive validity:

(D) A valid argument is an argument where if its premises are true, then its conclusion must be true.

(D') A valid argument is an argument where it is not possible for all of its premises to be true and its conclusion false.

Here are a few examples of deductively valid arguments

1. If Socrates is human, then Socrates is mortal

2. Socrates is a human.
3. Therefore, Socrates is mortal

1. All monkeys are primates
2. All primates are mammals
3. So, all monkeys are mammals

If you think about these two examples for a moment, it should be clear that there is no possible way for the premises to all be true and the conclusion false. The truth of the conclusion is guaranteed by the truth of the premises. In contrast, the following argument is not valid:

1. If Sue misses her plane, she will be late for the conference.
2. Sue is late for the conference.
3. Therefore, she missed her plane.

Again, to say that an argument is deductively valid is to say that it is impossible for all of its premises to be true and its conclusion to be false. To see why the last argument is not valid, try to think of a possible scenario that makes both of the premises true and the conclusion false. One scenario is where Sue catches her plane, but her cab from the airport gets stuck in traffic. Another would be where Sue makes her flight, but the plane is delayed due to bad weather. If we can think of any possible way for the premises of an argument to be true and its conclusion false, then we have shown that the conclusion does not deductively follow from the premises. That is, we've shown that the argument is not valid.

Our intuitive test for validity is to think about whether it is possible for the argument's premises to be true and its conclusion to be false. A key point to notice here is that validity is not directly about the truth or falsity of the premises or the conclusion. The concept of validity is really a concept about what is and isn't logically possible. A deductively valid argument may or may not have true premises. Consider this argument:

1. All planets are stars.
2. All stars are bodies that shine steadily.
3. All planets are bodies that shine steadily.

Both of the premises in this argument are false, but the argument is still valid. Suppose, contrary to fact, that the premises were true. The conclusion would have to be true if this were the case. Validity isn't about whether the premises or the conclusion are in fact true. It is only about whether the conclusion logically follows from the premises.

Given this, a deductively valid argument only provides one with a good reason for believing its conclusion *if its premises are true*. If a deductively valid argument has all true premises, we say that it is deductively sound. For an argument to be deductively sound is one way for it to pass

both steps for evaluating arguments. A deductively sound argument has all true premises that support its conclusion.

The deductive arguments we've looked at here are pretty intuitive. We only need to think about whether the conclusion could be false even if the premises were true. But most deductive arguments are not so obvious. Logic is the science of deductive validity. Philosophy has made some historic advances in logic over the past century. Bertrand Russell, who we got acquainted with in the last chapter, was among the key contributors to developments in logic early in the 20th century.

Logical Form

Since Aristotle, the first major logician, it's been recognized that deductive validity is a matter of an argument's logical form. We can display an argument's logical form by replacing all but the logically operative vocabulary with symbols (we'll use capital letters for this). So, consider the logical form of a few of our examples so far.

1. All planets are stars.
2. All stars are bodies that shine steadily.
3. All planets are bodies that shine steadily.

This argument has the following form:

1. All P are S
2. All S are B
3. All P are B

Any argument that has this logical form will be valid. So,

1. All fish are vertebrates.
2. All vertebrates are animals.
3. So, all fish are animals.

Remember, validity is just a standard of support. Validity does not assume true premises or a true conclusion. So even though it sounds a bit "off," this argument is also valid:

1. All red things are bricks,
2. All bricks are rocket ships.
3. So, all red things are rocket ships.

Of course, this argument sounds silly. Both premises are ridiculously false. But then any possible world where both premises are true would be a possible world where all red things are rocket ships. The argument is valid in virtue of its valid logical form. Now consider this familiar argument:

1. If Socrates is human, then Socrates is mortal

2. Socrates is a human.
3. Therefore, Socrates is mortal

This argument has the following logical form:

1. If H, then M
2. H
3. M

Similarly, any argument that has this logical form will be valid. Plug any declarative sentences you like in for H and M and you will have a valid argument. The premises might be false, or even absurd, but it will remain the case that any way the world could be that makes both premises true will also make the conclusion true. Once you appreciate how deductive validity is a function of the logical form of an argument, it soon becomes clear that a valid argument can be constructed for any possible conclusion, true, false, or completely absurd. So, for instance:

1. If pigs fly, then the oceans will dry up.
2. Pigs fly
3. Therefore, the oceans will dry up.

So, you might be wondering what the point of all this silliness is. It's partly to limber up your logical sense and help you recognize that logical validity is only about what follows from what, not about what is in fact true or false. Of course, the oceans aren't going to dry up. But if both premises were true, then the conclusion would follow logically and also be true. But there is a further point to the hypothetical silliness. The fact that the conclusion of the "pigs fly" argument is absurdly false is a good indicator that at least some of the premises of this valid argument are also false. And this is a very useful thing to recognize. To see this, let's look at another valid argument pattern that captures what we've just said about the pigs fly argument:

1. If P, then C
2. Not C
3. So, not P

This is a valid pattern of reasoning that we use routinely. For instance:

1. If I have milk, then it will be in the fridge
2. There's no milk in the fridge
3. So, I am out of milk.

Now notice how we used this pattern of reasoning in our analysis of the "pigs fly" argument. It is valid, which means that if its premises are all true, then its conclusion is true. But obviously, its conclusion isn't true. So, its premises are not all true.

Reductio ad Absurdum

I mentioned near the beginning of this chapter that arguments are multifunctional tools in inquiry. Arguments aren't always used directly to show the truth of some conclusion. As we've just seen, the concept of a valid argument can be used to tease out falsity in the premises. For instance, we might consider a claim that sounds pretty good and ask what follows from that claim deductively. What conclusion could we validly argue for on the basis of that claim? If we find that by deductively valid reasoning we can get from our claim that sounds pretty good to an absurd conclusion, then we have shown that our starting point, the claim that sounded pretty good, is false. This strategy is known as *reductio ad absurdum*, which is a handy bit of Latin for "reducing to absurdity." We can use this strategy to test an idea for problems by considering what follows from that idea by valid argument and making sure it doesn't lead to anything obviously false or absurd.

To illustrate *reductio ad absurdum*, let's consider a view we mentioned briefly in the first chapter of this text about the nature of morality. A view that many people find attractive: moral relativism. According to moral relativism, there are no objective moral standards, rather morality is relative to groups depending on what is considered right in that group. When we consider what follows from moral relativism deductively, we wind up with some pretty unsavory results. The first premise in this argument is just a statement of moral relativism as a view about the nature of morality. From here, bad things start to happen.

1. If a society considers something morally good, then it is morally good (relative to that society).
2. Nazi Germany considered the extermination of Jewish people good.
3. The extermination of Jewish people was good (relative to Nazi Germany).

The argument here is valid. It's logical form is a minor variation on a valid pattern we examined above. If the premises are true, then the conclusion must be true. This means that if the conclusion is false, at least one of the premises must be false. Note that the conclusion here is not about what the Nazis considered to be good, it's about what *is good* in the only sense that matters according to moral relativism. Since the conclusion of this argument is obviously false, not to mention horrible, and the second premise is a matter of historical fact, moral relativism must be false. Here we have reasoned validly from a view about the nature of morality that many people find attractive to a conclusion that is obviously absurdly false and horrible. A view about the nature of morality that has obviously and horribly bad logical consequences has got to be false. Moral relativism reduces to absurdity, *reductio ad absurdum*. We will get to examine moral relativism in greater detail when we get around to studying ethics. The point here is just to see how deductive argument is useful not just for getting at true conclusions, but also for teasing out false premises using the strategy we'd just identified as *reductio ad absurdum*.

Revealing Hidden Assumptions

Another very helpful function of valid argumentation is as an aid to revealing hidden assumptions. In everyday life we don't generally formulate formally valid arguments when we

give reasons for what we believe. And often these unmentioned premises are where our biases hide. A good understanding of deductively validity can help us bring these hidden assumptions and biases to light. For example:

1. Every story I hear about politician X on Facebook says he is doing a terrible job.
2. So, Politician X is doing a terrible job.

This is not a valid argument as it stands. We would have a valid argument if we added a premise as follows.

1. Every story I hear about politician X on Facebook says he's doing a terrible job
2. If every story I hear about politician X on Facebook says he's doing a terrible job, then politician X is doing a terrible job.
3. So, Politician X is doing a terrible job.

Now we have a valid argument, but adding the premise required to have a valid argument reveals a hidden assumption that, as some of you probably know, we have reason to doubt. Facebook uses AI and algorithms to customize what you see in order to maximize engagement. It turns out that anger is very engaging. As a result, Facebook tends to feed you news stories that will stoke your anger. If you have a history of clicking and commenting on stories that say awful things about politician X or others of his political persuasion, Facebook will load your news feed with more articles that say awful things about politician X. The same goes for Google, YouTube and most search engines and social media platforms. The reason Facebook keeps showing you stories about what a terrible job politician X is doing isn't that politician X is actually doing a terrible job. Rather its that the AI behind Facebook algorithms knows that stories like this will keep you glued to the platform, posting angry comments about politician X.

So, formulating deductively valid arguments brings our assumptions to light where they can be scrutinized for truth or reasonableness. A good understanding of deductive validity can be very useful in identifying and addressing our usually unspoken assumptions and biases (well, perhaps in the example we just considered, the bias lies mainly in the AI driven algorithms employed by Facebook).

Clarity

For reasons we just been discussing, a good understanding of deductive validity can help you clarify your own reasons and express them clearly to others. If you can recognize when an argument you find appealing has some deductive gaps in it, your understanding of validity will guide you in filling in those gaps. Assuming the argument is a good one, you will then have a clearer understanding of it and be able to express your reason more clearly to others. Of course, as just discussed, if your argument is not so great, your understanding of validity can alert you to this and perhaps guide you in formulating a better argument. All of this applies to understanding the arguments offered by others. When others formulate their reasons in incomplete, less than

valid ways, your understanding of validity can guide you in identify the questions you'd want to ask in reconstructing a more complete version of your friend's argument.

Charitable Interpretation

A good understanding of deductive validity will help you formulate the clearest and best possible arguments for your view. It will also help you formulate and appreciate the best possible arguments for views you oppose. Formulating the best possible interpretation of and argument for opposing views is what we call "charitable interpretation. In the dialectical spirit of cooperatively working towards truth and reasonableness, it is best to be charitable in filling out your reconstruction of another's reasons. While helping to clarify arguments is a kindness, this is isn't really the point of charitable interpretation. Finding faults with bad arguments for a view you disagree with doesn't really undermine that view, it just undermines bad arguments for that view. If you have good reasons for rejecting a view, you should aim to identify the flaws in the best possible arguments for the view you think is wrong. Trying to make the opposing view sound outrageous or ridiculous will only lead you into straw man attacks (see discussion of the fallacy below). The strongest argument you can offer against a view you oppose is not the argument that makes the view sound outrageous, but the argument that targets the best interpretation of the opposing view and the best possible arguments for it. Being a reasonable and effective critical thinker calls for charitable interpretation of opposing views and the arguments for them, not just out a sense of fair mindedness, good as that may be, but also in order to be the most effective critic of the view you oppose.

In the last few sections, I've tried to characterize a few useful functions for deductive argumentation beyond merely trying to give reasons for thinking that something is true. While sound arguments, arguments that are both valid and have all true premises, provide good reasons for accepting their conclusions as true, valid argument isn't only useful for directly getting at the truth. A good understanding of validity is useful for clarifying reasoning and bringing hidden assumptions to the surface. It can be useful for drawing our attention to false premises. And it can help us make our criticism of views we oppose more effective by targeting the best versions of those views. The path to knowing truths and understanding issues is usually not a straight paved sidewalk. It takes some skill to recognize the switchbacks and stay on the trail. A good understanding of deductive validity is powerful guide.

Inductive Strength

I won't have as much to say about inductive strength and cogency since you will already be more familiar with it from science classes and because philosophy trades more in deductive argument. Let's start with our example argument from above:

1. Sam is a line cook.
2. Line cooks generally have good of kitchen skills.
3. So, Sam can probably cook well.

This is a decent argument. The premises do support the conclusion. And yet it might be that both premises are true and the conclusion is false. Sam could be a brand new cook, hired because he's the manager's son, but also someone who has never cooked in his life. Many arguments give us good reasons for accepting their conclusions even when true premises don't guarantee the truth of the conclusion. This suggests that we need another standard of support for arguments that aim at giving us pretty good but not absolutely compelling grounds for accepting their conclusions. And this standard of support is called inductive strength. Here are two equivalent ways of defining inductive strength:

- (I) An inductively strong argument is an argument in which if its premises are true, its conclusion is probably to be true.
- (I') An inductively strong argument is an argument in which it is improbable that its conclusion is false given that its premises are true.

If you look again at the earlier definitions for deductive validity you will find a good deal of similarity. The only difference is in the use of the words "probably" rather than "must be" in the first definition, and "improbable" rather than "impossible" in the second. This is a big difference. As in the case of validity, when we say that an argument is strong, we are not assuming that its premises are true. We are only claiming that *if* the premises are true then the conclusion is likely to be true. Corresponding to the notion of deductive soundness, an inductive argument that is both strong and has true premises is called a *cogent* inductive argument. Unlike the case of deductively sound arguments, it is possible for an inductively cogent argument to have true premises and a false conclusion.

What makes an argument an inductive argument is that it is aiming at the standard of inductive strength. Similarly, what makes an argument a deductive argument is aiming at validity. Students frequently ask if an invalid deductive argument can be considered inductively strong. Generally, not. The targets are different. Missing the target of deductive validity doesn't make an argument inductively strong. Invalid deductive arguments are generally just bad arguments. Once in a while there will be a decent inductive argument that looks a bit similar, but not typically. Deductive and inductive refer to different kinds of reasoning.

Lots of good reasons for holding a belief fall short of the standard of deductive validity. The sort of reasoning you were taught as "the scientific method" is inductive reasoning. As it is taught in high school, the scientific method consists of formulating a general hypothesis and testing it against a large sampling of data. If the data is consistent with the hypothesis, then the hypothesis is considered confirmed by the data. Here a limited amount of evidence is taken to support a broader more general hypothesis. In the simplest case, inductive reasoning involves inferring that something is generally the case from a pattern observed in a limited number of cases. For instance, if we were to conduct a poll of 1000 Seattle voters and 600 of them claimed to be Democrats, then we could inductively infer that 60% of the voters in Seattle are Democrats. The

results of the poll give a pretty good reason to think that around 60% of the voters in Seattle are Democrats. But the results of the poll don't guarantee this conclusion. It is possible that only 50% of the voters in Seattle are Democrats and Democrats were, just by luck, over represented in the 1000 cases we considered.

When evaluating deductive arguments for validity we ask if it is possible for the premises to be true and the conclusion to be false. This is either possible or it isn't. Possibility does not admit of degrees. But probability does. The truth of the conclusion of an inductive argument can be probable to a greater or lesser degree. An argument either is or isn't valid. But inductive arguments can be more or less strong. We can identify a few factors that bear on the degree of strength an inductive argument has. One is how much evidence we have looked at before inductively generalizing. Our inductive argument above would be stronger if we drew our conclusion from a poll of 100,000 Seattle voters, for instance. And it would be much weaker if we had only polled 100. Also, the strength of an inductive argument depends on the degree to which the observed cases represent the makeup of the broader class of cases. So our inductive argument will be stronger if we randomly select our 1000 voters from the Seattle phone book than if they are selected from the Ballard phone book (Ballard being a notably liberal neighborhood within Seattle).

So far, we've only discussed inductive generalization, where we identify a pattern in a limited number of cases and draw a more general conclusion about a broader class of cases. Inductive argument comes in other varieties as well. In the example we started with about Sam the line cook, we inductively inferred a prediction about Sam based on a known pattern in a broader class of cases. Argument from analogy is another variety of inductive reasoning that can be quite strong. For instance, I know that my housecat is very similar to cougars in the wild. Knowing that my cat can jump great heights, it would be reasonable to expect that by analogy, or based on this similarity, cougars can jump well too.

There are further varieties of argument that aim at the standard of inductive strength, but we will discuss just one more in detail now. Abduction is inference to the best explanation. Detective work provides a good example of abductive argument. When Holmes discovers Moriarty's favorite brand of cigar and a bullet of the sort fired by Moriarty's gun at a murder scene, inference to the best explanation suggests that Moriarty was the killer. That Moriarty committed the murder provides the overall best explanation of the various facts of the case.

The 19th century American pragmatist and logician, Charles Sanders Peirce offers the *Surprise Principle* as a method for evaluating abductive arguments. According to the surprise principle, we should count one explanation as better than competing explanations if it would render the facts we are trying to explain less surprising than competing explanations. The various clues in the murder case are among the facts we want explained. The presence of the cigar and the bullet casing at the murder scene is much less surprising if Moriarty committed the murder than if the maid did it. Inference to the best explanation aims at strength. So a strong abductive argument in

this case needn't rule out the possibility that the murder was committed by Moriarty's evil twin who convincingly frames his brother. There might an argument against the death penalty lurking nearby. Inference to the best explanation is worth more attention than it often receives. This kind of reasoning is pervasive in philosophy and science, but seldom gets much notice as an integral part of the methods of rational inquiry.

Fallacies

Early on in the exploration of reasonableness we made a point of acknowledging basic human fallibility. Inquiry is not a linear path from absolute truth to absolute truth. Inquiry is a more typically a meandering path with frequent back tracking as we learn from our mistakes. Our conclusions, even when they support a healthy degree of confidence, remain always provisional. New evidence or argument may reveal previously unrecognized mistakes. Of course, learning from our mistakes does require that being able to recognize them. Many of the mistakes in reasoning we humans are prone to are well known. These are fallacies. A fallacy is just a mistake in reasoning. Assuming we've developed a decent understanding of what good reasoning looks like over the prior chapters, we should now be in a position to examine some common fallacies and understand why they are mistakes.

A fallacious argument fails to support its conclusion. This is all that we can conclude when we find that an argument contains a fallacy. Finding an argument to be fallacious does not in itself provide a reason for rejecting its conclusion. There might be other good arguments for that conclusion or good arguments against it. The value in fallacy spotting is that it gets tempting but bad arguments out of the way and thereby helps us get a clearer view on issues. Reasonable people won't want to believe false things for bad reasons or true things for bad reasons. Bad reasons distort our understanding of the world, often in ways that indulge biases or prejudices, as we'll see in a few examples below.

I will only discuss a choice selection of fallacies here. A full course in critical thinking would introduce you to many more and include lots of practice at identifying them, first in text book exercises, then "in the wild."

- ***Ad hominem***: This fallacy is known in Latin for "against the person." As the name suggests, *ad hominem* consists of attacking the proponent of a position rather than critically evaluating the reasons offered for the proponent's position. The reason *ad hominem* is a fallacy is that the attack on an individual is simply not relevant to the quality of the reasoning offered by that person. Attacking the person who offers an argument has nothing to do whether or not the premises of the argument are true or whether they support the conclusion. *Ad hominem* amounts to a way of changing the subject from whatever was at issue to potential flaws in the character or behavior of the person who was trying to reason about that issue. Part of what makes *Ad hominem* so effective is that people are generally quick to defend their honor.

Ad hominem is a particularly rampant and destructive fallacy in our society. It quickly turns the cooperative social project of inquiry through conversation into polarized verbal combat. This fallacy makes reasonable dialogue impossible while it diverts attention from interesting issues that often could be fruitfully investigated.

Here's an example of *ad hominem*: A car salesman argues for the quality of an automobile and the potential buyer discounts the argument with the thought that the person is just trying to earn a commission. We can imagine a situation where the salesperson is just trying to earn a commission and yet he is also making good arguments. So, consider a salesman who is just concerned with make lots of money. However, this salesperson is not very good at lying and manipulating people and decides that the best way to earn good commissions is to research his product carefully and then to only accept a sales position with the dealer that sells the very best. He then sincerely delivers good arguments for the quality of his product, makes lots of money, and dresses well. The customer who rejects his reasons for buying the car he sells on the *ad hominen* grounds that he is just trying to earn a commission misses an opportunity to buy the best. The moral of the story is just that the salesperson's motive is logically independent of the quality of his argument. The quality of an argument or an idea doesn't depend on who is offering it or what their motivation is.

- **False Dichotomy:** A dichotomy is an either/or choice where this is no third or fourth option. We've seen an example of a dichotomy in the contrast between the claim that there is intelligent life on other planets and the claim that there is no intelligent life on other planets. If one option is false then the other is true. There is no third or fourth possibility. On the other hand, when you go to a restaurant and you are trying to decide between the Impossible Burger or the Caesar Salad, you are probably not facing a dichotomy. You also have the option of having the salmon, or perhaps the fajita. The fallacy of false dichotomy is committed when we are presented with just two options as if these were the only possibilities when in fact there may be a third, fourth or more other possibilities.

So, here is a famous example of the false dichotomy fallacy. Shortly after 911, while building his case for invading Iraq, George W. Bush proclaimed, "You are either with us or you are with the terrorists." Some American's protested the invasion of Iraq, arguing that we did not have good reason to feel threatened by Iraq and that an unjust war would inspire more terrorism than it prevented. People who protested the Iraq war were neither with the terrorists, nor with the Bush administration. They shared the administration's goal of fighting terrorism, but doubted that invading Iraq was going to be an effective way of pursuing that goal. False dichotomy is a common strategy for dividing people into opposing camps by distracting attention from the middle ground where more productive conversation might be found.

- **Straw Man:** When soldiers fought with bayonets on their rifles, they would train by attacking straw men. Straw men are fairly easy to stab with a bayonet since they don't run

away or fight back. But then stabbing a straw man is no victory over an actual opponent. The fallacy of straw man is committed when someone criticizes an easy to attack distortion of an argument or idea rather than the actual view. Like many fallacies, this one can be committed deliberately or inadvertently. In our highly polarized social media environment, it is not uncommon for a disingenuous manipulator to deliberately broadcast a straw man attack (or some other fallacy) all the while knowing that his audience, lacking well-developed critical thinking skills, will fall for the manipulation and go on to propagate the bad argument unwittingly. This is often how propaganda works.

You may have heard a commonly propagated straw man attack committed against efforts to address climate change. Critics will often charge that people concerned about climate change are really just socialists looking to take our freedom away. There's a lot going on here and its worth pointing out the fallacies are gregarious. It is quite possible to commit more than one at a time. So, you might also notice an element of *ad hominem* in this example where reasons for taking climate seriously get ignored in favor of attacking the people trying to take climate change seriously. This sheds some light on the old quip that lies travel half way around the world before the truth gets its shoes on. It takes lots more work to diagnose and filter out fallacies than it does to commit and propagate them.

But aside from the *ad hominem* attack, the notion that people who want to see action on climate change are just big government lovers includes a straw man fallacy. It is easy and appealing to attack the socialist idea of government taking over the economy. It is not so easy to attack the idea that we have a serious problem in climate change and effective government action will be required to address it. Climate advocates are not arguing for socialism, a complete government take-over of the economy. They are arguing for government and business to work together to move us rapidly towards a sustainable economy, one that is based on renewable energy instead of fossil fuels, sustainable agriculture instead of deforestation, etc. Who owns and operates the industries of the future is simply not what is at issue, though many climate activists will be quick to point out the entrepreneurial opportunities in shifting to a more sustainable economy. In light of the existential risk we face in climate change, the policy measures called for are much harder to argue against than the straw man of widely despised socialism.

- **Hasty generalization:** The human brain has evolved to recognize patterns and project from these to unobserved instances. We instinctively expect things to continue to happen in accordance with the patterns we have observed. When we generalize from genuinely reliable patterns, our inferences can be inductively strong. But assuring the strength of our inductive generalizations requires that we generalize from ample evidence that is actually reflective of larger patterns in the world. In everyday life, we are highly prone to short circuiting this process and drawing generalizations too quickly from too little evidence, or evidence that is biased or distorted in some manner. When we do so, we generalize hastily and commit this fallacy.

Our fears and anxieties are often complicit in our hasty generalizations. When we hear a rustling in the bushes that sounds like it could be a bear or a mountain lion, the price of not jumping to this conclusion and being wrong (failing to infer that there is a mountain lion when there is one) is much higher than the price of making the inference and misfiring (inferring that there is a mountain lion when there is none). Evolution favors hastily inductive inference, much more so than generalizing methodically and scientifically. Where our fears are rational, this is all well and good. But fear is often not rational, and worse, our fears are easily manipulated. Hasty generalization on the basis of irrational or manipulated fear is the foundation of some of the worst injustices people perpetrate. Racial prejudice is a prime example.

The German Historical Museum in Berlin curates a vast collection of antisemitic propaganda tracing German history leading up to the Holocaust. An examination of this history quickly reveals that prejudice is often founded on hasty generalizations. Further, these hasty generalizations are largely built on manufactured evidence. The propaganda that stoked antisemitism was not typically based on fact. Fear is a powerful motivator both when it is credible and when it is not. Our own society's treatment of Black Americans provides many further troubling examples of the racial injustice based on hasty generalizations from biased or even fabricated representations. I'll discuss one example in the context of another inductive fallacy.

- **Spurious Correlation:** When we find a significant correlation between one condition and another, it is tempting to assume this indicates that one condition causes the other. Indeed, often it does. The high correlation between flipping the light switch and the room lighting up is explained by the former action causing the later condition. But a correlation between one condition and another doesn't always work this way. It can also, for instance, be that both conditions have a common cause. For example, night routinely follows day, but day does not cause night. The correlation we find in night following day is caused by the rotation of the planet as it orbits the sun.

Official crime rates among Black Americans are higher than they are among white Americans. The statistics here need to be understood in the context of an assortment racial biases in the criminal justice system. This is not just a matter of individual police officers being racially biased, though some are. There are also a number of systemic factors involved. Poor neighborhoods are more heavily policed and these tend to be more racially diverse. The crack cocaine epidemic that plagued Black communities was aggressively prosecuted and sent many Black people to prison. The current wave of opioid addiction that more often afflicts white communities is treated with more a more compassionate approach. So, there is a good

deal of institutional racial bias built into the official crime statistics. But even if we bracket these injustices, the correlation between crime and race in official statistics is spurious.

The gap in official crime rates between Black Americans and white Americans leads a significant number of people to the conclusion that Black people are just more criminally prone, as if race alone could explain this. But the racist conclusion is not well supported. The gap in crime rates closely mirrors the gap in unemployment rates between Black Americans and whites. Both official crime rates and official unemployment rates are higher among Blacks by similar factors. This suggests a causal explanation for higher crime rates among Black Americans that makes a good deal of sense and doesn't attribute innate criminality to Black Americans as racists would have us believe. Typically, people turn to crime only when they are deprived of decent opportunities in life. Regardless of race, people with good jobs and some prospects for a decent life have lots to lose and won't be very tempted to risk it all on criminal activity.

Correlations call for explanations. There is an inference to the best explanation involved in this. So, let's recall how the surprise principle discussed in the last chapter works. The explanation that makes the correlation we want to explain least surprising is the one that explains best and is thereby inductively confirmed. The idea that race somehow explains criminality is rather mysterious. There have been many racially motivated attempts to substantiate this idea and none have panned out. But higher crime rates among people who have been denied opportunities in life is not at all surprising. Inference to the best explanation strongly favors the idea that unemployment is a significant causal factor in crime over the idea that race somehow explains crime. So, inference to the best explanation indicates that the correlation between crime and race is spurious, not causal.

There are many more fallacies worth getting familiar with. I'll leave you to explore these on your own. The [The Fallacy Files](#) is a good place to start. I'll wrap up here with a brief discussion of confirmation bias.

Confirmation bias is the intellectual bad habit of endorsing just the evidence and argument that seems to support the view you already hold. This isn't exactly a fallacy because it isn't a specific kind of mistake in reasoning. We might think of confirmation bias as a meta-fallacy. It's the bad habit of trafficking in fallacious arguments for conclusions we like. Any fallacy can be involved in confirmation bias.

Confirmation bias is about what we should expect to find among people who lack strong critical thinking skills. People who don't know how to evaluate arguments have little else to go on except to prefer arguments that seem to confirm opinions they hold.

We all have good reason to avoid confirmation bias because it tends to undermine our credibility. Even if your view is well supported by good reasons, your presentation of it will be far less persuasive when you throw in a few shoddy arguments to boot. Your audience is likely to feel manipulated and to lose faith in your intellectual integrity. The only way to avoid confirmation bias is through cultivating your critical thinking skills; we do this by learning how to evaluate arguments and how to identify fallacies.

Review and Discussion Questions

1. How does dialectic differ from debate?
2. What is it for a claim to be true? How does this issue differ from that of determining whether a claim is true?
3. Explain our everyday concept of truth in terms of correspondence.
4. What is an argument? How do we evaluate arguments?
5. What does it mean for the premises of an argument to support its conclusion, and what are the two standards of support?
6. Explain the surprise principle and illustrate its use in evaluating an inference to the best explanation.
7. What is a fallacy? Learn about a fallacy or two on [The Fallacy Files](#) and report back.

Vocabulary

- Truth
- Dialectic
- Argument
- Valid
- Sound
- Strong
- Cogent
- Inference to the best explanation
- Fallacy

Exercises

Which of the following arguments are valid? Which are invalid?

A

1. James will get an A in philosophy if and only if he writes a good paper.
2. James got an A in philosophy

3. Therefore, he wrote a good paper.

B

1. If Ben writes a good paper, he will get an A in philosophy.
2. Ben got an A in philosophy
3. Therefore, he wrote a good paper.

C

1. If whales are mammals, then they are not fish.
2. Whales are fish
3. Whales are not mammals.

D

1. If the rapture has occurred, then either some of the cars on the highway will be unoccupied or all drivers are damned.
2. Some drivers are not damned.
3. None of the cars on the highway are unoccupied.
4. Therefore, the rapture has not occurred.

E

1. Some snarks are bandersnatches.
2. All bandersnatches are igglypoofs.
3. So, some snarks are igglypoofs

Answer the following questions. Give short explanations that reason from the definitions of the relevant logical concepts.

1. Does an argument provide a good reason for believing its conclusion if it is valid? Explain.
2. Can a valid argument have a false conclusion? Explain.
3. Can a sound argument have a false conclusion? Explain.
4. What is it for a statement to be valid? (trick question)

Which of the following arguments are inductively strong? Which are weak?

1. It has rained every day in the Darién Gap for the past 25 years. Thus, it will probably rain in the Darién Gap tomorrow.

2. People try on shoes before buying them. People drive cars before signing up for a three-year lease. People take a close look at travel information before committing to an expensive vacation. So, people should have sex with each other before committing to marriage.
3. Two teenagers were found writing graffiti on the school walls yesterday. Thus, all teenagers are delinquents.
4. A reliable study showed that 90 percent of Bellevue College's students want more training in critical thinking. Maria is a student at Bellevue College. So, Maria probably wants more training in critical thinking.
5. Upon landing at the SeaTac Airport, plane passengers saw broken buildings, large cracks in the runway, fire engines running about, and paramedics assisting injured people. The passengers concluded that an earthquake just occurred.

Answer the following questions. Give short explanations based on the definitions of the relevant concepts.

1. Explain how deductive validity and inductive strength differ.
2. Can the conclusion of an inductively cogent argument be false? Explain.
3. Must an inductively strong argument have true premises? Explain.

5. Ancient Philosophy

Our main focus in this chapter will be with the three major ancient Greek philosophers, Socrates, Plato, and Aristotle. It is in these thinkers that science and philosophy get started. But even these great minds were working in an intellectual tradition and it will be well worth a few minutes to appreciate the historical context in which they worked and the foundation it provided. So we will begin with a quick tour of Pre-Socratic thought.

The Presocratics

In the *Iliad* and the *Odyssey*, the early Ionian epic poet Homer offers a view of the world as under the influence of the Olympian gods. The Olympian gods were much like humans, capricious and willful. In the Homeric view of the world, human qualities are projected onto the world via human-like gods. Here explanation of the natural world is modeled on explanation of human behavior. This marks the world view of the epic poets as pre-philosophical and pre-scientific. However, even in the early epic poems we find a moral outlook that is key to the scientific and philosophical frame of mind. In Homer and in later Greek tragedy, we find stories of the grief that human hubris brings upon us. The repeated warnings against human pride and arrogance make a virtue out of humility. Intellectual humility involves recognizing the fallibility of human thought, in particular one's own. The willingness to submit one's own opinions to rational scrutiny is essential to moving beyond the realm of myth and into the realm of philosophy and science. Intellectual humility makes it possible to see the world and one's place in it as a matter for discovery rather than a matter of self-assertion.

The Milesians

The beginning of philosophy in ancient Greece is often given as 585 B.C., the year that the Milesian philosopher Thales predicted a solar eclipse. Thales brings a new naturalistic approach to explaining the world. That is, his proposed explanations for natural phenomenon are given in terms of more fundamental natural phenomenon, not in supernatural terms. The step away from supernatural myth and towards understanding the natural world on its own terms is a major development. Thales is interested in the fundamental nature of the world and arrives at the view that the basic substance of the world is water. His reason for thinking that water is fundamental is that of the four recognized elements - earth, air, fire and water - only water can take the form of a solid, liquid, or a gas. According to Thales, earth is really water that is even more concentrated than ice and fire is really water that is more rarified than steam. While his view sounds absurd to us, the significance of his contribution is not the specific answer he gives to the question of the ultimate nature of the world, but how he proposes to answer this question. Thales takes an important step away from projecting ourselves onto the world through myth and superstition and towards explanations that invite further investigation of the world as it is independent of human will.

Pythagoras (fl. 525-500 B.C.) traveled in Egypt where he learned astronomy and geometry. His thought represents a peculiar amalgam of hardnosed mathematical thinking and creative but rather kooky superstition. Pythagoras holds that all things consist of numbers. He saw mathematics as a purifier of the soul. Thinking about numbers takes one's attention off of particular things and elevates the mind to the realm of the eternal. Scientific thinking, on this view, is not so far from meditation. Pythagoras is responsible for the Pythagorean Theorem which tells us that the square of the hypotenuse of a right triangle is equal to the sum of the squares of the remaining sides. He also discerned how points in space can define shapes, magnitudes, and forms:

- 1 point defines location
- 2 points define a line
- 3 points define a plane
- 4 points define solid 3 dimensional objects

Pythagoras introduces the concept of form. The earlier Milesians only addressed the nature of matter, the stuff of the universe. A full account of the nature of the world must also address the various forms that underlying stuff takes. Form implies limits. For Pythagoras, this is understandable in numerical terms. Number represents the application of limit (form) to the unlimited (matter). The notion of form takes on greater sophistication and importance in the thought of Plato and Aristotle.

Pythagoras led a cult that held some rather peculiar religious beliefs. The more popular beliefs in the Homeric gods are not concerned with salvation or spiritual purification. There was the Dionysian religion, which sought spiritual purification and immortality through drunken carnal feasts and orgies. Pythagorean religious belief also aims at purification and immortality, but without the intoxication and sex. Pythagoras founded a religious society based on the following precepts:

- that at its deepest level, reality is mathematical in nature
- that philosophy can be used for spiritual purification
- that the soul can rise to union with the divine
- that certain symbols have a mystical significance
- that all brothers of the order should observe strict loyalty and secrecy

Members of the inner circle were strict communist vegetarians. They were also not allowed to eat beans. Pythagoras might have done well in Ballard.

The last of the Milesians we will discuss is Heraclitus. Heraclitus (544-484 B.C.) was born in Ephesus on the coast of Asia Minor. He is best known for his doctrine of eternal flux according to which everything undergoes perpetual change. "One can never step in the same river twice." The underlying substance of the world is fire or heat according to Heraclitus. This is the least

stable of the elements and explains the transitoriness of all things. Everything is a kindling or extinguishing of fire. While everything is in a continual state of flux, this change is not without order. Heraclitus saw Logos or rational order as essential to the world. Changes are injustices, which by natural necessity are redressed in further changes. Heraclitus held ethical views worth noting as well. The good life involves understanding and accepting the necessity of strife and change.

The Sophists

Most of early Greek philosophy prior to the Sophists was concerned with the natural world. The desire to explain an underlying reality required natural philosophers to speculate beyond what is observable and they lacked any developed critical method for adjudicating between rival theories of substance change or being. In this situation, it is easy to see how many might grow impatient with natural philosophy and adopt the skeptical view that reason simply cannot reveal truths beyond our immediate experience. But reason might still have practical value in that it allows the skilled arguer to advance his interests. The Sophists were the first professional educators. For a fee, they taught students how to argue for the practical purpose of persuading others and winning their way. While they were well acquainted with and taught the theories of philosophers, they were less concerned with inquiry and discovery than with persuasion.

Pythagoras and Heraclitus had offered some views on religion and the good life. Social and moral issues come to occupy the center of attention for the Sophists. Their tendency towards skepticism about the capacity of reason to reveal truth and their cosmopolitan circumstances, which exposed them to a broad range of social customs and codes, lead the Sophists to take a relativist stance on ethical matters. The Sophist's lack of interest in knowing the truth for its own sake and their entrepreneurial interest in teaching argument for the sake of best serving their client's interests leads Plato to derisively label the Sophists as "shopkeepers with spiritual wares."

One of the better known Sophists, Protagoras (481-411 B.C.), authored several books including, *Truth, or the Rejection* (the rejection of science and philosophy), which begins with his best-known quote, "man is the measure of all things, of those that are that they are, of those that are not that they are not." Knowledge, for Protagoras is reducible to perception. Since different individuals perceive the same things in different ways, knowledge is relative to the knower. This is a classic expression of epistemic relativism. Accordingly, Protagoras rejects any objectively knowable morality and takes ethics and law to be conventional inventions of civilizations, binding only within societies and holding only relative to societies.

Socrates

Socrates is widely regarded as the founder of philosophy and rational inquiry. He was born around 470 B.C., and tried and executed in 399 B.C.. Socrates was the first of the three major Greek philosophers; the others being Socrates' student Plato and Plato's student Aristotle.

Socrates did not write anything himself. We know of his views primarily through Plato's dialogues where Socrates is the primary character. Socrates is also known through plays of Aristophanes and the historical writings of Xenophon. In many of Plato's dialogues it is difficult to determine when Socrates' views are being represented and when the character of Socrates is used as a mouthpiece for Plato's views.

Socrates was well known in Athens. He was eccentric, poor, ugly, brave, stoic, and temperate. He was a distinguished veteran who fought bravely on Athens' behalf and was apparently indifferent to the discomforts of war. Socrates claimed to hear a divine inner voice he called his *daimon* and he was prone to go into catatonic states of concentration.

The conflicting views of the Ionian and Eleatic philosophers of nature encouraged skepticism about our ability to obtain knowledge through rational inquiry. Among the Sophists, this skepticism is manifested in epistemic and Moral Relativism. Epistemic relativism is the view that there is no objective standard for evaluating the truth or likely truth of our beliefs. Rather, epistemic standards of reasoning are relative to one's point of view and interests. Roughly, this is the view that what is true for me might not be true for you (when we are not just talking about ourselves). Epistemic relativism marks no distinction between knowledge, belief, or opinion on the one hand, and truth and reality on the other. To take a rather silly example, if I think it's Tuesday, then that's what's true for me; and if you think it's Thursday, then that's what is true for you. In cases like this, epistemic relativism seems quite absurd, yet many of us have grown comfortable with the notion that, say, beliefs about the moral acceptability of capital punishment might be true for some people and not for others.

Moral Relativism is the parallel doctrine about moral standards. The moral relativist takes there to be no objective grounds for judging some ethical opinions to be correct and others not. Rather, ethical judgments can only be made relative to one or another system of moral beliefs and no system can be evaluated as objectively better than another. Since earlier attempts at rational inquiry had produced conflicting results, the Sophists held that no opinion could be said to constitute knowledge. According to the Sophists, rather than providing grounds for thinking some beliefs are true and others false, rational argument can only be fruitfully employed as rhetoric, the art of persuasion. For the epistemic relativist, the value of reason lies not in revealing the truth, but in advancing one's interests. The epistemic and Moral Relativism of the Sophist has become popular again in recent years and has an academic following in much "post-modern" writing.

Socrates was not an epistemic or moral relativist. He pursued rational inquiry as a means of discovering the truth about ethical matters. But he did not advance any ethical doctrines or lay claim to any knowledge about ethical matters. Instead, his criticism of the Sophists and his contribution to philosophy and science came in the form of his method of inquiry.

As the Socratic Method is portrayed in Plato's Socratic dialogues, interlocutor proposes a definition or analysis of some important concept, Socrates raises an objection or offers counter examples, then the interlocutor reformulates his position to handle the objection. Socrates raises a more refined objection. Further reformulations are offered, and so forth. Socrates uses the dialectic to discredit others' claims to knowledge. While revealing the ignorance of his interlocutors, Socrates also shows how to make progress towards more adequate understanding.

A good example of the Socratic Method at work can be found in one of Plato's early Socratic dialogues, *Euthyphro*. Here is a link: <http://www.gutenberg.org/ebooks/1642>.

In Plato's dialogues we often find Socrates asking about the nature of something and then critically examine proposed answers, finding assorted illuminating objections that often suggest next steps. In this dialogue, Socrates and Euthyphro are discussing the nature of piety or holiness. Socrates and Euthyphro never conclusively discover what piety is, but they learn much about how various attempts to define piety fail. The dialogue works the same if we substitute moral goodness for piety. Understood in this way, *Euthyphro* provides a classic argument against Divine Command Theory, a view about the nature of morality that says that what is right is right simply because it is commanded by God.

Socrates would not have us believe our questions have no correct answers. He is genuinely seeking the truth of the matter. But he would impress on us that inquiry is hard and that untested claims to knowledge amount to little more than vanity. Even though Euthyphro and Socrates don't achieve full knowledge of the nature of piety, their understanding is advanced through testing the answers that Euthyphro suggests. We come to see why piety can't be understood just by identifying examples of it. While examples of pious acts fail to give us a general understanding of piety, the fact that we can identify examples of what is pious suggests that we have some grasp of the notion even in the absence of a clear understanding of it.

After a few failed attempts to define piety, Euthyphro suggests that what is pious is what is loved by the gods (all of them, the Greeks recognized quite a few). Many religious believers continue to hold some version of Divine Command Theory. In his response to Euthyphro, Socrates points us towards a rather devastating critique of this view and any view that grounds morality in authority. Socrates asks whether what is pious is pious because the gods love it or whether the gods love what is pious because it is pious. Let's suppose that the gods agree in loving just what is pious. The question remains whether their loving the pious explains its piety or whether some things being pious explains why the gods love them. Once this question of what is supposed to explain what is made clear, Euthyphro agrees with Socrates that the gods love what is pious because it is pious. The problem with the alternative view, that what is pious is pious because it is loved by the gods, is that this view makes piety wholly arbitrary. Anything could be pious if piety is just a matter of being loved by the gods. If the gods love puppy torture, then this would be pious. Hopefully this seems absurd. Neither Socrates nor Euthyphro is willing to accept that what is pious is completely arbitrary. At this point, Socrates points out to Euthyphro that since an

act's being pious is what explains why the gods love it, he has failed to give an account of what piety is. The explanation can't run in both directions. In taking piety to explain being loved by the gods, we are left lacking an explanation of what piety itself is. Euthyphro gives up shortly after this failed attempt and walks off in a huff.

If we substitute talk of God making things right or wrong by way of commanding them for talk of the gods loving what is pious in this exchange of ideas, we can readily see that Divine Command Theory has the rather unsavory result that torturing innocent puppies would be right if God commanded it. We will return to this problem when we take up ethical theory later in the course. While we don't reach the end of inquiry into piety (or goodness) in *Euthyphro*, we do make discernible progress in coming to see why a few faulty accounts must be set aside. Socrates does not refute the skeptic or the relativist Sophist by claiming to discover the truth about anything. What he does instead is show us how to engage in rational inquiry and show us how we can make progress by taking the possibility of rational inquiry seriously.

Apology

This dialogue by Plato is a dramatization of Socrates' defense at his trial for corrupting the youth among other things. Socrates tells the story of his friend Chaerophon who visits the Oracle of Delphi and asks if anyone in Athens is wiser than Socrates. The Oracle answered that no one is wiser than Socrates. Socrates is astounded by this and makes it his mission in life to test and understand the Oracle's pronouncement. He seeks out people who have a reputation for wisdom in various regards and tests their claims to knowledge through questioning. He discovers a good deal of vain ignorance and false claims to knowledge, but no one with genuine wisdom. Ultimately, Socrates concludes that he is wisest, but not because he possesses special knowledge not had by others. Rather, he finds that he is wisest because he recognizes his own lack of knowledge while others think they know, but do not.

Of course people generally, and alleged experts especially, are quite happy to think that what they believe is right. We tend to be content with our opinions and we rather like it when others affirm this contentment by agreeing with us, deferring to our claims to know or at least by "respecting our opinion" (whatever that is supposed to mean). We are vain about our opinions even to the point of self identifying with them (I'm the guy who is right about this or that). Not claiming to know, Socrates demonstrates some intellectual humility in allowing that his opinions might be wrong and being willing to subject them to examination. But in critically examining various opinions, including those of the supposed experts, he pierces the vanity of many of Athens' prestigious citizens. Engaging in rational inquiry is dangerous business, and Socrates is eventually brought up on charges of corrupting the youth who liked to follow him around and listen to him reveal people's claims to knowledge as false pride. The *Apology* documents Socrates' defense of his of behavior and the Athenian assembly's decision to sentence him to death anyway.

You will find the *Apology* in several formats here: <http://www.gutenberg.org/ebooks/1656>

Study questions for the Apology:

1. What are the ancient and the more recent charges brought against Socrates?
2. How does he answer the ancient charges?
3. According to Socrates, why would he not intentionally corrupt the youth?
4. Suppose Socrates unintentionally corrupted the youth. Should he be punished anyway for the negative impact of his actions? Explain your answer.
5. Explain the mission Socrates sets himself on in response to the the pronouncement by the Oracle at Delphi.
6. Socrates argues in a couple places that the worse man can not harm the better man. How does that argument go?
7. What does Socrates’ defense reveal about the values he lives by? What matters most to Socrates in life?
8. How does Socrates argue that the fear of death is irrational?
9. How could Socrates have avoided the death penalty?
10. Was his choice not to evade death an honorable one?
11. How did Socrates see his critical questioning of Athenians as beneficial to his city and its citizens?
12. Do you think Socrates was too hard on his fellow Athenians before his accusers came forward? Was he too hard on them during the trial and after the verdict?

Plato

Plato (429-347 B.C.) came from a family of high status in ancient Athens. He was a friend and fan of Socrates and some of his early dialogues chronicle events in Socrates’ life. Socrates is a character in all of Plato’s dialogues. But in many, the figure of Socrates is employed as a voice for Plato’s own views. Unlike Socrates, Plato offers very developed and carefully reasoned views about a great many things. Here we will briefly introduce his core metaphysical, epistemological and ethical views.

Metaphysics and Epistemology

Plato’s metaphysics and epistemology are best summarized by his device of the divided line. The vertical line between the columns below distinguishes reality and knowledge. It is divided into levels that identify what in reality corresponds with specific modes of thought.

Objects	Modes of Thought
The Forms	Knowledge
Mathematical objects	Thinking
Particular things	Belief /Opinion

Images	Imaging
--------	---------

Here we have a hierarchy of Modes of Thought, or types of mental representational states, with the highest being knowledge of the forms and the lowest being imaging (in the literal sense of forming images in the mind). Corresponding to these degrees of knowledge we have degrees of reality. The less real includes the physical world, and even less real, our representations of it in art. The more real we encounter as we inquire into the universal natures of the various kinds of things and processes we encounter. According to Plato, the only objects of knowledge are the forms which are abstract entities.

In saying that the forms are abstract, we are saying that while they do exist, they do not exist in space and time. They are ideals in the sense that a form, say the form of horse-ness, is the template or paradigm of being a horse. All the physical horses partake of the form of horse-ness, but exemplify it only to partial and varying degrees of perfection. No actual triangular object is perfectly triangular, for instance. But all actual triangles have something in common, triangularity. The form of triangularity is free from all of the imperfections of the various actual instances of being triangular. We get the idea of something being more or less perfectly triangular. For various triangles to come closer to perfection than others suggests that there is some ideal standard of “perfectly triangularity.” This for Plato, is the form of triangularity. Plato also takes moral standards like justice and aesthetic standards like beauty to admit of such degrees of perfection. Beautiful physical things all partake of the form of beauty to some degree or another. But all are imperfect in varying degrees and ways. The form of beauty, however, lacks the imperfections of its space and time bound instances. Perfect beauty is not something we can picture or imagine. But an ideal form of beauty is required to account for how beautiful things are similar and to make sense of how things can be beautiful to some less than perfect degree or another.

Only opinion can be had regarding the physical things, events, and states of affairs we are acquainted with through our sensory experience. With physical things constantly changing, the degree to which we can grasp how things are at any given place and time is of little consequent. Knowledge of the nature of the forms is a grasp of the universal essential natures of things. It is the intellectual perception of what various things, like horses or people, have in common that makes them things of a kind. Plato accepts Socrates’ view that to know the good is to do the good. So his notion of epistemic excellence in seeking knowledge of the forms will be a central component of his conception of moral virtue.

Ethics

Plato offers us a tripartite account of the soul. The soul consists of a rational thinking element, a motivating willful element, and a desire-generating appetitive element. Plato offers a story of the rational element of the soul falling from a state of grace (knowledge of the forms) and dragged

down into a human state by the unruly appetites. This story of the soul's relation to the imperfect body supports Plato's view that the knowledge of the forms is a kind of remembrance. This provides a convenient source of knowledge as an alternative to the merely empirical and imperfect support of our sense experience. Plato draws an analogy between his conception of the soul and a chariot drawn by two horses, one obedient, the other rebellious. The charioteer in this picture represents the rational element of the soul, the good horse the obedient will, and the bad horse, of course, represents those nasty earthly appetites. To each of the elements of the soul, there corresponds a virtue; for the rational element there is wisdom, for the willing element of the soul there is courage, and for the appetitive element there is temperance. Temperance is matter of having your appetites under control. This might sound like chronic self-denial and repression, but properly understood, it is not. Temperance and courage are cultivated through habit. In guiding our appetites by cultivating good habits, Plato holds, we can come to desire what is really good for us (you know, good diet, exercise, less cable TV, and lots more philosophy - that kind of stuff).

Wisdom is acquired through teaching, via the dialectic, or through "remembrance." Perhaps, to make the epistemological point a little less metaphysically loaded, we can think of remembrance as insight. A more general virtue of justice is conceived as each thing functioning as it should. To get Plato's concept of justice as it applies to a person, think of the charioteer managing and controlling his team; keeping both horses running in the intended direction and at the intended speed. Justice involves the rational element being wise and in charge. For a person to be just is simply a matter of having the other virtues and having them functioning together harmoniously.

Given Plato's ethical view of virtue as a matter of the three elements of the soul functioning together as they should, Plato's political philosophy is given in his view of the state as the human "writ at large." Project the standards Plato offers for virtue in an individual human onto the aggregate of individuals in a society and you have Plato's vision of the virtuous state. In the virtuous state, the rational element (the philosophers) are in charge. The willing element (the guardians or the military class) is obedient and courageous in carrying out the policies of the rational leadership. And the appetitive element (the profit-driven business class) functions within the rules and constraints devised by the rational element (for instance, by honestly adhering to standards of accounting). A temperate business class has the profit motive guided by the interests of the community via regulation devised by the most rational. The virtuous business class refrains from making its comfort and indulgence the over-riding concern of the state. Plato, in other words, would be no fan of totally free markets, but neither would he do away with the market economy altogether.

Plato's vision of social justice is non-egalitarian and anti-democratic. While his view would not be popular today, it is still worthwhile to consider his criticism of democracy and rule by the people. Plato has Socrates address this dialectically by asking a series of questions about who we would want to take on various jobs. Suppose we had grain and wanted it processed into flour. We would not go to the cobbler or the horse trainer for this, we'd go to the miller. Suppose we

had a horse in need of training. We obviously would not go to the miller or the baker for this important task, we'd go to the horse trainer. In general, we want important functions to be carried out by the people with the expertise or wisdom to do them well. Now suppose we had a state to run. Obviously we would not want to turn this important task over to the miller, the cobbler, or the horse trainer. We'd want someone who knows what he or she is doing in charge. Plato has a healthy regard for expertise. As Plato sees it, democracy amounts to turning over the ethically most important jobs to the people who have the least expertise and wisdom in this area. There is very little reason to expect that a state run by cobblers, millers, and horse trainers will be a virtuous state.

Aristotle

Aristotle is a towering figure in the history of philosophy and science. Aristotle made substantive contributions to just about every philosophical and scientific issue known in the ancient Greek world. Aristotle was the first to develop a formal system of logic. As the son of a physician he pursued a life-long interest in biology. His physics was the standard view through Europe's Middle Ages. He was a student of Plato, but he rejected Plato's other-worldly theory of forms in favor of the view that things are a composite of substance and form. Contemporary discussions of the good life still routinely take Aristotle's ethics as their starting point. Here I will offer the briefest sketch of Aristotle's logic and metaphysics. We will return to his ethics later in the course.

Logic

Aristotle's system of logic was not only the first developed in the West, it was considered complete and authoritative for well over 2000 years. The core of Aristotle's logic is the systematic treatment of categorical syllogisms. You might recall this argument from Chapter 2:

1. All monkeys are primates.
2. All primates are mammals.
3. So, all monkeys are mammals.

This argument is a categorical syllogism. That's a rather antiquated way of saying it's a two premise argument that uses simple categorical claims. Simple categorical claims come in one of the following four forms:

- All A are B
- All A are not B
- Some A are B
- Some A are not B

There are a limited number of two premise argument forms that can be generated from combinations of claims having one of these four forms. Aristotle systematically identified all of

them, offered proofs of the valid one's, and demonstrations of the invalidity of the others. Beyond this, Aristotle proves a number of interesting things about his system of syllogistic logic and he offers an analysis of syllogisms involving claims about what is necessarily the case as well.

No less an authority than Immanuel Kant, one of the most brilliant philosophers of the 18th century, pronounced Aristotle's logic complete and final. It is only within the past century or so that logic has developed substantially beyond Aristotle's. While Aristotle's achievement in logic was genuinely remarkable, this only underscores the dramatic progress of the 20th century. The system of symbolic logic we now teach in standard introductions to logic (PHIL 120 here at BC) is vastly more powerful than Aristotle's and while this system was brand new just a century ago, it is now truly an introduction, a first step towards appreciating a great many further developments in logic. Reasonably bright college students now have the opportunity to master deductive reasoning at a level of sophistication unknown to the world a mere 150 years ago. The methods and insights of modern symbolic logic are already so thoroughly integrated into contemporary philosophy that much of contemporary philosophy would not be possible without it.

Metaphysics

While Aristotle was a student of Plato's, his metaphysics is decidedly anti-Platonist. The material of the world takes various forms. Here it constitutes a tree and there a rock. The things constituted of matter have various properties. The tree is a certain shape and height, the rock has a certain mass. Plato accounts for the various forms matter takes and the ways things are in terms of their participating in abstract and ideal forms to one degree or another. Plato's metaphysics centrally features an abstract realm of eternal unchanging and ideal forms. Plato's forms are not themselves part of the physical spatio-temporal world. Aristotle rejects the theory of abstract forms and takes everything that exists to be part of the physical spatio-temporal world. It might thus be tempting to think of Aristotle as a materialist, one who thinks all that exists is matter, just atoms swirling in the void. Some pre-Socratic philosophers could accurately be described as materialists. But this would miss key elements of Aristotle's metaphysics. While Aristotle denies the existence of an abstract realm of eternal and unchanging ideal entities, his account of the nature of things includes more than just matter. Aristotle holds the view that form is an integral part of things in the physical world. A thing like a rock or a tree is a composite of both matter and form. In addition to matter, the way matter *is* gets included in Aristotle's metaphysics.

Among the ways things are, some seem to be more central to their being what they are than others. For instance, a tree can be pruned into a different shape without the tree being destroyed. The tree can survive the loss of its shape. But if it ceased to be a plant, if it got chipped and mulched, for instance, it would also cease to be a tree. That is to say, being a plant is *essential* to the tree, but having a certain shape isn't. An essential property is just a property a thing could not survive losing. By contrast, a property something could survive losing is had *accidentally*. Aristotle introduces the distinction between essential and accidental characteristics of things.

This was an important innovation. When we set out to give an account of what a thing is, we are after an account of its essence. To say what a thing is essentially is to list those ways of being it could not survive the loss of. My hair length is not essential to me, and neither is my weight, my having four limbs, or a given body/mass index. But my having a mind, perhaps, is essential to being me.

How a thing functions is a critical aspect of its nature in Aristotle's view. As an organism, I metabolize. As an organism with a mind, I think. These are both ways of functioning. For Aristotle, what makes something what it is, its essence, is generally to be understood in terms of how it functions. Aristotle's account of the essential nature of the human being, for instance, is that humans are rational animals. That is, we are the animals that function in rational ways. Functioning is, in a sense, purposeful. Aristotle would say that functioning is ends oriented. The Greek term for an end or a goal is *telos*. So Aristotle has a *teleological* view of the world. That is, he understands things as functioning towards ends or goals, and we can understand the essence of things in terms of these goal-oriented ways of functioning. We still understand people's actions as teleological or goal oriented. We explain why people do things in terms of their purposes and methods. Aristotle similarly understands natural processes generally as ends oriented. Even Aristotle's physics is fundamentally teleological. So water runs downhill because it is part of its essential nature to seek out the lower place.

Explanation: The Four Causes

What does it mean to explain something? If we'd like to have some idea what we are up to when we explain things, giving some account of explanation should seem like an important methodological and epistemological issue. In fact, the nature of explanation continues to be a central issue in the philosophy of science. Aristotle was the first to address explanation in a systematic way and his treatment of explanation structures and guides his philosophical and scientific inquiry generally. According to Aristotle, to explain something involves addressing *four causes*. Here we need to think of "causes" as aspects of explanation or "things *because of* which . . ." Only one of Aristotle's *four causes* resembles what we would now think of as the cause of something. Three of the four causes, or explanatory principles, are reflected in Aristotle's metaphysics and will be familiar from the discussion above. Part of explaining something involves identifying the material of which it is made. This is the *material cause*. Thales account of the nature of the world addressed its material cause. A further part of explaining something is to give an account of its form, its shape and structure. The chair I am sitting on is not just something made of wood, it is something made of wood that has a certain form. A complete explanation of what this chair is would include a description of its form. This is the *formal cause*. Pythagoras and Plato introduce the explanation of formal causes. The idea of a *final cause* refers to the function, end, or *telos* of a thing. What makes the chair I'm sitting on a chair is that it performs a certain function that serves the end or *telos* of providing a comfortable place to sit. Again, Aristotle sees final causes as pervasive in the natural world. So part of explaining what a tomato plant is, for instance, will involve giving an account of how it functions

and the goals towards which that functioning is aimed. Bear in mind Aristotle's interest in biology here. A complete biological account of an organism includes both its anatomy (its material and formal causes) and physiology (which involves functioning and final causes). The remaining cause (explanatory principle) is the one we can identify as a kind of cause in our normal sense of the word. The efficient cause of a thing is that which brings it into existence or gives form to its material. So, for instance, the activity of a carpenter is the efficient cause of my chair.

From Ancient to Modern Philosophy

The following two chapters focus on what is known as the modern classical period which runs roughly concurrent with the scientific revolution. About 2000 years elapse between the ancient Greek philosophy and the modern classical period. This chapter will end with a very brief sketch of some trends and developments over the course of those two millennium.

The rise and fall of Rome follows the golden age of ancient Greece. Greek philosophical traditions undergo assorted transformations during this period, but Rome is not known for making significant original contributions to either philosophy or science. Intellectual progress requires a degree of liberty not so available in the Roman Empire. Additionally, the intellectual talent and energy available in ancient Rome would have been pretty fully occupied with the demands of expanding and sustaining political power and order. Rome had more use for engineers than scientists, and more use for bureaucrats than philosophers. Christianity becomes the dominate religion in Rome after emperor Constantine converts in the 4th century A.D., Also in the 4th century, the great Christian philosopher [Augustine](#), under the influence of Plato, formulates much of what will become orthodox Catholic doctrine. After a rather dissolute and free-wheeling youth, Augustine studies Plato and find's much to make Christianity reasonable in it. With the rise of the Catholic Church, learning and inquiry are pursued largely exclusively in the service of religion for well over a millennium. Philosophy in this period is often described as the handmaiden of theology. The relationship between philosophy and theology is perhaps a bit more ambiguous, though. As we've just noted in the case of Augustine, much ancient Greek philosophy gets infused into Catholic orthodoxy. But at the same time, the new faith of Christianity spearheads an anti-intellectual movement in which libraries are destroyed and most ancient Greek thought is lost to the world forever.

Through the West's period of Catholic orthodoxy, most of what we know of Greek science and philosophy, most notably [Aristotle's](#) thought, survived in the Islamic world. What remains of the [complete works of Aristotle](#) covers subjects as far ranging as metaphysics, ethics, politics, rhetoric, physics, biology, and astronomy, and amounts to enough writing to fill 1500 pages in the fine print translation on my bookshelf. But even this consists largely of lecture notes and fragments. Most of his polished prose is lost forever.

The crusades were a series of conflicts between the Christian and Islamic world towards the end of the middle-ages. This conflict between Christianity and Islam was also an occasion for cultural exchange, and the Crusades led to the re-introduction of Aristotle and other ancient Greek scholarship to the west. Aristotle's philosophy and science was too carefully reasoned, systematic, and subtle to be dismissed as pointless pagan speculation. Instead, Christian thinkers in the west set out to understand Aristotle and interpret him a manner that would cohere with Catholic doctrine. [St. Thomas Aquinas](#) is the most famous philosopher to engage in this work of Christianizing Aristotle. He found ways to harness Aristotle's metaphysical arguments in the cause of advocating the existence of a Christian God.

Aristotle's views about the natural world quickly come to be received as the established truth in the Christian world. Aristotle's physics, for instance becomes the standard scientific view about the natural world in Europe. Aristotle also wrote about the methods of science, and he was much more empirical than his teacher Plato. Aristotle thought the way to learn about the natural world was to make careful observations and infer general principles from these. For instance, as an early biologist, Aristotle dissected hundreds of species of animals to learn about anatomy and physiology. The Scholastics who studied Aristotle obviously did not adopt the methods Aristotle recommended. But some other people did. Galileo, Leonardo da Vinci, and Copernicus were among the few brave souls to turn a critical eye to the natural world itself and, employing methods Aristotle would have approved of, began to challenge the views of Aristotle that the Scholastics had made a matter of doctrine. Thus begins the Scientific Revolution.

Where the Renaissance is the reawakening of the West to its ancient cultural and intellectual roots, the Scientific Revolution begins as a critical response to ancient thinking, and in large part that of Aristotle. This critical response was no quick refutation. Aristotle's physics might now strike us as quite naïve and simplistic, but that is only because every contemporary middle school student gets a thorough indoctrination in Newton's relatively recent way understanding of the physical world. The critical reaction to Aristotle that ignites the scientific revolution grew out of tradition of painstakingly close study of Aristotle. The scholastic interpreters of Aristotle were not just wrongheaded folks stuck on the ideas of the past. They were setting the stage for new discoveries that could not have happened without their work. Again, our best critics are the ones who understand us the best and the one's from whom we stand to learn the most. In the Scientific Revolution we see a beautiful example of Socratic dialectic operating at the level of traditions of scholarship.

Europe also experiences significant internal changes in the 16th century that pave the way for its intellectual reawakening. In response to assorted challenges to the authority of the Catholic Church and the decadence of 16th century Catholic churchmen, [Martin Luther](#) launches the Reformation. The primary tenet of the reformation was that faith concerns the individual's relation to God who is knowable directly through the Bible without the intermediary of the Catholic Church. The Reformation and the many splintering branches of Protestant Christianity that it spawns undermines the dogmatic adherence to a specific belief system and opens the way

for more free and open inquiry. The undermining of Catholic orthodoxy brought on by the reformation combined with the rediscovery of ancient culture in the Renaissance jointly give rise to the Scientific Revolution and, what we often refer to as the Modern Classical period in philosophy. The reawakening of science and philosophy are arguably one and the same revolution. Developments in philosophy and science during this period are mutually informed, mutually influencing, and intermingled. Individuals including Newton, Leibniz, and Descartes are significant contributors to both science and philosophy.

Review and Discussion Questions:

1. Explain Protagoras' epistemic relativism.
2. How does Socrates oppose epistemic relativism?
3. What is the Socratic Method?
4. How does Socrates respond to Euthyphro's suggestion that the pious is what is loved by all the gods? How does his response point us towards a critique of Divine Command Theory? What is the problem with the view that what is pious is pious because it is loved by the gods?
5. What are Plato's forms? Why does Plato take the forms to be the most real sorts of entities?
6. What is temperance and why is it a virtue in Plato's view?
7. How is Plato's vision of justice non-egalitarian and anti-democratic?
8. How do Plato and Aristotle's views on form differ?
9. What is the difference between essence and accident?
10. What does it mean to say that Aristotle held a teleological view of the world?
11. Explain Aristotle's four causes as principles of explanation.
12. What is the role of Aristotle's philosophy and science in leading to the scientific revolution?

5. The Rationalists

Is all of our knowledge based on the evidence of the senses, or is some of it justified by other means? This epistemological question about the foundations of knowledge is what separates Rationalism and Empiricism. According to **Rationalism** at least some knowledge can be had through reason alone. For rationalists, the paradigm example of knowledge acquired independent of sense experience is mathematics. Once we have the concepts required to understand mathematical propositions (like $2+2=4$), no experience is required to be justified in accepting their truth. They seem to be adequately known “through the light of reason.” **Empiricism**, on the other hand, takes all of our knowledge to be ultimately grounded in sense experience. Descartes was the first significant rationalist philosopher of the modern classical period. He rejects sense experience as a trustworthy source of knowledge early in his *Meditations*. Following Descartes, a number of other European philosophers develop rationalist philosophical systems. Leibniz and Spinoza are the most notable. Meanwhile, an empiricist tradition gets started in Great Britain. The three major empiricist philosophers are John Locke, Berkeley and David Hume. In this chapter we will focus on Descartes, Spinoza, and Leibniz, and we will take up the empiricists in the next chapter.

Descartes

Rene Descartes (1596 – 1650) lived during an intellectually vibrant time. European scholars had supplemented Catholic doctrine with a tradition of Aristotle scholarship, and early scientists like Galileo and Copernicus had challenged the orthodox views of the Scholastics. Surrounded by conflicting yet seemingly authoritative views on many issues, Descartes wants to find a firm foundation on which certain knowledge can be built and doubts can be put to rest. So he proposes to question any belief he has that could possibly turn out to be false and then to methodically reason from the remaining certain foundation of beliefs with the hope of reconstructing a secure structure of knowledge where the truth of each belief is ultimately guaranteed by careful inferences from his foundation of certain beliefs.

When faith and dogma dominate the intellectual scene, “How do we know?” is something of a forbidden question. Descartes dared to ask this question while the influence of Catholic faith was still quite strong. He was apparently a sincere Catholic believer, and he thought his reason-based philosophy supported the main tenants of Catholicism. Still he roused the suspicion of religious leaders by granting reason authority in the justification of our beliefs.

Descartes is considered by many to be the founder of modern philosophy. He was also an important mathematician and he made significant contributions to the science of optics. You might have heard of Cartesian coordinates. Thank Descartes. Very few contemporary philosophers hold the philosophical views Descartes held. His significance lays in the way he broke with prior tradition and the questions he raised in doing so. Descartes frames some of the

big issues philosophers continue to work on today. Notable among these are the foundations of knowledge, the nature of mind, and the question of free will. We'll look briefly at these three areas of influence before taking up a closer examination of Descartes' philosophy through his *Meditations of First Philosophy*.

To ask "How do we know?" is to ask for reasons that justify our belief in the things we think we know. Descartes' *Meditations* provide a classic example of the epistemological project of providing systematic justification for the things we take ourselves to know, and this remains a central endeavor in epistemology. This project carries with it the significant risk of finding that we lack justification for things we think we know. This is the problem of skepticism. *Skepticism* is the view that we can't know. Skepticism comes in many forms depending on just what we doubt we can know. While Descartes hoped to provide solid justification for many of his beliefs, his project of providing a rational reconstruction of knowledge fails at a key point early on. The unintended result of his epistemological project is known as the problem of Cartesian skepticism. We will explain this problem a bit later in this chapter.

Another area where Descartes has been influential is in the philosophy of mind. Descartes defends a metaphysical view known as dualism that remains popular among many religious believers. According to this view, the world is made up of two fundamentally different kinds of substance, matter and spirit (or mind). Material stuff occupies space and time and is subject to strictly deterministic laws of nature. But spiritual things, minds, are immaterial, exist eternally, and have free will. If dualism reminds you of Plato's theory of the Forms, this would not be accidental. Descartes thinks his rationalist philosophy validates Catholic doctrine and this in turn was highly influenced by Plato through St. Augustine.

The intractable problem for Descartes' dualism is that if mind and matter are so different in nature, then it is hard to see how they could interact at all. And yet when I look out the window, an image of trees and sky affects my mind. When I will to go for a walk, my material body does so under the influence of my mind. This problem of mind-body interaction was famously and forcefully raised by one of the all too rare female philosophers of the time, princess Elisabeth of Bohemia.

A whole branch of philosophy, the philosophy of mind, is launched in the wake of problems for substance dualism. Today, the philosophy of mind is merging with neuroscience, cognitive psychology, and information science to create a new science of mind. We are rapidly learning how material brains realize the processes of thought. Once again, Descartes has failed in a most fruitful way. We also see how undeserved philosophy's reputation for failing to answer its questions is. While many distinctively philosophical issues concerning the mind remain, the credit for progress will go largely to the newly minted science of mind. The history of philosophy nicely illustrates how parenthood can be such worthwhile but thankless work. As soon as you produce something of real value, it takes credit for itself. Later in a chapter on the

philosophy of mind we will examine some developments in this area since Descartes and get acquainted with a few of its contemporary issues including the nature of consciousness.

The final big issue that Descartes brought enduring attention to is the problem of free will. We all have the subjective sense that when we choose something we have acted freely or autonomously. We think that we made a choice and we could have made a different choice. The matter was entirely up to us and independent of outside considerations. Advertisers count on us taking complete credit and responsibility for our choices even as they very effectively go about influencing our choices. Is this freedom we have a subjective sense of genuine or illusory? How could we live in a world of causes and effects and yet will and act independent of these? And what are the ramifications for personal responsibility? This is difficult nest of problems that continues to interest contemporary philosophers.

Descartes' is also a scientific revolution figure. He flourished after Galileo and Copernicus and just a generation before Newton. The idea of the physical world operating like a clockwork mechanism according to strict physical laws is coming into vogue. Determinism is the view that all physical events are fully determined by prior causal factors in accordance with strict mechanistic natural laws. Part of Descartes' motivation for taking mind and matter to be fundamentally different substances is to grant the pervasive presence of causation in the material realm while preserving a place for free will in the realm of mind or spirit. This compromise ultimately doesn't work out so well. If every event in the material realm is causally determined by prior events and the laws of nature, this would include the motions of our physical bodies. But if these are causally determined, then there doesn't appear to be any entering wedge for our mental free will to have any influence over our bodily movements.

Now we will turn to Descartes' *Meditations* and examine how he comes to the positions just outlined. Here is a link to several of Descartes' writings including *Meditations on First Philosophy*: [EMT - René Descartes \(earlymoderntexts.com\)](http://EMT - René Descartes (earlymoderntexts.com))

The *Meditations*

Descartes project in his meditations is to carry out a rational reconstruction of knowledge. Descartes is living during an intellectually vibrant time and he is troubled by the lack of certainty. With the Protestant Reformation challenging the doctrines of the Catholic Church, and scientific thinkers like Galileo and Copernicus applying the empirical methods Aristotle recommends to the end of challenging the scientific views handed down from Aristotle, the credibility of authority was challenged on multiple fronts. So Descartes sets out to determine what can be known with certainty without relying on any authority, and then to see what knowledge can be securely justified based on that foundation.

In the first meditation we are introduced to Descartes' method of doubt. According to this method, Descartes goes through all of his beliefs, not individually but by categories, and asks whether there is any possible way that beliefs of this or that type can be mistaken. If so, they

must be set aside as doubtful. Many of these beliefs may ultimately be redeemed as knowledge, but they cannot serve as part of the secure foundation of indubitable beliefs from which his rational reconstruction of knowledge proceeds. Empirical beliefs, things that we believe based on the evidence of our senses, are set aside first. Our senses sometimes deceive us, as when an oar appears bent in water or a stranger in a crowd appears to be a friend. It won't do to say that we can reliably diagnose these cases and correct for mistaken appearances though because we also have experiences just like seemingly reliable sense experiences that are anything but in the case of dreams. How can we be certain that any of our seeming sense experiences of the external world aren't in fact dreams? How can we be certain that our whole life isn't a dream?

So sense experience is set to the side as uncertain and insufficient for justifying knowledge. Descartes then considers things we might know for certain by the light of reason, like mathematical claims. I seem to be about as certain in my belief that $2+2=4$ as I can be about anything. Is there any possible way I could be mistaken? Descartes here imagines a powerful demon that could deceive me into always thinking that $2+2=4$ when in fact this is not true. Is this a genuine possibility? Descartes allows that it is and considers all such knowledge had through reason doubtful as well.

Does anything remain? Are there any beliefs that can't be doubted, even given the hypothesis of a powerful evil deceiver? Descartes does find at least one. Even an evil deceiver could not deceive Descartes about his belief that he thinks. At least this belief is completely immune from doubt, because Descartes would have to be thinking in order for the evil deceiver to deceive him. In fact there is a larger class of beliefs about the content of one's own mind that can be defended as indubitable even in the face of the evil deceiver hypothesis. When I look at the grey wall behind my desk I form a belief about the external world; that I am facing a grey wall. I might be wrong about this. I might be dreaming or deceived by an evil deceiver. But I also form another belief about the content of my experience. I form the belief that I am having a visual experience of greyness. This belief about the content of my sense experience may yet be indubitable. For how could the evil deceiver trick me into thinking that I am having such an experience without in fact giving me that experience? So perhaps we can identify a broader class of beliefs that are genuinely indubitable. These are our beliefs about the contents of our own mind. We couldn't be wrong about these because we have immediate access to them and not even an evil deceiver could misdirect us.

The problem Descartes faces at this point is how to justify his beliefs about the external world based on the very narrow foundation of his indubitable beliefs about the contents of his own mind. And this brings us to one of the more famous arguments in philosophy: Descartes' "Cogito Ergo Sum" or "I think, therefore I exist." Descartes argues that if he knows with certainty that he thinks, then he can know with certainty that he exists as a thinking being. Many philosophers since then have worried about the validity of this inference. Perhaps all we are entitled to infer is that there is thinking going on and we move beyond our indubitable foundation when we

attribute that thinking to an existing subject (the “I” in “I exist”). There are issues to explore here. But bigger problems await Descartes, so we will just note this one and let it pass.

So far Descartes has only adequately justified his beliefs about the contents of his own mind and his own existence as a thinking being. Knowledge about any external reality or even truths of reason like $2+2=4$ remain in need of justification. To overcome skepticism about these matters, Descartes sets out to prove that God exists and is not an evil deceiver. Once the evil deceiver hypothesis is in check, Truths of reason and perhaps others may be yet be knowable. Not any argument for God’s existence and good nature will do, though. The trick for Descartes’ project of a rational reconstruction of knowledge is to prove the existence of a good God by reasoning only from those beliefs that he has identified as indubitable and foundational.

Descartes argument for the existence of a good God goes roughly as follows:

1. I find in my mind the idea of a perfect being.
2. The cause of my idea of a perfect being must have at least as much perfection and reality as I find in the idea.
3. I am not that perfect.
4. Nothing other than a good and perfect God could be the cause of my idea of a perfect being.
5. So, a good and perfect God must exist.

This argument simplifies the rather involved reasoning Descartes goes through in the *Meditations*. But it will do for diagnosing the fatal flaw in Descartes’ reasoning. Let us grant the validity of the argument and consider the truth of its premises. Keep in mind that to accord with the method Descartes has set for himself in carrying out a rational reconstruction of well grounded certain knowledge, all of the premises of this argument must be indubitable and foundational. Being a belief about the contents of his own mind, we can grant the certain truth of premise one. Though it is not as clear, premise three might arguably count as a foundational belief about the contents of Descartes’ own mind. An evil deceiver, being evil, would lack perfection found in Descartes’ idea of a perfect being. So as powerful as such a being could be, the cause of Descartes’ idea of a perfect being must be more perfect than any evil deceiver. Perhaps any being so perfect would have to be a good God.

But the fatal flaw for Descartes’ rational reconstruction of knowledge is the second premise. What are our grounds for thinking that the cause of something must have at least as much perfection as its effect? The idea of degrees of perfection and the notion that the less perfect can only be explained in terms of the more perfect is an idea that we find in Plato’s theory of forms. It will strike many of us as implausible or even incomprehensible. Just what is perfection supposed to mean here? And even once we’ve spelled this out, why think causes must be more perfect? It seems not at all uncommon for less perfect things to give rise to more perfect things (just consider my son, for instance). In any case, whether the second premise can be explained

and defended at all, the fatal flaw for Descartes' project is that it is not foundational. It is not an indubitable belief about the contents of Descartes' own mind, but rather a substantive belief about how things are beyond the bounds of Descartes' own mind. So Descartes' attempt to provide a rational justification for a substantive body of knowledge leaves us with an enduring skeptical problem. All we have immediate intellectual access to is the contents of our own minds. How can we ever have knowledge of anything beyond the contents of our own mind based on this? This is the problem of Cartesian skepticism.

Having diagnosed the fatal flaw in Descartes' project, we should briefly consider how his rational reconstruction of knowledge was to go from there. Given knowledge of God's existence and good nature, we would appeal to this to assure the reliability of knowledge had through reason and later also through the senses. God being the most perfect and good being would rule out the possibility of interference by an evil deceiver. We might still make mistakes in reasoning or be misinformed by the senses. But this would be due to our failure to use these faculties correctly. A good God, however, would not equip us with faculties that could not be trusted to justify our beliefs if used properly. This is a very cursory summary of the later stages of Descartes' attempted rational reconstruction of knowledge in his *Meditations*. But it will suffice for our purposes.

The mind-body problem

Descartes is a substance dualist. This is the metaphysical view that the world is made up of two fundamentally different kinds of substance: matter and spirit (or mind). In the Second Meditation Descartes motivates this view by arguing that there are distinguishing differences between the mind and the body. In particular, I can doubt the existence of my body but I can't doubt the existence of my mind. Is this a difference that justifies denying that the mind is in some sense identifiable with the body? If something is true of one thing and not of another, then we have conclusive grounds for thinking they are not one and the same thing. So if my favorite bike is red but the bike in my office is yellow, then the bike in my office is not identical to my favorite bike. Does this straightforward line of reasoning apply to the case of the mind and the body? The existence of my body is dubitable, but the existence of my mind is indubitable. Descartes would count this as a reason for denying that my mind is identical with my body. But consider this analogous argument:

1. Mark Twain is such that Joe thinks he is the author of Huckleberry Finn.
2. Samuel Clemens is not such that Joe thinks he is the author of Huckleberry Finn.
3. So, Mark Twain is not identical to Samuel Clemens.

Clearly the conclusion that Mark Twain and Samuel Clemens are not one and the same person does not follow in this case. As the Mark Twain argument looks like it is closely analogous to Descartes' argument for the non-identity of the mind and the body, it looks like Descartes' argument is not valid. The problem here is that the premises of both arguments concern someone's mental states about something. In Descartes' argument we have premises about what

he can or can't doubt. In the Mark Twain argument we have premises about what Joe does or doesn't believe. But people can fail to recognize true identity claims. Joe doesn't know that Mark Twain just is Samuel Clemens. Because of this, Joe can believe one thing about Mark Twain and something different about Samuel Clemens. But this doesn't show that Mark Twain and Samuel Clemens aren't identical. Likewise, it may be that the mind is identical with the body or some part of it, but since Descartes doesn't know this, he can believe one thing about the body (that its existence is doubtable) and something else about the mind (that its existence is not doubtable).

So far we have just offered a critical evaluation of one of Descartes' arguments for mind/body dualism. Now we will consider a serious problem for the view. When Descartes' considers how the substance of mind and body differ, he offers a view that should sound familiar from popular religious belief. On this view, the body is a physical object that exists in space and time and is subject to the laws of nature. The mind, being spiritual in nature, exists eternally in an abstract realm rather than existing in the physical realm of space and time. Further, the mind is not bound by mechanistic laws of nature, but it has free will that allows it to will or not will to do one thing or another. Descartes was both a believer in Catholicism and an active participant in the scientific revolution. He was among those who were developing a view of the natural world in which events occur in accordance with strict law-like regularities. A view of the natural world as functioning like a predictable clockwork mechanism was on the rise. And yet Descartes' Christian theology held that as a person created in the image of a divine being, he had free will through which he might choose to do one thing or another, perhaps most notably, to choose to accept the Catholic faith as true and be saved or not. His philosophical view is an attempt to reconcile these conflicting scientific and theological perspectives.

An unfortunate fact of history is that women in Descartes' time were rarely given a thorough education or allowed to participate fully in intellectual life. A notable exception is the case of the Princess Elisabeth of Bohemia. And she was among the first to notice serious difficulties in the substance dualism that Descartes advocated. The central problem has to do with mind-body interaction. Clearly things going on in the physical realm have an influence on the mind. Light reflecting off clouds and trees cause me to have the mental perception of a sunset. And likewise, mental phenomenon cause things to happen in the physical world. When I mentally will to preserve the image of the sunset in a picture, my body causes things to happen in the law governed physical realm. I reach for my camera. But how can a non-physical soul be affected by or effect events in the physical realm? If events in the physical realm are all transfers of physical energy happening at specific places and times, how can it be that the non-physical mind has any role to play in this? The problem gets all the more difficult when we take the physical world to be deterministic, governed by laws where each event is determined to happen by prior events in conjunction with mechanistic laws. Determinism in the physical realm would appear to leave no room for the non-physical mind to influence events at all. Contemporary philosophers who study the nature of the mind generally take these problems to be intractable and to constitute decisive objections to Descartes' substance dualism. More recent philosophy of mind has been mainly

taken the mind to be physical. And philosophers, along with neuroscientists and cognitive psychologists, are making tremendous progress in understanding how mental phenomenon can be understood in biological, physical terms. When we take a closer look at more recent developments in the philosophy of mind, however, we will find some arguments for denying that some mental properties, consciousness in particular, can ever be identified with purely physical properties or processes.

Study Questions for Meditations 1-3

1. Explain Descartes' method of doubt. What is Descartes' purpose in exercising this method?
2. Why can't Descartes be certain about beliefs he acquires through the evidence of the senses?
3. Why can't Descartes be certain about mathematical beliefs, like the belief that $2+2=4$?
4. What belief(s) does Descartes ultimately identify as indubitable?
5. Why can't an evil deceiver deceive Descartes about his belief that he thinks?
6. How does Descartes build up from the foundation of indubitable beliefs?
7. How does Descartes argue for the existence of God?
8. Given the existence of God, how does Descartes justify his beliefs based on reason and on the senses?

Spinoza

[Spinoza](#) was the most rigorous and systematic of the major rationalist philosophers. Where Descartes was confident that reason would vindicate the main tenets of his Catholic faith, Spinoza dared to follow reason into religiously more treacherous territory. Spinoza is alternately described as the "God intoxicated Jew" and as a heretical atheist. Spinoza's family had fled the Inquisition in Portugal for the relative religious tolerance of Amsterdam. At the age of 23, however, Spinoza was excommunicated from his Jewish synagogue for holding heretical beliefs. Spinoza knew religious persecution both as a member of a community of faith, and then as an individual cast out of a community of faith. Perhaps not surprisingly then, he becomes an early advocate for freedom of conscience and religious belief in his political writings.

Spinoza supported himself as a lens grinder until his early death at the age of 46. While he demonstrated immense courage in the face of excommunication and in philosophically challenging religious doctrines, he led a fairly quiet and reclusive life devoted to study and work in a state of relative poverty. His views were widely considered so threatening to established religion that his considerable influence went largely unacknowledged for a century or two after his death. Still, the next most significant thinker of this period after Descartes, Leibniz, is now seen by some as devising his philosophical system as an attempt to protect religious belief from the intellectual threat of Spinozism.

Descartes' method of doubt sets him on the project of finding epistemological foundations for knowledge. Descartes seeks to identify some knowledge as foundational in the sense of being able to justify the rest of our knowledge. By contrast, we might best understand Spinoza as seeking metaphysical foundations. Suppose the world is intelligible, that its nature can be understood rationally. Assuming this, what must the world be like? We might worry that this approach simply sidesteps epistemological worries about how we can know. But suppose that in exploring the assumption that the world is intelligible we find that all but one view about the nature of the world gets us mired in contradictions or intractable problems. We would then have grounds to accept the one coherent metaphysical account of the nature of the world as an instance of inference to the best (or perhaps, the only) explanation. Seeking coherent systematic explanation can, ultimately, yield justifying reasons.

This is just a suggestion for how to understand what Spinoza is up to in his masterpiece *The Ethics*. This strategy is not made explicit in the work itself. Rather, Spinoza's *Ethics* is written in a geometric style. He begins with a few definitions and axioms and the work proceeds by deductively proving an impressive array of further propositions. The propositions derived from his initial definitions give an account of God, the natural world (these turn out to be the same thing), the self, the nature of human freedom, the nature of the emotions, and the nature of the good life in-so-far as it is attainable for beings like ourselves. We might say of the entire system that it is elegantly consistent. But why accept its starting points? His initial definitions and axioms might strike us as arbitrary or even implausible (though his contemporaries would have found them pretty reasonable). The case for the system as a whole is that it is elegant and consistent while the alternatives are not. The axioms and definitions are not just arbitrarily preferred starting places; they are the starting places that allow us to give a clear coherent picture of God, the world, and the human condition.

It would help to see how Spinoza might make this case by understanding how his view of the world is offered in response to an alternative, Descartes', which did seem to lead to intractable problems. Recall Descartes' dualism, his view that the world contained two fundamentally different kinds of substance: matter and mind. The difficult problem for this view was to give some account of how mind and matter could interact in spite of being so different. But however that problem is to be solved, there is something further to be noticed. Any kind of mind body interaction will perforce involve mutual limitations on each. If through a mental act of will I cause some change in the material realm, then the material realm is limited in that it can't be other than I have willed it. Likewise, if the material world has some effect on my mind, then my mind is similarly limited.

Now consider the idea of God. Spinoza defines God as a being that is infinite, where being infinite entails being unlimited. The only way that any substance could be absolutely unlimited is for there to be no other substances that could possibly limit it. So, argues Spinoza, there is only one substance and it is both God and nature. Every facet of the world is a mere part of this one substance, God/nature. And everything we do and experience is a limited manifestation of the

essence of God. Every aspect of our lives, everything we think and do, is an expression of God/nature's essence which is uncaused and necessary. For this reason, nothing we do or experience could possibly be any different. This settles the matter of free will, though not quite in the way Descartes would hope.

Our perception of the world as including many distinct things and minds other than our own is a confusion of ours, or, as Spinoza would put it, an "inadequate idea." The true nature of the world is singular. There is only one thing in existence, and it is both God and all of nature. God/nature being the one existing substance is self-sufficient. Since it depends on nothing and is affected by nothing, everything about God/nature is necessary. God/nature, being infinite and perfect in all respects, has an infinite number of aspects, or attributes. Our existence as human beings presents us with only two of these, the attribute of thought and the attribution of extension (physical spatio-temporal existence).

While thought and extension, which we experience as mind and body, are attributes of God, our idea that there is some interaction between the two is a further confusion according to Spinoza. The mind and the body are really one and the same. We are limited modification of God/nature. One of the ways we are limited is in only being aware of two of the infinite attributes of God, thought and extension. The idea that the mind and the body are different and interact is a confusion of ours that we suffer due to thinking of ourselves sometimes under one attribute, thought, and at other times under another, extension. In thinking about ourselves, we are in position much like Joe who thought of a particular individual in one way, as Mark Twain, and also in a different way, as Samuel Clemens. Spinoza's view is that mind and body are one and the same limited modification of God, understood on one hand through the attribute of thought and on the other through the attribute of extension. A better way to put this might be to just say that the mind is the idea of the body.

We are finite and imperfect "modes" of the attributes of thought and extension. As such limited and imperfect beings, we see ourselves as separate from many other things. Being ignorant of the causes of things, including the determination of our own wills, we imagine that things might have been otherwise. But everything happens of necessity. So Spinoza's answer to the problem of free will and determinism is to deny that we have free will. This doesn't mean, however, that there is nothing to say about how to live well. Living well, according to Spinoza, involves coming to terms with our limitations and the way things must be as a matter of necessity. And the way to do this is through better understanding ourselves, the world (God/nature) and our position in the world. The good life, for Spinoza, is one organized around the intellectual love of God/nature.

There is one kind of freedom that we might aspire to in all of this, and it is the kind of freedom that can be had through the intellectual love of God/nature. The freedom we can have is freedom from the tyranny of our passions, our emotions. Our hopes and fears are passions that make us anxious and insecure when we fail to understand their causes and our own place in nature. A

better understanding of the necessity of all things, which for Spinoza is just the intellectual love of God/nature, is the one therapy open to us in addressing the insecurity and anxiety that comes with human vulnerability and mortality. Knowledge of how to live one's life is established after the manner of a proving a theorem of geometry in Spinoza's *Ethics*. Coming to understand his demonstration of how to live well will itself be an exercise in living well.

Leibniz

Among quite a few other things, [Leibniz](#) was an important mathematician. He and Newton vied for credit for discovering the calculus of infinitesimals. He was also politically active as an advisor to assorted rulers and aristocrats. Like Descartes, Leibniz was, at least publicly, religious. His grandest political ambition was to see the Christian church re-unified (recall that Protestants had broken off from Catholics over the prior few centuries). Leibniz was arguably the first to have imagined anything like information technology. Among his grand ambitions was to formulate a universal symbolic language for science and philosophy that would be rigorously rule driven and free of all ambiguity. He even got as far as constructing a calculating machine, though not a very reliable one.

Leibniz' metaphysical views seem pretty exotic at first glance. Leibniz took the world to consist of monads. Each monad is simple and indivisible. But monads are not merely physical, like atoms. Each monad would include both a physical aspect and a mental aspect. Physical objects are made up of monads that are also minds, just particularly dim-witted ones. Monads appear to interact with each other. We seem to influence each other and make things happen in the physical world. But according to Leibniz there is no actual interaction between monads. Instead, monads exist in a harmony that is pre-established by God. As a result, like an element in a spectral image or a droplet in a cloud, each monad carries in it a reflection all of creation.

If this seems to be a rather exotic picture of the world, let's review the problems Leibniz is trying to negotiate in the wake of Descartes and Spinoza. The problem of mind/body interaction looms large after Descartes. If mind and body are distinct kinds of substances, then it is very hard to see how either can have any influence on the other. Leibniz metaphysics handles this problem neatly by making his substances, monads, have mind as an integral part. We needn't worry about mind-body interaction if mind and body are already unified. Next, bear in mind the theologically challenging aspects of Spinoza's monism. In taking there to be just one substance, Spinoza identifies God with all of nature and denies that people have any existence distinct from God/nature. God is not personal on this view. God/nature is really nothing like us at all. Spinoza's God is so unlike the traditional God of Christianity that Spinoza is widely deemed to be an atheist. Worse, in taking humans to be mere parts of a self-caused and hence necessary God, we lack free will entirely on Spinoza's view. Leibniz is eager to provide a philosophical route to avoiding Spinoza's atheism and denial of free will. To avoid atheism, and in particular a variety of atheism where people are mere parts of an impersonal God/nature, Leibniz needs to posit a plurality of substances. Monads fit the bill. In order to preserve free will, which is also

central to Christian theology, Leibniz needs for the substances that are mind not to be causally determined by other substances. The pre-established harmony of monads is his means of achieving this. But while Leibniz thereby avoids causal determinism, he seems to be saddled with a kind of theological determinism instead. Everything that happens, including every choice you make, will have been determined by God.

Leibniz was both intrigued and repelled by Spinoza's thought. The two met for a few days while Leibniz was ostensibly on a diplomatic mission in Amsterdam. As much as Leibniz abhorred Spinoza's views, he couldn't dismiss Spinoza's carefully reasoned and systematic response to Descartes' thought. As a result, Leibniz devotes a considerable amount of creative intellectual energy to finding some way to avoid Spinoza's heretical conclusions. Such was the influence of the outcast Jew of Amsterdam.

Review and Discussion Questions

1. Explain Descartes' method of doubt. What is Descartes' purpose in exercising this method?
2. Why can't Descartes be certain about beliefs he acquires through the evidence of the senses?
3. Why can't Descartes be certain about mathematical beliefs like the belief that $2+2=4$?
4. What belief(s) does Descartes ultimately identify as indubitable?
5. Why can't an evil deceiver deceive Descartes about his belief that he thinks?
6. How does Descartes build up from the foundation of indubitable beliefs?
7. How does Descartes argue for the existence of a good God?
8. How does Descartes' argument for God fail?
9. Given the existence of a good God, how does Descartes justify his beliefs based on reason and on the senses?
10. How does Descartes argue for the distinction between mind and body? How does this argument fail?
11. Explain Spinoza's Monism.
12. How does Spinoza's view of God differ from more traditional theological perspectives?
13. How does Spinoza handle the mind/body problem?
14. How does Spinoza handle the matter of free will?
15. What sort of freedom can humans aspire to on Spinoza's view?
16. What are monads and how does Leibniz hope monad will help him avoid Spinoza's theologically controversial views?

6. The Empiricists

Empiricism, you might recall, is the view that all of our knowledge is ultimately acquired through by sense experience. The empiricist philosophical tradition comes to fruition in Great Britain over the course of the 17th and 18th centuries. We will discuss three major empiricist thinkers: John Locke, George Berkeley, and David Hume. We'll consider the first two briefly and focus more closely on Hume.

John Locke

John Locke (1632 –1704) is better known now for his political thought than his Empiricism. Locke spent time in Holland as a young man and his political thought was probably influenced significantly by Spinoza who had died only recently. Locke argued against the divine right of kings to rule and instead defended a liberal egalitarian political philosophy on which people have equal and natural rights to liberty. Liberty, in Locke's thought, should be understood as being free from domination by others. Liberty is not in Locke's view being free to do whatever one pleases. For starters, if everyone is to be free from domination, then it follows that nobody is free to dominate. Locke also offers the classic justification for property rights as an extension of our self-ownership. So property rights are seen as natural extensions of our human liberty. The point of government is just to secure our natural liberties to the highest degree possible on Locke's view. So government is legitimate only when it is limited to this role. This view should sound familiar. Locke's political philosophy was influential with the founding fathers of the U.S. Thomas Jefferson in particular was a close student of Locke's political thought. We will return to Locke at the end of the course when we take up political philosophy. But for now, we'll say a little about his epistemology.

Locke develops his empiricist epistemology in his *Essay Concerning Human Understanding*. Locke's approach is to examine the origins of the contents of the mind. Early in this work he argues against innate ideas. The mind starts off as a *tablula rasa*, a blank slate. All of our ideas have their origin in experience. Simple ideas, say of solidity and figure, are acquired through the senses, and from these we form complex ideas, say the idea of a dog, through the capacities of the understanding. The details of this account raise a number of challenging questions. We might think of Locke as launching a research program for developing an empiricist account of the mind rather than spelling out a fully developed view.

Locke thinks that some of the impressions we get from sense experience are genuinely similar to how things are objectively in the world. Our sense experience of the shape of things, for instance, reflects the ways things really are according to Locke. Locke refers to the qualities where there is a resemblance between our experience and the way things are as **primary qualities**. Shape, motion or rest, and number are a few of the primary qualities. Other aspects of

our sense experience don't resemble the qualities in their objects. The taste of an apple, for instance, is not really in the apple. What is in the apple is just a power to produce the experience of a certain flavor. But we have no grounds for thinking that this power as it exists in the apple resembles in any way the sense experience we have of its taste. Locke calls qualities where our sense experience doesn't resemble the qualities that give rise to our experience **secondary qualities**. Our knowledge of the external world, then, is based entirely on our experience of the primary qualities. Empiricism, as we will see in the case of later empiricists, especially Hume, tends to place sharp limits on what is knowable.

While all experience depends on having simple ideas had through sense experience, Locke does not take experience to be limited to these. We also have experience of the operations of the mind in building up complex ideas out of simple ideas. Once you have some simple ideas through sense experience, you also have an experience of yourself and of your mental operations on those simple ideas. So given simple ideas through experience, the operations of the mind become a source for further ideas. Locke thinks knowledge of the self, God, mathematics, and ethics can be derived from this additional internal source of experience. Hume, as we shall see, is not so optimistic.

George Berkeley

George Berkeley (1685-1753) is best known for arguing for Idealism on empiricist grounds. In metaphysics, Idealism is the view that there is no physical substance underlying our sense impressions of the world. Rather, the world consists entirely of ideas. Your mind is just a bundle of impressions, and there is nothing in the world except for so many minds having their various perceptions. Berkeley defends this as the view that best accords with common sense in [*Three Dialogues between Hylas and Philonous*](#).

Berkeley's argument attacks Locke's distinction between primary and secondary qualities and argues that all of our sense impressions are mere appearances and that we have no grounds for thinking that any of them bear any resemblance to the way things are. Since we lack any empirical experience of the underlying substances in which qualities inhere, we have no empirical reason to suppose underlying substances even exist. All we have access to are our sense impressions, and these are mental things, ideas. So all we can claim knowledge of are our ideas beginning with our sense impressions, the most basic ideas.

Berkeley also argues that positing underlying substances do no significant explanatory work. So, the common sense empiricist view ought to be that we live in a world of ideas that lacks any underlying physical substance. This startling view might make us wonder what happens to my desk when I leave the room and cease to perceive it. Does it pop out of existence when I leave and then pop back into existence just as it was when I return to my work? This would be most peculiar. Berkeley argues that the objects of our everyday life do have an enduring existence when we are absent. They continue to exist as ideas in the mind of God. Given this appeal to the

mind of God to explain the continued existence of things we aren't actively observing, we might argue that positing underlying substances does some explanatory work after all and charge that Berkeley has only substituted one unobservable theoretical posit, God, for another, underlying substances.

David Hume

Of the philosophers discussed here, David Hume (1711-1776) has probably had the greatest influence on contemporary philosophy. The twentieth century begins with a movement known as Logical Positivism that tests the limits of Empiricism. The Empiricism of the Logical Positivists is heavily indebted to Hume, as we'll see in the following chapter of Philosophy of Science.

Hume's empiricist epistemology is grounded in his philosophy of mind. Hume starts by asking what we have in the mind and where these things come from. He divides our mental representations into two categories, the relatively vivid **impressions**, these include sensations and feelings, and the less vivid **ideas** which include memories and ideas produced by the imagination.

What distinguishes impressions from ideas in our experience is just their vividness. The picture of the mind Hume offers is one where all of our beliefs and representations are cooked up out of basic ingredients provided by experience. Our experience gives us only impressions through sense experience and internal impressions like feelings. From this we generate less vivid ideas. Memories are merely faint copies of impressions. Through the imagination we can generate further ideas by recombining elements of ideas we already have. So through impressions we get the idea of a lizard and the idea of a bird. We can then generate the idea of a dragon by imaginatively combining elements of each. In cooking up new ideas from old ideas, the imagination is guided by associating relations like resemblance, contiguity (next-to-ness) and cause and effect. So, for example, an impression of a grapefruit might lead me to think of an orange due to their similarity. The thought of my bicycle might lead me to think of the ride I plan to go on this weekend. Through the association of cause and effect, my idea of a struck match leads me to the idea of a flame. The last of these principles of association, cause and effect, turns out to be faulty for reasons we will examine shortly.

The imagination is not merely a source of fancy and fiction. The imagination also includes our ability to understand things when we reason well in formulating new ideas from old ones. Think of our imagination here as just our capacity for making mental images, or, more broadly, representations. Our imagination is involved as much in representing things accurately as it is in representing things fancifully.

Hume empiricism is the view that all of our knowledge of the world is ultimately based on sense experience. To get a clear picture, we'll need to dive in to a bit more detail. Hume distinguishes

our beliefs into two categories, matters of fact and relations among ideas. Matters of fact are beliefs about how the world is. His empiricism consists in the idea that all matters of fact are ultimately justified *a posteriori*, which just means based on sense experience. Beyond matters of fact, we also have beliefs that are understood *a priori* or independent of experience. These are mere relations among ideas, conceptual truths like “All bachelors are unmarried.” Relations among ideas, which we can know *a priori*, tell us nothing about the world. Knowing that all bachelors are married gives you no reason to believe that David is a bachelor, for instance. A *a priori* justification of relations among ideas gives us no knowledge of how things are in the world, it only gives us an understanding of how ideas are related to each other.

We can summarize Hume’s empiricism in this table:

Matters of fact (This includes all our knowledge of the world)	Relations among ideas (This concerns only things that are conceptually true or false)
justified <i>a posteriori</i> (on the basis of sense experience)	Justified <i>a priori</i> (independent of sense experience)
For example: David is a bachelor	For example: All bachelors are unmarried

On this view, if a belief can’t be justified *a posteriori* as matter of fact or *a priori* as a relation among ideas, then it can’t be justified. Skepticism will be the result whenever something can’t be justified in one of these two ways. This strategy, known as Hume’s Fork, is deployed in arguing for a variety of skeptical results. Hume argues that we can’t have knowledge of causation, the rationality of inductive argument, morality, or even the existence of an external world by clever applications of Hume’s Fork. First, we’ll examine the application of Hume’s fork to the cases of causation and induction. Then we’ll consider a few further skeptical results Hume argues for.

Let’s take note of the logical form of Hume’s Fork. According to Hume’s empiricism, if we can justify knowledge of something, it must be justified either *a posteriori* through sense experience, if it is a substantive matter of fact, or it can be justified *a priori* as a relation among ideas, if it is merely a conceptual truth that makes no substantive claim about the world. If something can’t be justified in one of these two ways, then we reach the skeptical conclusion that we can’t have knowledge of it. Here’s the argument strategy of Hume’s Fork stated more formally:

1. If knowledge of X is justifiable, then either X can be justified *a posteriori* or X can be justified *a priori*.
2. X can’t be justified *a posteriori*
3. X can’t be justified *a prior*
4. So, knowledge of X is not justifiable (skeptical conclusion)

Now consider just the logical form:

1. If P, then Q or R
2. Not Q

3. Not R
4. Therefore, not P

Recall that a valid argument is one where if the premises are true, then the conclusion must be true. Alternatively, a valid argument is an argument where there is no logically possible way for all the premises to be true and the conclusion false.

Can you see the validity of the argument strategy behind Hume's Fork?

Hume's Skepticism about Causation

When we examine our everyday idea of causation, Hume says we find four component ideas:

- the idea of a constant conjunction of cause and effect (whenever the cause occurs, the effect follows).
- the idea of the temporal priority of the cause (the cause happens first, then the effect).
- the idea of causes and effects being contiguous (next to each other) in space and time.
- the idea of a necessary connection between the cause and the effect.

So, for instance, the idea that striking a match causes it to light is made up of the idea that whenever similar matches are struck (under the right conditions), they light, plus the idea of the striking happening first, and the idea of the striking and the lighting happen right next to each other in time and space, and, finally, the idea that the striking somehow necessitates or makes the match light.

So far, we have just examined a relation of ideas that is knowable *a priori*. Our idea of causation consists in these four more basic ideas of constant conjunction, temporal priority of the cause, contiguity in space and time and a necessary connection.

Now let's consider these component ideas and ask whether they all have an empirical basis in corresponding sense impressions. We do have sense impressions of the first three: the constant conjunction of cause and effect, the temporal priority of the cause, and the contiguity of cause and effect. But Hume argues that we lack any corresponding empirical impression of necessary connections between causes and effects. We don't observe anything like the cause making the effect occur. As Hume puts the point,

When we look about us towards external objects, and consider the operation of causes, we are never able, in a single instance, to discover any power or necessary connexion; any quality, which binds the effect to the cause, and renders the one an infallible consequence of the other. We only find, that the one does actually, in fact, follow the other. (*An Enquiry Concerning Human Understanding*, Section VII)

Now we have explored one prong of Hume's Fork and found that while three of the component ideas in our complex idea of causation have a basis in sense experience, the fourth, that of a necessary connection between cause and effect, is not something we can know *a posteriori* on the basis of sense experience.

Turning to the second prong of Hume's fork, can we have *a priori* knowledge of causes necessitating their effects. This won't work either. There is nothing in the idea of a cause (say, the striking of the match) that contains the idea of its effect (the lighting of the match). We can imagine a world where the first event produces something very different, or nothing at all. There is nothing conceptually absurd in the idea of a world where struck matches bloom like flowers rather than igniting in flame. Necessary connections between causes and effects aren't mere relations of ideas that can be known *a priori*. For this reason, Hume denies that we have rational grounds for thinking that causes do necessitate their effects.

So, we have identified a component of our idea of causation, the idea of a necessary connection, the sense in which the cause "makes" the effect happen, which can't be justified either *a posteriori* through sense experience or *a priori* as a conceptually true relation among ideas. The result is skepticism about causation. We can't have knowledge of any event causing another.

Maybe this isn't so bad, since we can still have knowledge of the other three components of our idea of causation. Maybe this will be enough for science. Bertrand Russell thought so at one time. But things are about to get worse, because in the absence of necessary connections between causes and effects, another application of Hume's fork will lead us to skepticism about the rationality of inductive reasoning.

Hume's Skepticism about Induction

Closely related to Hume's skepticism about causation is Hume's skepticism about inductive reasoning. Inductive argument, in its standard form, draws a conclusion about what is generally the case, or what will be the case in some future instance, from some limited number of specific observations. The following is an example of a typical inductive argument:

1. Every observed sample of water heated to well over 100 C has boiled.
2. Therefore, whenever water is heated to well over 100 C, it boils.

Here we make an inference from the cases we have observed to a broader class of all the times water has been heated to 100C. So inductive arguments like the one above are not deductively valid, which is OK because they aren't aiming at validity. They are aiming at inductive strength. But then what justifies the inference from the premise to the conclusion of an inductive argument?

Hume suggests that every inductive argument has a principle of induction as a suppressed premise, and it is this principle of induction that renders the inference from premises to conclusion rational. This principle of induction tells us roughly that unobserved instances tend to

follow the pattern of observed instances. So, according to Hume, inductive arguments really go something like this:

1. Every observed sample of water heated to over 100 C has boiled.
2. (Unobserved cases tend to follow the pattern of observed cases)
3. So, whenever water is heated to over 100 C, it boils.

Of course, the argument still isn't valid, but that's not what we are aiming for in induction. Given the hidden second premise - our principle of induction - we can reasonably hold that the premises taken together give us good grounds to accept that the conclusion is probably true. However, if this principle of induction, now given as premise 2, is to render inductive inferences rational, then we need some grounds for thinking that it is true. Here is we are about to skewer the principle of induction with Hume's Fork.

Our principle of induction is a substantive claim about how the world is. The world might have been completely random with no patterns we can discern between different kinds of events. Since our principle of induction is a substantive claim about the world that might have been false, it cannot be justified *a-priori* as a mere relation among ideas.

Next, since our principle of induction is a claim that is supposed to hold generally, not just in a specific instance, an *a posteriori* justification for our principle of induction based on empirical evidence would itself be an inductive generalization. But this won't work since the rationality of inductive argument depends the principle of induction, the very thing we are trying to justify. Any inductive argument for the principle of induction would assume just what it is trying to show as a premise. This would be circular reasoning. So, Hume concludes, we have no rational grounds for accepting inductive inferences.

Think about the ramifications of Hume's skepticism about induction. If inductive argument is not rational, then we have no reason at all to think the sun will rise tomorrow. Here we aren't worried about improbable possibilities like the sun getting blown to bits by aliens before tomorrow morning. Lack of certainty is not the issue here. Hume's argument against the rationality of inductive reasoning implies that all of our experience of the sun regularly rising gives us no reason to think its rising tomorrow is even likely to happen. If this sounds crazy, then we have a problem because it is not easy to find a defect in Hume's reasoning. This is why philosophers speak of the *Problem of Induction*. Very few are prepared to accept Hume's skepticism about induction. But in the two and a half centuries that have passed since Hume died, we have yet to settle on an entirely satisfactory solution to the problem of induction. We'll consider the problem of induction further when we take up the Philosophy of Science in the next chapter.

Morality

Some of Hume's skeptical results are not so surprising given his Empiricism. Hume is skeptical about objective moral truths, for instance. We don't get to observe rightness and wrongness in the way we can see colors and shapes, for instance. Suppose we witness some boys torturing a dog. We might see them strike the dog with a stick and we might hear the dog yelp. But we have no sense impression of the wrongness of how the boys are treating the dog. The idea that there are objective moral truths, according to Hume, is a mistaken projection of our subjective moral sentiments.

Doubts about morality on the grounds that we can't prove moral claims true or false on the basis of evidence are now commonplace. This much is in line with Hume's thinking. But people often appeal to disagreement about moral matters as support for moral skepticism, and Hume would not go along with this. There are moral issues where people disagree like whether the death penalty is justifiable, the morality of abortion in various cases, whether drug use should be prosecuted as a crime. Points of disagreement like these get a good deal of attention, but they are actually the exceptions to much broader moral agreement. We generally think it is generally good to be honest. Being prejudiced or manipulative is widely recognized as morally bad. Nobody thinks is morally OK to torture innocent puppies just for fun. We all think that theft, fraud and assault are morally wrong. Hume thought agreement about morality, not disagreement, was the striking fact that calls out for explanation. Hume thinks we can find an explanation for broad agreement on morality in our generally held moral sentiments. By and large, we all value kindness and disvalue cruelty. Hume thinks such sentiments are built into human psychology. Our moral judgments, on Hume's view, aren't claims about what is objectively true. Rather they are expressions of our subjective but generally shared moral sentiments.

So, Hume would not be impressed with the popular opinion that morality is subjective because people have different moral views (we might also note that people have differing opinions about all sorts of things that are clearly objective matters). His skepticism about morality is based entirely on his empiricism and our lack of sense impressions of right and wrong.

The External World

All of our reasoning about the external world is based on the idea of causation. Our beliefs about the external world are based on the idea that things going on in the external world cause our sense impressions. So, Hume's skepticism about causation undermines our reason for thinking there is a world beyond our mind.

More generally, our evidence for what we can know begins with our impressions, the mental representations of sense experience. We assume that our impressions are a reliable guide to the way things are, but this is an assumption we can't rationally justify. We have no experience beyond our impressions that could rationally certify that our impressions correspond in any way to an external reality. Our assumption that our impressions do correspond to an external reality is a rationally unsupportable product of our imagination.

God

Unlike Locke and Berkeley, Hume's rigorous Empiricism leads him to skepticism about religious matters. To avoid censorship or persecution, critics of religious belief in the 18th century exercised caution in various ways. Hume's earliest challenge to religious belief, an essay on miracles, was removed from his early work, his *Treatise of Human Nature*, and published only in his later *Enquiries Concerning Human Understanding*. In this essay, Hume argues that the belief in miracles can never be rational. A miracle is understood to be a violation of the laws of nature resulting from Divine will. But, argues Hume, the weight of the evidence of our experience overall will always give us stronger reason to mistrust our senses in the case of a seemingly miraculous experience than to doubt the otherwise consistently regular course of events in our experience. Testimony by others of miracles is on even shakier ground.

No testimony is sufficient to establish a miracle, unless the testimony be of such a kind, that its falsehood would be more miraculous than the fact which it endeavors to establish.
(*Enquiries Concerning Human Understanding*, Section 10)

Among educated people in the 18th century, religious belief was thought to be supported not just by Divine revelation, but by our experience of the natural world as well. When we look to the natural world we find impressive harmony in the natural order of things. The various species all seem well suited to their environments and ecological stability is maintained by the various roles organisms play in their environments. To the discerning mind in search of an explanation, the order and harmony we find in the world looks very much like the deliberate work of a Divine creator. This line of thought is known as the Argument from Design. Hume's last work, his posthumously published *Dialogues of Natural Religion*, aimed to undermine many arguments for the existence of God, including the Design Argument.

According to Hume, the Design Argument is a weak argument by analogy. We have reason to think that machines are the product of human design because we are familiar with their means of production. But we have no analogue in the case of the universe. We have not observed its creation. The alleged similarity of the universe to machines designed by humans is also suspect. We do find regularities in nature, but only in the small corner of nature we are familiar with. The regularity, order, and harmony we do find don't provide enough of the appearance of design to warrant positing an intelligent designer, according to Hume. But suppose we do think the natural world bears the marks of a designer's craftsmanship. The only sorts of designers we are familiar with are people like us. But that doesn't tell us much about what sorts of being could be designers of complex harmonious systems. So even assuming we find the appearance of design in nature, we have little grounds to think that it is the product of a personal god or any sort of entity we can relate to.

Charles Darwin's theory of evolution by natural selection provides a naturalistic account of the appearance of design in life forms. Thanks to providing a developed naturalistic alternative to the hypothesis of design by a Divine creator, Darwin probably had the greater impact in

undermining the design argument for the existence of God. Darwin cites Hume as among his major influences, and there are a number of passages in Hume's writing that foreshadow insights that Darwin developed.

The Self

Descartes didn't hesitate to infer the existence of himself from the certainty of his thinking. And it seems obvious to most of us that having thoughts implies the existence of a subject that thinks. Hume is more cautious on this point.

For my part, when I enter most intimately into what I call *myself*, I always stumble on some particular perception or other, of heat or cold, light or shade, love or hatred, pain or pleasure. I never can catch *myself* at any time without a perception, and never can observe anything but the perception." (*Treatise*, 1.4.6.3)

The contents of our immediate experience are just particular impressions and ideas. But we have no experience of any single unified self that is the subject of those experiences. The idea of a self, including the idea of the self as a soul, is a fanciful projection from our experiences. All we can say in an empirically grounded way of ourselves, according to Hume, is that we are just a bundle of experiences.

We've just given the briefest sketch of how Hume reaches his assorted skeptical conclusions. There are many further arguments and objections to consider, but hopefully we've covered enough to give you an appreciation for how carefully a strict and carefully reasoned Empiricism leads to a variety of skeptical conclusions. Hume's skepticism about causation and induction may be the most surprising. We often hold up science as the paradigm of human intellectual achievement, and we tend to think of science as pretty empirical. Yet Hume's strict Empiricism seems to undercut science on the key notions of causation and induction. Perhaps scientific inquiry is not as strictly empirical as Hume's epistemology. Or perhaps, as some have argued, science can get along fine without induction or causation. Still, if we are not comfortable with Hume's skepticism about causation and induction, this might be cause to reconsider his Empiricism. And perhaps also the skepticism about morality it seems to invite.

Review and Discussion Questions

1. Explain Locke's distinction between primary and secondary qualities.
2. While Locke thinks that having any thoughts depends on sense experience, he doesn't think sense experience is the only source of our ideas. Explain.
3. Explain Berkeley's Idealism.
4. How is Berkeley's Idealism a response to Locke's epistemology?
5. Explain Hume's view of the contents of the mind.
6. What are ideas and impressions? How does Hume distinguish these?
7. How does the imagination form new ideas, according to Hume?

8. How is Hume's Empiricism grounded in his philosophy of the mind?
9. How does Hume analyze our notion of the cause-effect relation? Explain Hume's skepticism about causation.
10. Why is Hume skeptical about the rationality of inductive argument?
11. Explain Hume's skepticism about morality. How does he argue for this view?
12. Why does Hume deny that we can have knowledge of an external world?
13. Why does Hume doubt we could ever have reason to believe in miracles?
14. What is the Design Argument for the existence of God?
15. Explain Hume's objections to the Design Argument.

7. Philosophy of Science

In high school you were probably introduced to something misleading called *the* scientific method. According to this picture of science, science proceeds by asking a question, formulating a hypothesis, designing an experiment to test the hypothesis, and analyzing the results to reach a conclusion. The experiment should be repeatable and the hypothesis is only considered well supported if our experimentation yields plenty of data in support of it. When we find plenty of data supporting our hypotheses, the pattern of reasoning employed is basic induction by enumeration where we generalize or predict based on observed patterns.

While this model does describe a frequently employed method in science, it's misleading to think of this as *the* scientific method. The disservice done to the actual practice of science by this bit of high school curriculum is really quite egregious. It's as if you were shown how to play a C major scale on the piano and then told "there you go, that's how to make music. That's the method." In actual practice, scientists employ a variety of methods that involve a broad range of patterns of reasoning, both inductive and deductive. Testing hypotheses often involves things like hunting for clues, diagnosing unexpected results, engineering new ways of detecting evidence, and a great many things beyond designing experiments and generalizing based on the results of these. The support for a hypothesis is often a matter of inference to the best explanation rather than inductive generalization. Sometimes the best analysis of data seeks alternative explanations for data anomalies that do not fit with predictions rather than automatically counting such data as evidence against a hypothesis.

Investigating the messy, gritty details that drive actual scientific practice is where the real action in the philosophy of science is today. Explaining how science advances human understanding of the world often requires a close examination of what's going on in actual scientific practice. It is not uncommon for philosophers of science to describe their work as something like the science of science. Methods are not to be prescribed up front by the philosophical lords of epistemology. Rather, in contemporary philosophy of science we look to science to see what methods actually work, and then try to better understand the significance of these.

Over the past few chapters, we have covered a couple of classic skeptical problems. In the wake of Descartes and Hume you might worry that we can't know much at all. Lots of people are willing to just let the matter rest there and think we can only have so many subjective opinions, even about scientific matters (witness, for instance, the skepticism generated by deniers of climate science). It's hard, however, to take this uncritical skepticism seriously in the face of the truly impressive achievements of science over the past few centuries. Looking at these achievements, it seems we have powerful evidence for our ability to figure things out and get at many truths. So, the suggestion I want to make at the outset of this chapter is that the way to address the skeptical problems raised by Hume might be to examine more closely the methods by which we seem to attain knowledge and begin to sort out how they work in practice. In this

chapter we will trace a few developments over the course of the 20th century with an eye to better understanding how the philosophy of science has developed into what it is today. We will start with Logical Positivism, a broad empiricist movement of the early 20th century.

Logical Positivism

The contemporary analytic tradition in philosophy that is now globally well-established got its start with Logical Positivism at the beginning of the 20th century in Vienna. Logical Positivism can be understood as Empiricism, heavily influenced by Hume, and supercharged with powerful new developments in symbolic logic. The system of logic that we now teach in college level symbolic logic courses (PHIL& 120 at BC) was developed just over a century ago in the work of Gotlob Frege, Bertrand Russell, and Albert North Whitehead for the purpose of better understanding the foundations of mathematics. In *Principia Mathematica*, Russell and Whitehead made a case for analyzing all of mathematics in terms of logic (together with set theory). According to the argument of *Principia Mathematica*, mathematical truths are not truths justified independent of experience by the light of reason alone. Rather they are derivable from logic and set theory alone. Merely logical truths are trivial in the sense that they tell us nothing about the nature of the world. Any sentence of the form ‘Either P or not P’, for instance, is a basic logical truth. But, like all merely logical truths, sentences having this form assert nothing about how the world is. Logic doesn’t constitute knowledge of the world; it is merely a tool for organizing knowledge and maintaining consistency.

Mathematics had long served as the rationalist’s paradigm case of knowledge justified through reason alone. So, we can make a powerful case for Empiricism by showing that math is really just an extension of logic. It remains debatable whether Frege, Russell, and Whitehead succeeded in showing this, but their attempt, and especially the powerful new system of logic they developed in making this attempt, constituted a powerful blow against Rationalism and inspired a group of empirically minded philosophers and scientists in Vienna to employ the same logical tools in analyzing and clarifying philosophical issues in science. As we will see, their ambitions were even grander since they also argued that much of what was going on in philosophy at the time was literally meaningless.

While we have moved on from most of the views about science promoted by the Logical Positivists, this movement represents a thorough house cleaning in philosophy. Logical Positivism marks the end of the grand system building we can see in Plato and Descartes. Philosophy since generally proceeds by addressing fairly specific problems. Speculative metaphysical theories, often seen as expressing the spirit of a culture or nation (Heidegger, for example) are thoroughly repudiated. Philosophy after the Logical Positivists proceeds as a disciplined branch of inquiry not so different in kind from science, but differing only in subject matter. While philosophers usually don’t engage in empirical experimentation (though sometimes they do), philosophers from here on are generally sensitive to the findings of empirical science and how these frame, raise, or resolve philosophical issues.

We will consider three central projects taken on by the Positivists in developing their Empiricist view of scientific knowledge. These are the demarcation problem, the problem of distinguishing science from non-science, developing a view about what a scientific theory is, and giving an account of scientific explanation. The Positivists utilize the resources of symbolic logic in each of these projects.

The Demarcation Problem

Among the main tasks the Positivists set for themselves was that of distinguishing legitimate science from other rather suspect fields and methods of human inquiry. Specifically, they wanted to distinguish science from religion, metaphysics, and pseudo-science like astrology.

The speculation and lack of clarity in most 19th century philosophy often made philosophy more representative of national or cultural viewpoints and values than straight inquiry into philosophical issues and ideas. Logical Positivism took such thinking to be literally meaningless as a representation of any shared reality. In doing so, they were taking a stance against thinking behind the rising tide of nationalism and fascism that took root in Europe, Germany in particular, and ultimately led to Nazism and the World War II. Politically, many members of the Vienna Circle that started Logical Positivism were liberal or socialist. Ethnically, many were Jewish or of Jewish ancestry. Most members of the Vienna Circle ultimately became refugees around the time Nazi Germany annexed Austria in the years leading up to World War II. Many wound up in the US or Britain.

19th century German metaphysics involved attempts to reason about such obscure notions as “the absolute,” or the nature of “the nothing.” Such metaphysics needed to be distinguished from genuine science. We had also seen appeal to obscure empirically suspicious entities and forces in Aristotelian science such as the “vital force” to explain life, or the “dormative virtue” a mysterious power of substances like opium to cause sleep. Such mysterious forces needed to be eliminated from genuine scientific discourse.

While metaphysics and talk of obscure forces in science were to be distinguished from genuine science, the Positivists needed to preserve a role for unobservable theoretical entities like atoms and electrons. The rejection of metaphysics and obscure forces must not undermine the legitimate role for theoretical entities.

The Positivists employed Empiricism in their proposed solution to the demarcation problem. Empiricism, as we know, is just the view that our sense experience is the ultimate source of justification for all of our factual knowledge of the world. The Positivists extend Empiricism to cover not just the justification of knowledge, but the meaningfulness of language as well. In particular, the Positivists advanced what they called The Verificationist Theory of Meaning (VTM). We can formulate the VTM as follows:

- A sentence is meaningful only if we can specify observable conditions that would verify it as true or false.

The kind of meaning we are concerned with here is representational. That is, the VTM is a standard of meaningfulness for language that aims to represent the world as being one way or another. There is other language, like love poetry, that isn't really concerned with what is true or false. Love poetry is more concerned with evoking feelings, and the Positivists aren't opposed to us using language for things other than describing what is true or false. They just aim to get clear on the difference and their proposal for doing this is the VTM.

The VTM has it that a sentence counts as meaningful only if we can specify the observable conditions under which it would count as true or false. This view is aimed at distinguishing empirically respectable scientific language from nonsense. So, we have a view on which science is distinguished as meaningful while pseudo-science, religion, poetry etc. are, strictly speaking, meaningless. Likewise, most of philosophy turns out to be meaningless as well. Not only will obscure 19th century German metaphysics turn out to be meaningless, but talk of free will, immaterial substances, and all of ethics will likewise turn out to be meaningless. The only legitimate role left for philosophers, according to the Logical Positivists, will be the logical analysis of scientific discourse. Being meaningless, religion, pseudo-science, most of philosophy, literature etc. is neither true nor false. While these things cannot be true or false, according to Positivists' criteria for meaningfulness, they may provide helpful expressions of human emotions, attitudes towards life, etc. That is, poetry, literature, religion, and most philosophy will be merely so much comforting or disturbing babble, mere coos, squeals, or screams.

Significant progress is made by paying close attention to the meaningfulness of scientific discourse. But the Verificationist Theory of Meaning eventually falls apart for a number of reasons including that it turns out not to be meaningful according to its own criteria. Amusingly, we can't provide an empirical test of truth or falsity for the claim that a claim is meaningful only if we can provide an empirical test for its truth or falsity. That is, according to the Verificationist Theory of Meaning, the term "meaningful" turns out to be meaningless. Logical Positivism remained a powerful influence in philosophy through much of the 20th century and it did serve to weed out some pretty incomprehensible metaphysics. But I can now happily report that other important areas of philosophy, notably ethics and metaphysics, have recovered from the Positivists' assault on philosophy from within. While we have moved on from most of the views propounded by Positivists, the more enduring influence has been to raise the standards of clarity and logical rigor in philosophy. This is necessary for philosophy to proceed as a kind of inquiry, unfettered by nationalist or culturally bound bias.

Theories

Understanding the Logical Positivist view of theories requires that we say a few things about formal languages. The symbolic logic developed in Russell and Whitehead's *Principia Mathematica* is a formal language. Computer languages are also formal languages. A formal language is a precisely specified artificial language. A formal language is specified by doing three things:

- identify the languages vocabulary.
- identify what counts as a well formed expression of that language.
- give axioms or rules of inference that allow you to transforming certain kinds of well formed expressions into other kinds of well formed expressions.

Scientific theories are formal languages according to the Positivists. We can understand what this means by considering the component parts of a scientific theory and how these map on to the elements of formal languages just given. A theory consists of the formal language of first order predicate logic with quantifiers (the logic developed first by Frege and then in greater detail by Russell and Whitehead) supplemented with observational vocabulary, correspondence rules that define theoretical terms in terms of observational vocabulary, and statements of laws like Galileo's laws of motion, Newton's law of universal gravitation etc. All of the non-logical vocabulary of a scientific theory is definable in observational terms. Well formed expressions in scientific discourse will be only those expressible in terms of formal logic plus the vocabulary of science. The rules of inference in scientific discourse consist only of the rules of inference of logic and math plus scientific laws.

The Logical Positivist's view of what a theory is has since been deemed overly formalized. There are numerous legitimate theories in science that can't be rendered in a formal system. Consider theories in anthropology or geology for instance. Nevertheless, the idea of a theory as a formal system is a powerful one and it remains the gold standard in many sciences. Linguistics has "gone computational" in recent years, for instance. The most ambitious scientific undertaking in all of human history, the science of climate change, also aims to render theory and explanation in formal systems through massive and intricately detailed computer models of climate change. In fact, roughly speaking, we can consider a theory formalizable when it can be comprehensively modeled on a computer. Computer programs are paradigm examples of formal systems.

A further more general lesson we might take from the Positivist's view of theories addresses a very commonplace misunderstanding of what a theory is. People commonly think of theories as just claims that lie on a scale of certainty being somewhat more certain than guesses or hypotheses, but rather less certain than established matters of fact. This is really a terrible misunderstanding of what a theory is. It is commonly invoked in fallacious attempts to discredit science, as when people dismiss evolution or climate change science as "just a theory." Such comments reveal a basic misunderstanding of what theory is. For something to count as a theory has nothing to do with our level of certainty in its truth. Many scientific theories are among the best established scientific knowledge we have. A few years ago, for instance, some scientist claimed to have observed a particle in a particle accelerator travelling faster than the speed of light. It made the news and caused a bit of excitement. But those in the know, those who understand Einstein's special relativity and the full weight of the evidence in support of it patiently waited for the inevitable revelation that some clocks had been mis-calibrated. Einstein's special relativity is right and we know this with about as much certainty as we can know

anything. In the other direction, there are lots of genuine theories that we know full well to be false. Aristotle's physics would be one example. Having very much or very little confidence in something has nothing to do with whether it is properly called a theory.

So if it's not about our degree of confidence, what does make something a theory? What makes something a theory is that it provides a general framework for explaining things. The Positivists didn't discover this, but their idea of a theory as a formal system illustrates the idea nicely. Theories generally consist of a number of logically interconnected principles that can be mutually employed to explain and predict a range of observable phenomenon. Bear this in mind as we consider the Positivist's view of scientific explanation.

Explanation

According to the Deductive Nomological model of explanation developed by the Logical Positivist, Carl Hempel, a scientific explanation has the form of a deductively valid argument. The difference between an argument and an explanation is just their respective purposes. Formally, arguments and explanations look alike. But the purpose of an explanation is to shed light on something we accept as true, while the purpose of an argument is to give us a reason for thinking something is true. Given this difference in purpose, we call the claim that occupies the place of the conclusion the explanandum (it's the fact to be explained), and the claims that occupy the place of the premises the explanans (these are the claims that, taken together, provide the explanation). In a scientific explanation, the explanans will consist of laws and factual claims. The factual claims in conjunction with the laws will deductively entail the explanandum. For example, consider this explanation for why a rock falls to the earth:

1. $F = GM_1M_2/r^2$, Newton's law of universal gravitation which tells us that massive bodies experience a force of mutual attraction that is proportionate to their mass and inversely proportionate to the distance between them.
2. $F=MA$. This is the force law, which tells us that force equals mass times acceleration.
3. The rock has mass of 1 Kg.
4. The earth has a mass of 5.97219×10^{24} kilograms.
5. The rock was released within the gravitational field of the earth.
6. No forces prevented the rock from falling to the earth.
7. The rock fell to the earth.

Recall that deductive logic is part of every theory, every explanatory framework. The first two claims in this explanation are statements of law from Newtonian physics. The remaining four are statements of fact. Taken together, these six claims deductively entail the explanandum, that the rock fell to the earth. This should illustrate how theories function as explanatory frameworks.

One very useful thing Hempel's account of explanation does is alert us to the argument-like structure of developed explanations. The basic idea here is that a complete explanation should

include all of the facts involved in making the fact to be explained true. These will include both particular facts relevant to the specific fact we want explained and general principles (scientific laws in the case of scientific explanations) that belong to a broader framework for explanation. A fully developed explanation reveals a logical relationship between the fact we want to explain, other relevant facts and connecting principles like laws of nature.

Hempel's account of explanation faced a number of problems that have helped to refine our understanding of scientific explanation. We won't address them here except to mention one because it's amusing. Consider this explanation:

1. Men who take birth control pills do not get pregnant.
2. Bruce is a man and he takes birth control pills.
3. Bruce is not pregnant.

This seems to meet all of the positivist's criteria for being an explanation. But aside from being silly, it's at least not a very good explanation for why Bruce is not pregnant. Problem cases like this suggest that purely formal accounts of explanation like Hempel's will fall short in sorting which facts are relevant in an explanation.

There is also a more general lesson I'd like you to take from the positivist's account of explanation. For your entire career as a student you've been asked to explain things, but odds are nobody has ever really explained what it means to explain something. Personally, I don't think I had ever given a thought to what an explanation was until I encountered the Deductive Nomological account in a Philosophy of Science class. But now you've been introduced to a model of explanation. You may not find it fully applicable to every academic situation you encounter. But if you try to make use of it by thinking of explanations as having a developed argument like structure, you might find grades in many of your classes improving significantly.

We mentioned at the outset that Logical Positivism was very much influenced by Hume's Empiricism. You will recall that Hume argued for some surprising skeptical results. The Logical Positivists adopted one of two strategies for dealing with this. On some issues it was argued that Hume's skeptical conclusions were acceptable, while on others Hume's skepticism was regarded as a problem yet to be solved. As an example of the first strategy, Bertrand Russell, though not a Logical Positivist himself, wrote an influential paper in which he argued that science can proceed as usual without any reference to the notion of causation. Skepticism about necessary causal connections was deemed not to be problematic. Skepticism about induction was more difficult to accept. So the early 20th century saw a variety of sometimes colorful but generally unsuccessful attempts to resolve the problem of induction. And this brings us to Karl Popper.

Karl Popper

Karl Popper was a philosopher in Vienna during the reign of Logical Positivism, but he was not himself a Positivist. Popper is best known for his contributions to the problem of induction and the demarcation problem. In both cases his views were critical of the Logical Positivists.

Conjecture and Refutation

As you will recall, Hume argues that inductive arguments fail to provide rational support for their conclusions. His reason for taking induction to be irrational is that every inductive argument assumes that unobserved events will follow the pattern of observed events and this assumption cannot be supported either deductively or inductively. No purely deductive support can be given for this principle of induction because it is not a mere truth of logic. And any inductive argument offered in support of the inductive principle that unobserved cases will be like observed cases will be circular because it will also employ the very principle of induction it tries to support as a premise.

Popper accepted Hume's conclusion that inductive inference is not rationally justifiable. He takes the problem of induction to have no adequate solution. But he rejects the further conclusion that science therefore yields no knowledge of the nature of the world. With Hume, Popper holds that no number of cases offered as confirmation of a scientific hypothesis yields knowledge of the truth of that hypothesis. But just one observation that disagrees with a hypothesis can refute that hypothesis. So while empirical inquiry cannot provide knowledge of the truth of hypotheses through induction, it can provide knowledge of the falsity of hypotheses through deduction.

In place of induction, Popper offers the method of conjecture and refutation. Scientific hypotheses are offered as bold conjectures (guesses) about the nature of the world. In testing these conjectures through empirical experiment, we cannot give positive inductive reasons for thinking that they are true. But we can give reasons for thinking they are false. To see how this works, let's look at the pattern of reasoning employed in testing a scientific hypothesis using induction on the one hand, and Popper's deductive method of conjecture and refutation on the other. First, in designing an experiment, we determine what we should expect to observe if the hypothesis is true. Using induction, if our observation agrees with our expectation, we take the hypothesis to be inductively confirmed. The pattern of reasoning looks like this:

1. If H, then O
2. O
3. Therefore, H

This pattern of reasoning is not deductively valid (generate a counterexample to see for yourself), and as an inductive argument it faces the problem of induction. So this pattern of reasoning fails to provide us with rational grounds for accepting H as true. But suppose that when we carry out our experiment, we observe "not O." In this case our pattern of reasoning looks like this:

1. If H, then O

2. not O
3. Therefore, not H

This pattern of reasoning is deductively valid. To see this try to suppose that the premises are true and the conclusion is false. If the conclusion were false, then 'H' would be true. And, given this and the truth of the first premise, 'O' would follow. But 'O' contradicts 'not O' which is asserted by the second premise. So it is not possible for the premises to be true and the conclusion false. In other words, the pattern of reasoning here is deductively valid.

The latter is the pattern of reasoning used in the method of conjecture and refutation. It is a deductively valid pattern that makes no use of inductive confirmation. It should now be clear how Popper's method of conjecture and refutation works and how empirical inquiry making use of this method can provide us with knowledge of the world (or rather, how the world isn't) while avoiding the problem of induction.

According to Popper, there is no rational methodology or logic for evaluating how scientists come up with hypotheses. They are just conjectures and no amount of evidence is capable of inductively confirming hypotheses in the sense of giving us positive reason for thinking our hypotheses are true. Evidence in agreement with a hypothesis never provides it with inductive confirmation. If all the evidence is in agreement with a hypothesis, we can say that it is "corroborated." To say that a hypothesis is corroborated is just to say that it has survived our best attempts at refutation. But contrary evidence can decisively refute hypotheses.

Demarcation through Falsifiability

The demarcation problem is the problem of distinguishing science from other things, from poetry to religion to obscure metaphysics. Popper offers an alternative to the Positivist's verificationist theory of meaning in addressing this problem. The Positivist's solution to the demarcation problem had the downside of denying that we can assert as true that it is wrong to torture innocent babies just for fun. Popper's view of the matter avoids this unsavory consequence.

Popper's method of conjecture and refutation suggests his criterion for distinguishing science from non-science. For it to be possible to refute a hypothesis requires that there be possible observations that would give us grounds for rejecting the hypothesis. We can only scientifically investigate hypotheses that take observational risks, those that are exposed to the possibility of being shown false through observation. That is, we can take a hypothesis to be scientific if and only if it is falsifiable. For a hypothesis to be falsifiable we must be able to specify possible observational conditions that would be grounds for rejecting the hypothesis as false. But this does not mean that that it will be proven false or that it can be shown to be false (either of these confusions would lead to the absurd view that a claim is only scientific if it is false). Let's look at some examples to make this clear.

Consider the hypothesis that all crows are black. We can specify observable conditions under which we would count this as false. Namely, seeing a white crow, or a green one. Being able to

specify the observational conditions under which we would reject this hypothesis doesn't mean that it false. Suppose the hypothesis is true. It is still a claim that takes risks in the face of observation because we know that some possible observations would refute it. So the hypothesis that all crows are black is falsifiable.

Now consider claims made by astrology. These are typically formulated in such a vague way that any eventuality could be interpreted as affirming the astrologer's predictions. If there are no possible observations that could refute astrology, then it is not scientific. Some astrologers might make specific and concrete predictions. These might get to claim that they are being scientific on Popper's view, but to the degree that astrologers do take risks of being refuted by observation, they have been refuted too often.

Political ideologies often fail to pass the falsifiability test. Popper was especially critical of Marxism which was very popular with the Viennese intellectuals he knew in his youth. Marxists seemed to have an explanation for everything. The inevitability of Marxist revolution was illustrated by its rising popularity in much of Europe. But if Americans, for instance, were not rebelling against their capitalist oppressors it was only because they had yet to see how alienating capitalism is. The conditions for revolution just weren't yet ripe. But they will be, says the confident Marxist. Popper's key insight was that a theory that can explain everything that might happen doesn't really explain anything. It is empty.

Today, Popper might make the same criticism of very different political ideologies. If free markets don't fix every problem, the libertarian can always complain that this is only because they have not been allowed to function freely enough. If government doesn't fix every problem, the big government liberal can always complain that big government hasn't been empowered enough (when we get around to political philosophy we will find reason to doubt that there are very many liberals that really fit this stereotype). Extreme views are only made plausible to their fans by elaborate schemes of excuses for why they don't work out as well as they should. Popper would say that in politics as in science, we need to try things where we can honestly examine the consequences and hold ourselves accountable when they don't go well by trying something else.

Auxiliary Hypotheses

Here we will describe an objection to Popper's method of conjecture and refutation that will set the stage for introducing the views of Thomas Kuhn. According to Popper, we make progress in science by refuting false conjectures. We never have inductive grounds for holding that proposed scientific hypotheses and explanations are true, but we can narrow in on the truth by eliminating the falsehoods. Our hypotheses lead us to expect certain observations. If we do not observe what we expect to observe, then we have non-inductive grounds for rejecting our hypothesis. Again, the pattern of reasoning followed in eliminating false hypotheses through scientific inquiry looks like this:

1. If H, then O

2. Not O
3. Therefore, not H

This is the deductively valid pattern of reasoning known as *modus tollens*. However, we rarely get to test hypotheses in isolation. Typically, our expectation of a given observation is based on the hypothesis we are interested in testing in conjunction with any number of background assumptions. These background assumptions are the auxiliary hypotheses. If we take into account the auxiliary hypotheses, the pattern of reasoning used in Popper's method of conjecture and refutation looks like this:

1. If H and AH, then O
2. Not O
3. Therefore, not H

But this argument pattern is not valid. The observation (not O) might indicate the falsity of one of the auxiliary hypotheses (AH) rather than the falsity of (H), the hypothesis we set out to test. What this tells us is that the implications of other than expected observations are always ambiguous. When our observations don't accord with our expectations it tells us that at least one of the assumptions or hypotheses that lead us to expect a given observation is false. It may be the hypothesis we set out to test, or it may be one of our auxiliary hypotheses. But unexpected observations don't tell us which is false.

Here's a nice example of auxiliary hypotheses at work in everyday reasoning. Our hypothesis is that Hare is faster than Tortoise. This hypothesis leads us to expect that Hare will win a race against Tortoise. But suppose that, contrary to our expectation, we observe Tortoise winning the race. The hypothesis that Hare is faster than Tortoise is not thereby falsified because of the presence of a number of auxiliary hypotheses. Among these auxiliary hypotheses are the following: (i) Hare did not stop in the middle of the race for a snack, (ii) Hare did not get run over while crossing the road, (iii) Hare did not get eaten by Coyote during the race, (iv) Hare did not get entangled in a philosophical discussion about the rationality of scientific methods with his friend Gopher before crossing the finish line. When Tortoise crosses the finish line first, that tells us that either Tortoise is faster than Hare or one of these or many other auxiliary hypotheses is false. But Tortoise winning doesn't tell us which. The unexpected observation thus fails to cleanly refute our hypothesis.

Thomas Kuhn

The Positivists and Karl Popper offer attempts to describe and develop rational methods for scientific inquiry. In so doing, they offer *normative* theories of scientific practice. That is, they offer views about how scientific inquiry should proceed and what counts as good scientific practice. Kuhn's philosophy of science is inspired by the history of science and seeks to describe how science actually develops. Kuhn's undertaking is not aimed at revealing universal norms of

rational scientific practice. But his views have been taken by some to imply that the development of science is not guided by general norms of rationality, at least at crucial revolutionary periods of theory change.

Kuhn describes three stages in the development of a science. The first stage is called “pre-paradigm science.” In pre-paradigm science, people seeking to understand an observed phenomenon share no common stock of background theory. Each inquirer essentially starts from scratch. Under these circumstances, very little progress is made. We have nothing resembling a tradition that can be passed from one person on to her students for further development and investigation. The various theories of the nature of the world proposed by pre-Socratic philosophers might be considered an example of pre-paradigm physics.

At some point, someone develops an account of the observed phenomenon that has enough substance and explanatory power to attract the attention of a community of individuals who will then carry on inquiry along the proposed lines. This marks the beginnings of normal science. Kuhn calls the sort of account of the observed phenomenon that is required for this to happen a paradigm.

A paradigm consists of the following four things:

1. A body of theory including laws: For instance, the basic laws of motion.
2. Background metaphysical assumptions: For instance, that there is an external world and that our senses provide a reasonably reliable guide to its nature, that we share common objects of perception, etc.
3. Values: Here we have in mind primarily epistemological values including norms of rationality. The idea here is that a paradigm tells you what counts as a phenomenon that requires explanation and provides a standard for what counts as an adequate explanation of that phenomenon.
4. Exemplars: These are textbook applications of the theory to the phenomenon it is intended to explain. Classical physics is taught through exemplars that include applying Newton’s laws to swinging pendulums and forces exerted on springs.

Normal science, the second of Kuhn’s three stages, is carried out within a paradigm. Working within a paradigm, the scientist normally accepts the core elements of the paradigm as dogma. The scientist’s job in the stage of normal science is to work out the details of the paradigm without calling into question the central laws of the paradigm, or the epistemic standards it presupposes. In the normal stage, we can think of science as puzzle solving. Investigators are not advancing bold new theories, but applying the accepted theoretical framework in new and novel sorts of cases. During normal science, a paradigm gets worked out in detail.

In the course of normal science, problems that resist resolution with the paradigm often arise. If these “recalcitrant” problems remain long enough, they become what Kuhn calls *anomalies*. As the details of a paradigm get worked out, the anomalies become harder and harder to ignore.


Researchers in need of projects may focus more and more scrutiny on the remaining anomalies. Continued and intensified but unsuccessful attempts to resolve anomalies can give rise to a crisis in normal science. Such a crisis makes it possible to call into question core elements of the paradigm that had been previously held dogmatically.

Persistent anomalies in a science can provoke a crisis in which the paradigm itself is called into question. In this atmosphere, it is possible for scientists to propose and win wide acceptance for significant changes in the theoretical framework. Until persistent anomalies provide a crisis, however, the social conditions aren't ripe for revolution. Even if someone had great revolutionary ideas, they simply won't get a hearing with the community since it is committed to working out the details of the standing paradigm. Revolutions in thinking can't happen until the community is convinced that the old paradigm is irrevocably broken. When this does happen and an appropriate alternative to the old paradigm is developed and proposed, then and only then can what Kuhn calls a scientific revolution happen. In a scientific revolution, the scientific community abandons one paradigm in favor of another.

Once a new paradigm takes hold in the scientific community, normal science is resumed, the details of the new paradigm begin to get worked out and normal science continues until a new batch of anomalies emerges and provokes the next crisis.

A key insight of Kuhn's is that science is a community effort. We often hold a "great genius" vision of the history of science where the fabulous insights of very special individuals are what drive science forward. Kuhn would say this is a distorted picture. The great geniuses like Newton or Einstein can only launch a revolution in scientific thinking when a broader community of inquirers have prepared the field and created the conditions for the germination of the seeds of a revolution in thinking. The history of science needs to be understood in terms how these broader communities progress to the point where revolutionary thinking is called for and can be fruitful. The great insights and discoveries never happen in a social vacuum.

Kuhn thinks that the paradigm shift that occurs in the course of a scientific revolution is comparable to a gestalt switch as in the duck/rabbit image below.



Seeing this image as a duck blocks out seeing it as a rabbit. Something similar happens in the case of a paradigm shift. In a paradigm shift one drops one conceptual framework in favor of

another. When we grasp and evaluate the claims made in normal science, we do so in the context of acceptance of a paradigm. Kuhn suggests that the very meaning of the claims made in paradigm-based normal science can only be comprehended relative to the conceptual framework of that paradigm. A result of this is that from the perspective on one paradigm, we are never really in a position to evaluate the claims of normal science under a different paradigm. In this sense, paradigms are said to be incommensurable (lacking any common measure or independent standard of comparison).

It is tempting to see the cycle of normal science and revolutionary science as a Popper style process of conjecture and refutation at the level of paradigms. However, Kuhn maintains that paradigms are never exactly refuted by intractable anomalies. Rather, when the scientific community enters a period of crisis and an attractive alternative to the old paradigm emerges, the community gives up on the old paradigm and adopts the new one. Paradigms are not so much refuted as abandoned. This raises serious questions about whether paradigm shifts in scientific revolutions can be understood as rational processes. They would seem not to be if we think of human rational processes as in some way rule driven like logical rules of inference. But we might instead take Kuhn to be revealing a richer view of human rationality.

On Kuhn's view, the methods and standards of science get articulated and refined through periods of normal science and are liable to undergo bigger shifts in periods of scientific revolution. What counts as good scientific inquiry and investigation cannot be specified independent of its history. We figure out what works as we encounter new challenges. The history of science reveals the practice of science to be dynamic and adaptive. Creativity and resourcefulness go into the hard-earned advances in our understanding of the world.

The broader moral of this story is that we should be highly suspicious of any attempt to boil the methods of science down to any specific series of steps. Rather, a good understanding of the many methods of science can only be had through a study of its history, its successes, and its failures. And even at this, our appreciation of the methods of science must remain open ended. The story of science is far from finished, and so our understanding of its methods is likewise incomplete.

Review and Discussion Questions

1. How does the development of more powerful symbolic systems of logic boost Empiricism at the beginning of the 20th century?
2. Explain how the Logical Positivists extend Empiricism to the theory of meaning.
3. How is the verificationist theory of meaning used to address the demarcation problem?
4. What is a theory according to the Positivists?
5. Explain what's wrong with the view that theories are just very well supported hypotheses that are still not so certain.
6. What does it mean to regard a theory as an explanatory framework?

7. How do the Logical Positivists understand explanation?
8. How would Popper resolve (not solve) the problem of induction?
9. How does Popper address the demarcation problem?
10. Explain how auxiliary hypotheses challenge Popper's method of conjecture and refutation.
11. Explain pre-paradigm science. Why is little lasting progress made at this stage of science?
12. What is a paradigm?
13. How does normal science under a paradigm proceed?
14. What is an anomaly?
15. What conditions are necessary for a scientific revolution?
16. What does it mean to speak of competing paradigms as incommensurable?
17. How are the methods of science sensitive to its history?

8. Philosophy of Mind

The philosophy of mind has seen tremendous progress since Descartes proposed his dualist view of mind and body. Contemporary philosophical analyses of mental states and processes are among the key components of a rapidly emerging new science of mind, cognitive science. Philosophers of mind, along with psychologists, information scientists, and neuro-scientists have begun to work out detailed explanations of how our physical brains realize and carry out the functions of many mental states. In this chapter we will cover some of the progress philosophy of mind has contributed over the past century. As we will see by the end, some challenging philosophical questions about the nature of mind persist.

Descartes' Dualism

As with many topics, modern philosophy of mind begins with Descartes and soon moves on. Descartes' dualism holds that the mind is composed of a fundamentally different kind of substance than the body. Bodies are composed of matter that exists in space and time, and behaves in accordance with laws of nature. Minds, however, are spiritual in nature according to Descartes. Their existence is not spatio-temporally bound, and unlike physical stuff, minds have free will. The critical faults in Descartes' view were quickly spotted by Princess Elisabeth of Bohemia. The central problem lies in accounting for how the mind and the body can have any influence on each other. Clearly the physical world has effects on the mind, as when I perceive things. And it seems equally obvious that the mind has effects in the physical world, as when I act on my will. But if mind and body are so completely different, it is hard to see how this can happen at all. How does something that exists outside of space and time have any influence over the body that exists in space and time? How can the behavior of my causally determined body be influenced by a freely willing mind? The problem of mind/body interaction is a major stumbling block for Cartesian dualism about the mind and the body.

Descartes' correspondence with Princess Elisabeth can be found in a heavily edited form here: <https://www.earlymoderntexts.com/authors/descartes> The full original text of Descartes' correspondence with Elisabeth includes some ornate literary flirtation in the first couple letters. By today's standards, Descartes could be accused of sexual harassment for much of his flattery of Elisabeth. But the patronizing flattery pretty much disappears in later letters as Descartes comes to appreciate the seriousness of the Elisabeth's incisively argued challenge. Elisabeth provides a brilliant illustration of how to deal most effectively with patronizing behavior whether it is of sexist variety or some other kind: be competent and this will show that you deserve to be taken seriously.

Empiricism leads to Logical Behaviorism

Some of the main questions for the philosophy of mind are metaphysical questions about the nature of minds and mental states. “What is the mind?” quickly proves to be too big a question. We might say that for a being to have a mind is just for it to have mental states like beliefs, desires, perceptions, memories, emotions, and so forth. And this leads us towards somewhat more tractable questions like “What is a belief (desire, memory, perception. . .)?” This is the helpful turn taken with Logical Behaviorism.

Following Hume in the 18th century, the philosophy of science takes a sharp empirical turn in the latter 19th and early 20th century. During this time, what is scientifically knowable is taken to be limited to what can be defined in observable terms. This puts the mind and psychological phenomena generally on epistemically shaky ground. Mental states like beliefs, desires, perceptions, and anxieties are not the sorts of things we can examine under a microscope. If all things knowable are supposed to be knowable through sense experience, then it begins to appear that minds and mental states are not knowable.

The philosophical behaviorism of Gilbert Ryle is an attempt to salvage talk of minds and mental states and make such talk empirically acceptable. Mental terms like belief or fear can often be associated with observable behavior. Anger and fear, for instance, often seem to be observable. Suppose we identified the mental state of being angry with displaying angry behavior. On this proposal, anger just is stomping around, cursing a lot or generally throwing a fit. The obvious problem here is that some people can be angry without displaying it and some people, good actors for instance, can engage in convincingly angry behavior even though they aren’t really angry. Or to take another example, my desire for chocolate ice cream might be observable in my rummaging around in the freezer, or it might not be observable at all because the usual behaviors are checked by my (also unobservable) desire to shed a few pounds. So mental states like anger or many beliefs and desires sometimes show in terms of behaviors, but perhaps only under the right sorts of conditions. To make mental states empirically respectable and yet avoid the obvious problems we’ve seen in identifying mental states with observable behaviors, Ryle proposed to analyze mental states as **dispositions** to behave.

We are disposed in one way or another when we would behave a certain way given certain conditions. The behavior is not the disposition itself, but a **manifestation** of the disposition. The **disposition** can be identified in terms of a certain kind of “if. . .then. . .” statement. To help get clear on the idea, consider simpler physical dispositions like solubility or flexibility. To say that a spring is flexible is not to say that it is currently flexed. It is rather to say that if you were to stress it in the right way, then it would absorb the force placed on it and bend. To say that sugar is soluble is not to say that it is dissolved. But it is to say that if you were to submerge it in water (under the right conditions), then it would go into solution. So dispositions are described in terms of stimulating conditions and responses or manifestations. Ryle’s idea is that talk of mental states, like beliefs, desires, perceptions, or emotions can be fully explained as talk of very complex dispositions where the stimulating conditions and the manifestations are observable conditions and behaviors. So, my desire for chocolate ice cream might be understood as a

complex disposition to exhibit behaviors like rummaging around in the freezer if I think I'll find chocolate ice cream there, and I'm not too worried about my weight, and If this project works out, then we can understand talk of mental states in terms of empirically respectable stimulus-response dispositions.

The project of defining talk of mental states in terms of observable environmental stimuli and behavioral responses faces a number of difficult challenges though. We normally understand simple physical dispositions as being grounded in some further physical basis. Sugar is soluble in water because of its molecular structure, for instance. For sugar to be soluble is for it to have this underlying chemical structure that causes it to go into solution when submerged in water. In his eagerness to avoid positing unobservable mental states, Ryle wanted us to understand talk of dispositions merely as defining mental terms in terms of empirically respectable stimulus response "if . . then . . ." claims. He wanted to avoid positing any unobservable states of the brain, for instance, as the basis of mental dispositions. So Ryle's talk of dispositions is limited to mere "if . . then . . ." statements without any appeal to underlying states of the brain.

A second problem is that while we might be able to formulate plausible stimulus response conditionals for some mental state terms like fear or anger, in many cases the subtle links between stimulus and response that we might associate with a belief or a desire are simply too complex to allow for an analysis of the mental state talk in terms of observably defined disposition talk. What "if . . then . . ." claim, for instance, analyzes talk of my belief that my brother lives in Arizona?

A distinct problem, one that will continue to dog subsequent theories in the philosophy of mind, is the problem of conscious experience. However Ryle's project works out, we could imagine some kind of mindless robot that satisfies all of the relevant stimulus-response dispositions we associate with beliefs, desires, and emotions. And yet the robot has no mind and so no conscious experience. When we think of our own case at least, our subjective conscious experience seems to be quite central to having a mind. This is an issue we will return to after considering a few other 20th century approaches to understanding the mind.

Here is a link to the IEP entry on Behaviorism and some related movements in the Philosophy of Mind including material eliminativism: <http://www.iep.utm.edu/behavior/>

The Brain State Identity Theory

Ryle's behaviorism attempts to make talk of mental states empirically respectable by defining mental terms in terms of observable conditions and behaviors. One concern raised about this approach was that mental state terms are to be understood entirely in terms of observable things going on outside the person. This seems to take the mind out of the person. There is no place in behaviorism for any account of our inner lives or even the notion that my beliefs and desires are

in some sense in me or part of me. The Brain State Identity Theory, most ably advanced by J. J. C. Smart, goes some ways towards remedying this defect (though Ryle would not have counted it as a defect). The Brain State Identity Theory proposes that mental states are identical with brain states. Contrary to Descartes' dualism, the Identity Theory takes mind to be a physical thing. Namely, it takes the mind to be identical with the brain. For this reason the Identity Theory as a **physicalist** view of the mind.

According to the Identity theory, the belief that the state of Washington is in the Pacific Northwest just is a certain neuro-chemical state of the brain. Note that a great many people share this belief. When we speak of the belief as a view about what is true, one that might be shared by many people, we are speaking of a belief **type**. My belief that the state of Washington is in the Pacific Northwest is just one **token** of that shared belief type. This distinction between types and tokens is important to understanding what the identity theory says. The Identity Theory holds that mental state types are identical with brain state types. So, the identity theory holds that for anyone to have the belief that the state of Washington is in the Pacific Northwest is just for them to have that same specific neuro-chemical property. A plausible example of such mental state/brain state type identity is that pain just is C-fibers, a certain kind of neuron, firing.

We have scientific evidence that very roughly points in the direction of something like the Identity Theory. Cases of localized brain injuries indicate that different parts of the brain carry out different functions. People who suffer lesions in specific areas of the brain tend to find specific mental functions impaired while other functions are left perfectly intact. It is through analyzing such cases that we began to map areas of the brain according to the functions they perform.

In the Identity Theory we have a significant point of intersection between the philosophy of mind and the science of mind. Philosophical speculation has given rise to a great many scientific hypotheses. Here we have an example of how this can happen. We have a theory about the nature of mental states that turns out to be empirically testable. The Identity Theory says that mental state types are identical with brain state types. Types are properties, so this view tells us that all of your mental properties are physical properties of your brain. We have learned a great deal about how brains store and process information since this hypothesis was popular. The science of mind is not yet mature, but it is well past its infancy and the broad outlines of how brains work are more or less in place. What the science tells us is that different brains store and process the same information in very different ways. That is, the Identity Theory is wrong. My belief that the state of Washington is in the Pacific Northwest involves many properties of my brain. But your belief that the state of Washington is in the Pacific Northwest involves your brain having different properties. So, in the case of the Identity Theory, we have philosophers of mind proposing a scientific hypothesis that subsequent evidence showed to be false. But the rejection of type-type identity points us towards functionalism.

Functionalism

Belief is the spring of action. We explain people's actions in terms of their mental states. People do what they do because of what they believe, desire, fear, hope for, and so forth. The behaviorists were on to something in thinking about mental states in terms of dispositions to behave. But recall that the behaviorists were looking for a way to analyze talk about mental states entirely in terms of observable things like behavior. They wanted to avoid talking about unobservable things going on in the head. Talk of dispositions for the behaviorists was not talk of underlying and possibly unobservable brain states that give rise to behavior. Rather it was merely talk of tendencies that might allow us to understand mental state terms as synonymous with complex "if. . . then. . ." conditional statements.

Functionalists would understand talk of dispositions more in line with the way we usually do when we take a disposition to be based on some underlying state, as when we take the springs flexibility to be based on the physical properties of the steel it is made of.

According to **functionalism**, to be in a mental state is to be in some underlying brain state that realizes a certain functional role. The functional role can be understood as a complex of dispositions. But, unlike in the case of the Identity Theory, we now take dispositions to be grounded in underlying physical states.

An important point to understand about functional roles is that they are **multiply realizable**. Note that dispositions like flexibility might be grounded in very different causal bases in different cases, like the molecular structure of steel in the case of a watch spring or the biology of muscles and tendons in the case of your calf. Likewise, functional roles can be realized by very different underlying physical mechanisms. Consider the case of mouse traps. For something to be a mouse trap is for it to function effectively in a way that traps mice. But this is not to say anything specific about the physical properties of a mouse trap. Many different sorts of things, things of different designs and materials, can perform the functional role of trapping mice.

Similarly, in the case of mental state as functional states, the functionalist allows that a variety of different underlying brain states might realize the same functional role and constitute the same type of mental state. So, we share a mental state when share the belief that Obama was president in 2012, but that same mental state might be realized by differing brain states in my brain and yours. This in fact is what the neuro science bears out. When we examine how brains process and store information, we find that different brains handle the same information in different ways.

Here is a link to further reading on functionalism: <http://www.iep.utm.edu/functionism/>

Consciousness and Property Dualism

Functionalism remains the reigning paradigm for analyzing mental states. But not all that is mental yields to functional analysis. In particular, subjective conscious experience resists analysis in terms of functional roles. You might recall that consciousness was among the defining characteristics of Descartes' immaterial minds. The difficulty of understanding conscious experience in terms of functional roles and physical processes leads David Chalmers to another kind of dualism: **property dualism**. On this view, though the world consists of just one kind of substance, matter, that substance has fundamentally different kinds of properties including those we can regard as purely physical, like mass, charge and so forth, and other kinds of properties, like consciousness, that are irreducibly mental. Let's start by thinking about how consciousness is special and especially difficult to analyze in terms of physical properties.

Functionalism gives us a promising approach for understanding some kinds of mental states in terms of physical states that fulfil functional roles. According to functionalism, for me to believe that my cat is sleeping on the sofa only requires that my brain be in some state that plays an appropriate functional role. I can't specify the functional role completely, but it might include walking softly when I go to refill my tea, not playing loud music on the stereo, saying "no" if my wife asks me if the cat is outside, etc. The state of my brain that fulfills this mental functional role might be one that can be entirely specified in physical terms. It is just the state of having certain connections between networks of neurons activated in certain ways. With enough neurophysiology, we could completely describe this brain state in terms of chemical and electrical properties. A great many kinds of mental phenomena might yield completely to such functional explanation in purely physical terms. Cognitive scientists have already made tremendous progress at understanding memory, shape recognition, belief, and desire in terms of functional roles that have purely physical bases. But then there is our subjective conscious experience, what it is like for me to perceive something, for instance, or how I experience desiring something, believing something, remembering something.

Consciousness does not yield to functional analysis in the same way. An interesting kind of thought experiment suggests that consciousness can't be understood in purely functional terms or in terms of physical properties and processes at all. First, we need to talk about zombies. The zombies we are familiar with from horror stories are easily recognizable. They walk in menacing dull-witted ways in spite of broken legs and open wounds. They are the reanimated dead. This is not at all like philosophical zombies, the beings that populate philosophical zombie thought experiments in the philosophy of mind. The idea of a philosophical zombie is the idea of being that functions exactly like a conscious person in every observable respect. The only difference between a philosophical zombie and a normal person is that the philosophical zombie lacks conscious experience. Imagine a physical duplicate of yourself, a doppelganger that is functionally indistinguishable from the actual you. It looks the same, acts the same, gives the same replies you would give to questions and the same responses to stimuli. It is just as subtly expressive as you in every conceivable way because it is functionally just like you. Your mother or your lover could never tell the difference. The only difference there is, is that the zombie lacks

the conscious experiences that you have. There is simply nothing it is like to be your zombie doppelganger.

There is philosophical debate about whether such a being is metaphysically possible. There don't appear to be any logical contradictions involved, but that may not settle the issue. However, if such a zombie is possible, this possibility would demonstrate something interesting. Since your zombie doppelganger is exactly like your conscious self in every physical and functional respect down to the atomic level, yet differs from you mentally because it lacks conscious experience, the mere possibility of such a being would show that whatever consciousness is, it can't be understood in terms of functioning or the kinds of physical biochemical properties that ground your functioning.

Chalmers thinks philosophical zombies are possible, so consciousness can't be understood purely in terms of physical properties or the functional processes they ground. He instead proposes that we understand some properties of minds, like consciousness, as fundamentally mental properties that are not reducible, even in principle, to physical properties. While no distinct kind of non-physical substance is proposed, Chalmers is offering a kind of dualism we now call property dualism. Property dualism in the philosophy of mind is the view that among the primitive most fundamental properties of our world, there are both basically physical properties and basically mental properties.

Here is David Chalmers' clear and accessible paper "Facing up to the Problem of Consciousness": <http://consc.net/papers/facing.pdf>

Here is the IEP reference article on consciousness: <http://www.iep.utm.edu/consciou/>

Here is a collection of SEP entries on Mind, largely edited by David Chalmers: <http://consc.net/guide.html>

9. Love and Happiness

In this chapter we will begin to study things that matter, things that are important. We have had brief passing encounters with ethical issues in prior chapters, but beginning with this chapter and for the remainder of this text we will be concerned with issues that are, at least broadly speaking, ethical. We will begin with the things that matter to us individually, the things we love. Of course, different people love different things, people, and activities, so our starting point has to do with things that are good in a highly subjective way. But after thinking some about the nature of love, we will turn our attention to the good life later in this chapter. It's tempting to think that happiness and the good life are, like love, highly subjective. But notice that we can love, prefer, and pursue things that are also quite self-destructive. Between love and happiness, it is quite possible for us to be at odds with ourselves. Indeed, this is the stuff of tragedy. So perhaps what will make us happy and lead to a flourishing life isn't so subjective after all.

In subsequent chapters we'll examine the nature of morality generally, some theories of morally good action, and finally social justice. It might be tempting to think we move from the more to the less subjective in this sequence of topics. But subjectivity and objectivity are not the organizing principle I have in mind in taking things in this order. Rather, our own sphere of concern, what matters to us or what we love can be more narrowly focused on our subjective desires or it can encompass a broader realm beginning with our own well-being and proceeding to concern for others, respect for persons generally, and ultimately concern for the various often nested communities we are part of, all the way from the homeowner's association up to the biosphere. Maturing as a person likewise involves moving beyond the narcissistic self-centered sphere of concern we have as infants and towards an appropriately broader sphere of concern. Our introduction to ethics is organized around the idea that developing towards moral maturity is partly a matter of successively expanding our sphere of personal concern.

Love

Love comes in many varieties. A few varieties of love identified in ancient Greece continue to provide useful points of orientation. The Greek terms for these are *Philia*, *eros* and *agape*. *Philia* is friendship (this word is also the root of "philosophy" literally translated as the love of wisdom). *Eros* refers to erotic love, and *agape* we are most familiar with through the Christian tradition as a kind of universal love for all people. *Agape* is the sort of love that God has for all people and it also provides the foundation for Christian ethical precepts.

The classic account of *Philia* comes from Aristotle who takes friendship to be a concern for the good of another for her own sake. In friendship we adopt the good of another as a good of our own. It's important that we understand this as expanding our sphere of concern beyond ourselves. Concern for another just because of some benefit she will bring to us is not genuine friendship. Cultivating a relationship with someone because you think it will improve your social

standing or help you land a job is not really love in the sense of friendship. This is the significance of having concern for another *for their own sake*. Given this view, we can see the cynical view that everyone is ultimately motivated only by narrow self-centered self-interest as entailing the non-existence of *Philia* or friendship. Cynicism seems a rather sad and lonely attitude.

Friendship is not, on Aristotle's view, opposed to self-interest. It is common to think that when we come to genuinely care for another we do so at the expense of self-interest. Love, on this popular view often involves a measure of self-sacrifice for the sake of another. But this popular view is at best a bad distortion of Aristotle's view of friendship. This is because Aristotle takes love in the sense of friendship to involve an expansion of our own sphere of concern to include the good of another. This is not the refocusing of it away from ourselves. Of course, there will be conflicting desires among friends. But among friends these aren't mere conflicts between their individual wills. Rather, when I love my friend in the sense of *Philia* and I want one thing while my friend has a competing desire, I experience this as an internal conflict of my own will, and perhaps my friend does, too. It might not be obvious to either of us which movie we should see on our night out together. But the question of my self-interest versus my friend's dissolves in our mutual concern for each other for his or her own sake. The salient issue becomes what movie *we* should see together.

We often suppose that loving another means feeling good about that person. But love is emotionally more complicated than that, and Aristotle's account of *Philia* sheds some light on this. A person who cares about you for your own sake will generally feel good about you when things are going well. But another emotional manifestation of whole-hearted love might include disappointment when if you make an irresponsible choice, or anxiety if you are at risk in some way. This makes perfect sense on Aristotle's account.

Loving things

If philosophy is genuinely a case of *Philia*, then we should be able to make sense out of talk of loving things other than persons. We do commonly talk of loving chocolate, loving this or that band, or loving our house. In most cases this probably shouldn't be taken literally. Love is not mere appreciation, preferring, or desiring. So, saying I love the new Spoon album isn't saying that I care for it for its own sake or that I have adopted its interests as my own. It doesn't have interests that I'm in any position to adopt as my own. So, talk of loving things other than persons is often merely metaphorical. But is it always metaphorical?

Philia requires that a thing have a good of its own that we can adopt as part of the good for ourselves. We don't ordinarily regard things other than persons and relatively sophisticated animals as having a good of their own. My computer has value only in that it is useful to me. We refer to this kind of value as *instrumental value*. This is the sort of value a thing has because it is instrumental to satisfying other ends that we have. Typically we deny that non-sentient things have any value beyond their usefulness to us. In this frame of mind we are prone to think the

young hot-rodder who takes his vintage Mustang to have a good of its own, a value that is intrinsic to it and not merely instrumental to him and its admirers, as suffering from a kind of delusion. But perhaps this doesn't tell us so much about whether non-sentient things can have the kind of value that can make them appropriate objects of love as it does about the shortcomings of our ordinary state of mind. It might be worth considering the matter from a more creative state of mind and pondering the relationship between an artist and his or her art.

As a listener I can admire or enjoy a Spoon song or a Rachmaninov piano concerto, but I'm not in much of a position to adopt the good of the work as a good of my own. However, Britt Daniel and his colleagues are in a position to adopt the good of a Spoon song as a good of their own. Rachmaninov could very well be concerned with the aesthetic quality of his piano concerto for its own sake. The cynic will snidely remark that artists only aim to please their audience for the sake of drawing praise and honor on themselves. But I think the cynic fails to understand the artist and the experience of creating art. Practicing artists are typically not too concerned with reviews and prizes while they are actively creating. The concert pianist invests most of her time and concern in playing well. The adulation of an audience may be icing on the cake, but what really matters to the serious artist is the art. The creative activity is not a mere means to some further end, but it's absorbing, even all consuming. The artist is concerned with playing well, dancing well, making a beautiful and functional building, cooking well. We might worry that this is all just so much self-indulgence. The cynic might claim that artists are just doing what they want to, and if they are really lucky, others might like it too. But again, this doesn't really do justice to the nature of creative activity. When I bake an apple pie, I might be inclined to do so this way or that, and I might be perfectly happy with the results. But I don't get to set the standard of apple pie goodness. If I'm to be serious about my baking, I have to learn from people who know. Then I have to practice a lot, recognize my mistakes, and learn from them. In this process I have to sublimate my own inclinations and aspire to standards of excellence that lie well beyond my self-interest narrowly conceived. Really baking well requires a kind of aspiration and devotion that goes beyond self-indulgence. In recognizing this we can see the potential for love in creative activity. Creative activity can involve expanding one's sphere of concern to include the goodness of some activity or product for its own sake and this is the essence of *Philia*. Art, I'd suggest, is distinguished in part by the loving devotion of the artist. Now, depending on your inclination, listen to some [Rachmaninov](#) or some [Spoon](#) and see if you get my point.

Self-Esteem

What if we apply Aristotle's classic treatment of *Philia* to ourselves? The result is just that to love ourselves is to adopt the good for ourselves as a good of our own. Broadly speaking, to love yourself is just to care about what is best for you. We haven't yet said much about what is best for you. But let's suppose for now that what is best for you is subjective at least in some aspects and can be understood in terms of what you love. If what's best for me is what's good for the things I love, then to love myself is just to love what I love. This is pretty much the view of self-love

advanced by Harry Frankfurt in his essay, “The Dear Self,” the last chapter of his book *The Reasons of Love*.

The idea that to love yourself is just to love what you love sounds kind of poetically appealing, but a significant worry about this application of Aristotle’s view of *Philia* to the self is that it also appears to be trivial. After all, how could you fail to love what you love? If loving what you love is all there is to self-esteem this would seem to make poor self-esteem logically impossible. This would rob the idea of self-esteem of meaningful content. Frankfurt appreciates this problem and addresses it in an interesting way. He argues that we can fail to love what we love by being half-hearted. Sometimes we are at odds with ourselves in ways that undermine our love for the things we love. To take an all-too-common example, many of us both love our health and at least like things that aren’t so good for our health. Our appetite for unhealthy food frustrates and undermines our love of being healthy. So, we are half-hearted and at odds with ourselves about the prospect of going to the gym and having tofu for dinner. To have low self-esteem on Frankfurt’s account just is to have a divided will that leaves us half-hearted about the things we love. So we might love our bodies, but not whole-heartedly when we hold ourselves to unrealistic standards of physical beauty. Too many of us love our lovers, but not whole-heartedly because we still wish they were somehow more ideal. Or we might love our work, but not whole-heartedly if we feel it is under-appreciated or if it has too much drudgery attached to it. On Frankfurt’s account, these are all examples of ways in which we might suffer from low self-esteem. To love yourself is nothing more than to love your friends and family, your community, your activities, and projects whole-heartedly. To love yourself is to wholeheartedly love what you love.

This might sound pretty good. Enough so that it can be easy to miss just how dramatic a departure Frankfurt’s account of self-love is from conventional popular wisdom. We are frequently told that we have to love ourselves before we can love others. And in this conventional wisdom, loving ourselves just means feeling good about ourselves or thinking we are perfectly fine the way we are. But this is the narcissistic approach to self-esteem, a self-referential approach that is continually and simultaneously perpetuated and exploited in our consumer culture. Pop psychology tells us that we can’t care for others until we care for ourselves and consumerism makes sure that we are never quite done taking care of ourselves. This view is so deeply ingrained in our culture that it can be hard to penetrate even with pretty clear and compelling argument. What Frankfurt is recommending, perhaps without enough fanfare, is that the popular way of thinking about self-esteem gets things backwards. Leading a meaningful life and loving yourself is a matter of whole-heartedly caring about other things. There is no reference to feeling good about yourself in Frankfurt’s account of self-love. If things go well, feeling good about yourself might be the result of whole-heartedly loving what you love. But trying to feel good about yourself is exactly the wrong starting place.

Erotic Love

Even the most subtle minds are often overly tempted by the lure of simplification. So, it is not too surprising to hear smart people speaking of erotic love as nothing more than friendship plus sex. This view has the attraction of reducing erotic love to just a special variety of *Philia*. But the world has seen plenty of lovers that for one reason or another can't or don't have sex. And we are familiar enough with the notion of "friends with benefits" that aren't cases of erotic love.

We may also struggle here with the unfortunate fact that the word "erotic" has acquired a seedy connotation over the past century or two and now often serves as a code word for "X Rated." This is somewhat worse than a distortion of the word's traditional meaning. Erotic love does involve desire, attachment, and passion that is focused on a person, but this is not exhausted by the desire for sex. It's not even clear that this kind of love entails desire for sex. So, it's probably best to try to examine erotic love on its own terms, first and then maybe somewhere down the line think about how it relates to *Philia* or friendship.

The classic work on erotic love is Plato's [*Symposium*](#). This dialogue is a literary masterpiece as well as an interesting philosophical discussion of enduring themes on love. Do we search for our ideal other half in love? Do we love for reasons? And if so, what of the individuals we love? Do they matter except for the qualities we find loveable in them? These issues remain relevant in the very active contemporary literature on erotic love.

Erotic love is traditionally thought of as the kind of love that involves passionate longing or desire. This would appear to make erotic love self-centered and this seems to be at odds with the idea of *Philia* where another is valued for his or her own sake. A developed account of *eros* might resolve this apparent tension. And some further reflection on passionate longing might motivate this. If erotic love is hopelessly mired in selfish desire, then we might deem it a bad thing, nothing more than a euphemism for lust. But this would entirely miss what many people seek and sometimes find in erotic loving relationships.

When we desire something, we generally have our reasons. There is something about it we appreciate. When we are attracted to and desire some person, it may be because of this person's wit, beauty, or some other quality we find charming. Socrates makes this point in Plato's *Symposium* and it becomes the first step towards a highly impersonal view of *eros*. We might love an individual for their beauty, but this is just a step towards loving beautiful people generally and ultimately to loving beauty itself. As Socrates sees it, this is all for the good as our attention and love is drawn ever closer to the most real and divine of things, the form of goodness itself. Attachment to a particular individual is not the proper aim of erotic love and may even be a hindrance. The view of erotic love voiced by Socrates in the *Symposium* becomes refocused on God in the thought of Augustine with the result that in some veins of the Christian tradition, proper erotic love becomes passionate devotion to God. When focused on a person, it is dismissed as mere sinful lust, a misguided *eros* focused on less than worthy beings.

Christianity aside, the Socratic conception of erotic love is much broader than personal love. Any passionate aspiration can fall under the scope of the erotic on this broad view. An artist's passionate devotion to creative activity might count as erotic even when it has nothing to do with sexuality per se.

Freud offers a kind of inversion of this view. All creative aspiration is erotic, but Freud sees erotic aspiration as essentially sexual. When our sexual longings get thwarted or repressed, they surface in other kinds of creative activity. So, Socrates would say that aspiration generally is erotic and not necessarily sexual. Freud would also say all aspiration is erotic and still indirectly sexual.

Socrates' view of erotic love in the *Symposium* is a highly intellectualized view that most people simply can't relate to. The contemporary literature on the erotic love has framed the problem with this impersonal view of erotic love in new ways. Robert Nozick, for instance, has pointed out that if erotic love for another is focused on qualities we find charming or desirable, then it should make sense for us to "trade up" whenever we find another individual who has those qualities to a higher degree or those qualities plus others we find charming. Indeed, in the shuffle of immature relationships we see this happen often enough. He dumps her for someone hotter, or she dumps him for someone cooler. And granted some adults never quite outgrow this behavior. But as erotic loving relationships go, we see something deficient in this "trading up" behavior. We are inclined to say that these are rather sad cases where some of the people involved don't really know how to love. And we are inclined to say this precisely because there is something superficial about loving the just the qualities we find attractive to the exclusion of the individual person that might have some of those qualities. So, insofar as passionate longing or desire is focused on qualities we find attractive or charming, we seem to be missing most of what we find valuable in loving relationships. Perhaps what we do prize is a mix of *Philia* and *eros*. But just saying this hardly solves the problems of erotic love, since it should be clear now how there is liable to be some tension between the two.

The Ideal Union

Nozick proposes a model of love as a kind of union. In Nozick's version of the union model lovers form a "we" which is a new and different kind of entity, something more than just the sum of two individuals. We might be on to Nozick's idea of a "we" when we think of lovers as couples. Being part of a couple changes how we relate to the rest of the world. The IRS now wants to hear from "us" every year. We now socialize with other couples as a couple. I might be known as her partner to some and she will be known as mine to others.

Nozick is hardly the first to think of erotic love as a kind of union. The first was probably, once again, Plato, who has Aristophanes offer a colorful telling of the myth of [the origins of love](#) at the outset of the *Symposium*. In this story people were once two-headed eight-limbed round beings who upstaged the gods in their joyful vitality. To instill a bit of humility, the gods split them in two, and since then, erotic love has been the attempt by us incomplete halves to find our

other half and rejoin, if only temporarily. Less mythical versions of a union model of erotic love have been articulated by several contemporary philosophers of love.

Critics of the union model often see a metaphor run amuck. To think of the couple as a new entity distinct from the individuals that form it obscures the underlying reality. In fact, lovers are autonomous individuals making their own decisions. My selfishness and my self-sacrifice remain relevant in the relationship, but they are impossible to conceptualize on the union model. Our individual tastes, desires, and transgressions get dissolved in a “we.” Love as a kind of union sounds appealing as an ideal, but it may shed only limited light on the nature of our relationships and attitudes, even when these are at their best.

However philosophical theories of love as a union work out, many of us seek partners we think will be ideal complements to ourselves. The dream of a “soul mate” has powerful appeal. In the grip of this vision, we often find ourselves projecting what we want to see on to others who probably don’t actually live up to our desires. This may be a pretty good description of infatuation. A classic literary expression of “the birth of love” is given by the French writer Stendhal (1822 *On Love*). Stendhal describes falling in love as a process of crystallization, referring to how a twig left in a salt mine for a period of time will be retrieved covered in salt crystals. Similarly, our perception of our beloved is laced with projections of our own imaginative desire. In infatuation, our imagination gets the best of us and presents a distorted picture of another. The prospects for disappointment are built into such high expectations. If you haven’t personally fallen victim to the cycle of infatuation, disillusionment, and heartbreak yourself, I’m sure you know others who have.

Perhaps the stumbling block of distorting imagination is just a practical problem and the quest for one’s soul mate can be redeemed if only we can get a clear picture of who really is ideal for us. But there are other stumbling blocks built into the quest for one’s soul mate. A problem inherent to this quest is that, except for the searching, it puts one in a totally passive position, expecting another to conform to one’s own needs and desires. There is a tendency towards narcissism in this view when it encourages us to limit our sphere of concern to our own desires. This passivity renders us vulnerable to dependency and disappointment.

Another problem with the ideal union vision of love is that people are not just packages of qualities and capacities. People are active, dynamic, malleable beings that have their own will, grow in their own way, and have their own experience of the world. The person with the qualities you like might not have them tomorrow. Or you might come to prefer different qualities. The person you admire has his or her own desires and will, and fixating on what you want in a partner will render you ill equipped to be responsive to the autonomy and agency of another person.

Our challenge at this point is to find a way of understanding erotic love that is not both selfish and self-defeating. As easy as it might be to fall victim to cynicism, we should find hope in the

many cases of people who do find mutually enriching loving relationships. It might help to turn our attention away from what we want and towards trying to understand how erotic love enriches the lives of lovers when it does.

Bestowal

Notice how lovers affect each other. A kindness from a lover isn't just pleasant, it can also improve the beloved. A sincere complement isn't just acknowledgement of something attractive or admirable in us, it amplifies that attractive or admirable quality. When we value something or someone, we generally appraise that thing or person positively. Further, doing so can make that thing or person more praiseworthy. Through valuing something we *bestow* value on it. The marketplace provides a simple illustration of this in a very straightforward sense. If lots of people value a house when it goes on the market, its market value increases. It will fetch a higher price as a result. We bestow objective market value through subjectively appraising things highly. Popularity and attraction works something like this in a superficial way. A person's popularity can be substantially boosted by a few people deeming him or her to be likable. But this is not the kind of bestowal of value that makes loving relationships so enriching to the lives of lovers.

The marketplace example is just supposed to illustrate how valuing something can make it more valuable. Popularity and attractiveness are fleeting things that are as liable to mask vices and insecurities as contribute to flourishing. Much more significant kinds of bestowal are at work in loving relationships. When people care about each other they have much greater impact on each other. And this is not just a matter of degree. We affect those we care about in very different ways. Valuing an admirable quality in someone we love is a way of cultivating that quality. The contemporary philosopher of love Irving Singer has made *bestowal of value* central to his treatment of erotic love. Singer's central idea is that through valuing another in a loving relationship we create value and bestow it on our beloved. Loving another is not just a feeling on this view, it's a creative activity. Loving another and being loved brings out the best in us and improves the quality of our lives. It might seem obvious enough when put this way. What Singer would underscore is that through loving, we create and bestow value. Of course, as most of us will recognize, the value we bestow on people we care about is not always positive. We can tear down our lovers by being overly critical or unkind. So, the first rule of being a good lover is to try to be charitable and kind. Love, when it goes well, is the cultivation of value and goodness in the person we love.

Happiness

Let's start with the idea of something mattering. First notice the difference between something mattering to us and something mattering for us. Almost anything could matter to someone. All it takes for something to matter to someone is for that person to be concerned with it. Mattering to or being important to someone is subjective. Stamp collecting might matter to one person but not another. Football matters to some people but not others. The difference lies entirely in what the various parties are concerned with, prefer, or value.

Mattering for is another matter. Eating well and getting exercise matter for your health whether you prefer to do these things or not. Participating in caring relationships matters for your psychological well-being and this is so even for introverted people who enjoy their solitude. The notion of mattering for is not entirely subjective. And what matters for you is not relative to you in the way that what matters to you is. But mattering for is relational in a different way. There is a sense in which the idea of something mattering for you is incomplete. Things matter for your health, for your psychological well-being, your happiness, your marriage, your career, your projects, or for the quality of your life. You don't get to just pick and choose which things matter for your health or for your psychological well-being. For these things at least, what matters for you is largely settled by what and who you are. It remains an open question whether what matters for our happiness is up to us or subjective in the way that what matters to us is subjective.

Do I get to pick and choose what matters for my happiness? Is what matters for my happiness a question of what matters to me? Assuming the good life is the happy life, the question we have before us now is whether or not happiness and the good life are subjective and relative to our values and preferences the way that what matters to us is. Popular opinion would seem to make short work of these questions and straightaway affirm the subjectivity of happiness and the good life. Surely different people enjoy different things depending on their preferences and values. And nobody gets to decide what I enjoy or prefer but me. So, concludes this line of argument, happiness for me and the good life for me are up to me.

Cultural norms of individualism and liberty probably feed a bias towards subjective ways of thinking about happiness and the good life. The American way of life at the beginning of the 21st century does seem to include an implicit view of the good life. We are indoctrinated into that view of the good life by advertising and media from a very early age, and our peers, similarly influenced, reinforce the programming. We are used to being referred to as consumers rather than citizens, or simply people. And consumerism might be as good a name as any for the philosophy of the good life that is standard issue in our culture. Consumerism as a philosophy of the good life tells us that what is good for us is just getting what we want. It's a seductive view. Who could possibly object? We all want what we want after all.

No matter how seductive the conventional wisdom is, if we want to address questions about happiness and the good life philosophically, we had better resist the urge to settle them by wishful thinking. A bit of critical thinking should get us past our culturally ingrained bias and at least suggest a more subtle and interesting appreciation of these issues. For starters, let's look at cases where people really do get everything they want and ask whether these make for plausible examples of happiness and the good life. The spoiled child comes to mind. The spoiled child, by definition, is the child that always gets what he or she wants. But spoiled children are typically not very pleasant or happy. Closer to home, we are all familiar with instances where we get what we wanted and then find that we aren't as pleased as we'd hoped. Even when we are well satisfied, we are generally not pleased for long. The next even more pressing want waits just

around the corner and we are dissatisfied again until we are briefly sated by its attainment. Getting what he or she wants doesn't seem to help the spoiled child either. Part of the problem might be that having your every wish indulged provokes insecurity. The spoiled child becomes completely and passively dependent on the parent who indulges, and the stakes are ever higher as the desires become more pressing. Clearly, we can get what we want and still not be happy. This should be a pretty clear indicator that what we want is not a perfectly reliable guide to what will make us happy.

Perhaps we have said enough to debunk consumerism as a plausible theory of the good life. But in doing so we ran roughshod over an issue that might be worth exploring. We often suffer from internal conflicts between two or more of the many things that matter to us. It might be that losing a few pounds and being physically fit matter to me. And yet when the dessert cart comes around, I give in to the temptation for chocolate. What I want at the moment might diverge from what really matters to me. Of course, if I want the chocolate cake, then there is a sense in which it matters to me as well. So, we are conflicted. Our wants also change. We get drawn in and, in our craving, we neglect other things that also matter to us. Our wants don't just conflict, they seem to jostle and vie for the privilege of commanding our will. So, when we choose among the various things we want, like losing some weight or enjoying a piece of cake, we might begin to wonder whether there is some wise rational executive function in our mind that can systematically bring our competing desires into line with each other. Perhaps there is, but the effectiveness of this rational deliberative function varies significantly from person to person, and from period to period in the lives of the same person. It seems that among the people we know some do better and some do worse at resisting the temptation of the moment and staying motivated by what matters to them most. Even in our own lives, most of us can identify times when we exercised self-control more effectively than others.

So, here's what we have so far: There can be conflicts among the things that matter to us and some things matter to us more than others. We can do a better or worse job of resolving the various conflicts in favor of the things that matter more to us. When we do well at this we have willed rationally and exercised self-control. When we fail, we fall victim to weakness of will. Based on this, it should be clear that self-control is a good thing. That is to say it's a virtue. Self-control empowers us to act most effectively on what matters most to us. We are now in position to articulate another view of the good life, one that still doesn't appeal to any external standards and makes the good life a function of what matters to us, but doesn't simply make the good life a question of whatever I want or choose. The good life on this model is one where we reflectively weigh the various things that matter to us in a way that makes it possible to resolve conflicts among them in favor of the things that matter more to us and then exercise the virtue of self-control in formulating the will to act in accordance with those things that matter to us most. On this view, consumerism takes a step in the right direction by looking to what matters to us, but then fails to articulate a model for resolving conflicting values and desires and misses the virtues of rational deliberation and self-control in adjudicating these.

The next big question should be how do we determine what matters most to us? How do we settle the conflicts among our competing desires? Is this simply a matter of our choice? If so, then the whole structure we just articulated might be at risk of collapsing. If what matters to us most can be read off of what we choose, then the distinction between exercising self-control and being weak-willed simply collapses. Suppose we say that if I choose the chocolate cake, that can only be because that's what matters most to me. If what matters to us is simply a question of what we choose, then there can be no such thing as weakness of will or self-control. So, formulating a plausible view of the good life seems to require that we somehow reach beyond our subjective preferences, but it is not yet clear just how. Should some things matter more to us than others? It begins to look like we need to recognize some substantive difference between what matters to us and what matters for us. But it is not yet clear how to do so.

One thing does seem clear on the question, however: we don't want to be told what is good for us. Settling on what is good for us shouldn't be a matter of acquiescence to some authority, whether it is our parents, some tyrant, or the tyranny of popular opinion. If we don't get to decide what matters for us, then it won't do to have anybody else deciding for us. The possibility that remains open is that determining what matters for us is not a matter of anybody deciding, but instead a matter of us figuring it out. If this suggestion is on the right track, then questions about happiness and the good life are not subjective, that is, they aren't matters for us or anybody else to just decide. Rather, they are objective in the way that scientific truths are. We have to investigate, discover, reason well, and figure out what is good for us. Enter Aristotle.

The *Nicomachean Ethics*

When we considered the consumerist conception of happiness and the good life, we spoke of identifying what would make us happy. Notice that this way of thinking about happiness puts us in a passive position. Something outside of us does something to us, and it makes us happy. All that is required of us is to be fortunate enough to be in a position to receive this wonderful benefit. By contrast, for Aristotle, happiness is active. Things external to us might help or hinder, but ultimately, for us to be happy just is for us to be active in the right sorts of ways.

Aristotle identifies leading the good life with being happy. But happiness in the sense he has in mind is not just feeling happy or being in a happy mood. Moods and feelings are things that come and go in our lives. They are temporary states of mind. Aristotle is not interested in moods so much as what it means to live well. So, we are after the idea of an excellent life. The Greek term Aristotle uses is *eudaimonia* and this might be best translated as living well and doing well. So when Aristotle identifies the good life with happiness, he has something more enduring and emblematic of a life in mind than just feeling good.

You might recall that Aristotle has a teleological view of the world. That is, everything has an end or a goal towards which it strives. He is inclined to understand the nature of things in terms of how they function in pursuing the ends towards which they are oriented. In this spirit, Aristotle would take goodness to be something we naturally aim at, something we are oriented

towards by nature. So, for Aristotle, the idea of the good life is understood in a naturalistic way. Aristotle conceives of ethics in a way that blends seamlessly into his broader paradigm for understanding the natural world. Goodness is an integral aspect of the natural world. What is good for a thing can be understood in terms of that thing realizing its *telos*.

The good life, conceived of as happiness in the broader enduring sense, is a goal or an end for a person's life. But it's an end of a particular kind in that it is sought for its own sake, not as a means to some further end. Aristotle refers to ends like this as final ends. In more contemporary language we might speak of things that are pursued for their *intrinsic value*, the value had "in itself" as opposed to things that are pursued for their *instrumental value*, their value in the sense of being useful as a means to other ends. Money, for instance, has instrumental value, but no intrinsic value. It's a useful instrument for attaining other things of value like clothes or food. And these also may have only instrumental value toward yet further goals. The value of clothing is to keep us comfortable and make us look good. But clothes don't have a value of their own independent of their usefulness towards these other ends.

The idea of things having instrumental value seems to presuppose that some things have value just for their own sake. Otherwise, we seem to have a regress of value where many things are valuable as means to further and further ends, but at no point do any of these ends have any value of their own. So, to make sense out of anything having any sort of value, it seems there would have to be some things that have intrinsic value or value in themselves. In the broadest sense, goodness is an end that has "to be pursued" built into it. Thus goodness, for the ancient Greeks, was a natural and obvious theoretical posit, needed to make sense out of any sort of talk of value. For humans, the kind of goodness that matters is the good life. So ethics in general is concerned with how to live well, how to lead an excellent life.

In the idea of flourishing, we have at least one familiar notion that should help us better understand how Aristotle sees the good life. Think about what it is for the vegetable plants in the garden to be flourishing. The flourishing tomato plant is one that grows vigorously without disease and is well on its way to achieving its natural end, growing lots of sweet ripe tomatoes. In line with his teleological view of the natural world, Aristotle has it that the good for any sort of thing can be understood in terms of fulfilling its natural function well and thereby realizing its *telos*. So what then is the unique function of humans in terms of which our essence can be understood?

It seemed we had a pretty good idea of what it means for a tomato plant to flourish. Roughly it is for it to take in nourishment and grow. Biologically we might say its function is to photosynthesize, converting nutrients and CO₂ into lots of sugar and oxygen (and, ultimately, tilth). But we are essentially different from plants, so our function must be different as well. Aristotle entertains the idea that our function might be to satisfy our appetites. This much seems in line with the consumerist idea of the good life. But Aristotle rejects this too since it fails to separate us from barnyard animals. Perhaps as infants we are similar to animals in functioning

only to satisfy our appetites, but then we outgrow this similarity. We might now see the consumerist conception of the good life as infantilizing us since it appeals only to how we function as infants, getting our appetites satisfied. But for Aristotle, how we function beyond the developmental stage of small children is important to understanding our *telos* as human beings. Ultimately Aristotle settles on our rational capacities. He takes the function of the human being to lie in exercising our rational capacities because these are the ways of functioning that are unique and special to humans. Humans are distinguished from others sorts of being by their ability to function rationally. For Aristotle, the human being essentially is the rational animal. The ability to reason is what sets us apart from other animals and this is what defines us.

Since for Aristotle, what's good for us is not something we get to choose for ourselves, his idea of the good life might seem much less flexible and personalizable than the consumerist conception of the good life. But the apparent flexibility of the consumerist conception might be just that, only apparent. On the consumerist conception of the good life, the preferences that fix what is good for you are just given. What we want is taken as the starting point for thinking about what is good for us. For this reason, the consumerist philosophy affords no means of critically evaluating our wants. We just want what we want; that's all there is to it. In our contemporary consumerist culture, any challenge to the aptness of our wants is received as grounds for offense, where our freedom to choose is compromised by someone else telling us what we should want. However, our desires are quite malleable. Our tastes are typically acquired and usually this happens without much critical reflection. Advertisers, and political pundits among others know this well. The most powerful institutions in our culture put immense and sophisticated effort into shaping and manipulating our desires. We are free to choose what we want as consumers, but only after our wants have been engineered with care by others that aren't really concerned about what's actually in our interest. In practice, the supposed freedom and flexibility of the consumerist conception of the good life is more illusion than reality.

On the other hand, Aristotle's view of the good life as the life of actively exercising one's rational capacities might be more flexible than it appears at first. Interpreted narrowly, Aristotle offers a highly intellectualized view of the good life. The good life is the life of the philosopher/scientist. It would not be unreasonable to suspect a bit of professional bias in Aristotle's idea of what it is to live well. We might object that some people would rather work in the garden, ride bicycles, or practice yoga than just do philosophy all the time, and that this is a good way to live too. A reply that Aristotle can offer here (the reply I think he ought to offer) is to say, "very well, and any of these activities will contribute to your flourishing only if you engage your rational capacities and do them in thoughtful and inquisitive ways." Many crafts, arts, and skills can be cultivated in ways that exercise and develop our rational capacities. A life spent working in the garden, riding bikes, doing yoga, or working as a plumber can be a flourishing life on this more liberal interpretation of Aristotle's account. There are details to work out here concerning just what ways of life will exercise and cultivate our uniquely human rational capacities. But more generally, perhaps I can understand the good life as the active life

of exercising and developing our uniquely human rational, capacities whatever specific endeavors and activities I ultimately identify as serving that end.

Once we have a thought-out idea of what the good life is, there remains the issue of how to go about leading such a life. Through critical reflection on our nature and capacities we might discover (as opposed decide or choose) where our genuine interests lie. But then how do we bring ourselves to act in our best interest. What if it turns out that we don't desire what is best for us? Are we then just fated for misery? Aristotle doesn't think so. There is a degree of flexibility in our inclinations and preferences and we have some ability to shape these over time. On the consumerist conception, what matters for us is set by our desires and the theory of the good life is made to conform to them. On Aristotle's view, the theory of the good life is developed according to what matters for us and this is set by the sort of being we are. So, living well is a matter of bringing our desires into line with our interests. If Aristotle's idea that we don't get to simply choose what is best for us still seems at all stifling, your sense of personal autonomy might be replenished in appreciating how we are empowered to shape our tastes and preferences and gradually bring them into line with what we can learn about our interests.

We are creatures of habit. While this often presents an obstacle to acting on our considered interests, habit is also the means available to us for shaping our lives for the better. Recognizing that making some a change would be good for us typically doesn't result in our immediately preferring it. Many of us, for instance, recognize that getting more exercise and eating better would be good for us. But thinking that more exercise would improve our lives doesn't automatically result in feeling the urge to go for a run. Habituation, however, can bring our preferences and urges into line. People that regularly go for runs like to go for runs. Good habits can be as addictive as bad ones. And once we establish a good habit, that becomes what we prefer and enjoy the most. For Aristotle, the power you have to shape your life for the better lies in your ability to shape your habits. Once we have figured out what really is in our best interest, the key to being happy and living well is to mold our inclinations, preferences, and pleasures through habituation. The good life, which is also the virtuous life, will be the most pleasant life because it is the life where our pleasures cohere rather than clash with our interests. The good life is one where we have resolved the conflicts in our inclinations and pleasures and we no longer have things that matter to us fighting against things that matter for us. The virtuous person can wholeheartedly pursue what pleases her most because this will be well aligned with what is best for her. The affinity between Aristotle's advice on how to live well and Frankfurt's account of self-esteem should be easy to see here. Both would say living well is largely a matter of getting your desires, inclinations, and motivations to hang together in a unified coherent way. Where Frankfurt and Aristotle will differ is just in how that unified will gets oriented. Frankfurt would have our considered best interest be determined by what we love. Aristotle sees our considered best interest as settled by our nature as rational animals.

For Aristotle, to be virtuous is to have habitually established inclinations and preferences for actively exercising our human rational capacities. Virtue aims at flourishing. Habit, on this view,

is quite literally character building. This way of thinking about virtue stands in sharp contrast to more popular conceptions where to be virtuous involves lots of sacrifice. We suffer from a Christianized notion of virtue that is more often than not associated with self-denial. To be virtuous in the popular sense means something like not overindulging in cheesecake or sex. But we are concerned with the idea of virtue as a kind of excellence. When Aristotle talks about virtue, he is just talking about the excellent character traits a person might have. What makes a character trait a good one is that, on the whole, it contributes better to a flourishing life than contrasting traits. So, life in accordance with virtue promotes human flourishing, and for this reason it is also likely to be the most pleasant.

Happiness, however, requires more than just virtue. It also requires some degree of good fortune. A person with a virtuous character who is also in a coma is not really flourishing. Likewise, a virtuous person who lives in a community of not so virtuous people faces a significant obstacle to flourishing. Living in community of fools might provide very limited opportunities for exercising one's rational capacities. There would be no one to talk philosophy with for starters. More seriously, disputes could not be settled reasonably, but only through vicious maneuvering for dominance. Extreme poverty can be an obstacle to flourishing. Being always anxious about where your next meal is coming from could make leading the active life of the rational element a difficult proposition. But extreme affluence and luxury could present its own obstacles since they offer endless distractions, draw your attention to trifles, and ultimately render you passive and weak. How much and what kind of good fortune does leading the good life require? Perhaps we can't give a very precise answer, but it might do to say that we require enough good fortune to give us ample opportunity to exercise our rational faculties.

Here is an excellent translation of Book 1 of Aristotle's *Nicomachean Ethics* as a PDF:
<http://catdir.loc.gov/catdir/samples/cam032/99036947.pdf>

Here is the complete *Nicomachean Ethics* in a good, but older translation:
<http://www.gutenberg.org/ebooks/8438>

Review and Discussion Questions

1. Explain Aristotle's view of *Philia*.
2. What is it to love one's self on Frankfurt's account?
3. How is love of the self the purest form of love? How does Frankfurt handle the two complications for his analysis of self-love he raises at the end of section 9?
4. How is self-love different from self-indulgence? What does Frankfurt mean by "wholeheartedness" and how does this bear on self-love.
5. How does Frankfurt see self-love as related to morality?
6. How does he see it as related to the meaningfulness of one's life? Do think he is right on these points?
7. How does Frankfurt's analysis of love as whole-heartedness square with Aristotle's view of virtue and happiness?

8. Explain the apparent tension between *eros* and *Philia*.
9. Explain the view of erotic love voiced by Socrates in the *Symposium*.
10. Explain Nozick's union model of love. What problems does it raise?
11. What problems are inherent in the idea that the search for love is a quest for one's ideal other half or "soul mate"?
12. How does the cycle of infatuation, disillusionment, and heartbreak work? Explain the role of the imagination in this.
13. What failings do the consumerist conception of the good life and the vision of love as a quest for one's "soul mate" share?
14. How do we bestow value in a loving relationship?
15. How do we understand what matters for us and what matters to us on the consumerist conception of the good life?
16. What sort of evidence can be brought against the consumerist conception of the good life?
17. How can rational deliberation and self-control help the subjectivist view of the good life?
18. What is weakness of will, and what does the possibility of weakness of will tell us about how we can determine what matter most to us?
19. What is the good life according to Aristotle?
20. Explain the ways in which the good life is passive on the consumerist conception and active on Aristotle's view.
21. Explain instrumental and intrinsic value.
22. On Aristotle's view, once we know what the good life is, how do we lead it?
23. How does Aristotle understand virtue?
24. What sorts of things can be aids or obstacles to our flourishing?

10. Meta-Ethics

Ethics is inquiry into what is good. Ethics isn't concerned with describing how things are, but with saying how things ought to be. That is, ethics isn't descriptive, it is normative. By contrast, when anthropologists examine the moral traditions of various cultures, they are generally aiming only to describe, explain and understand the ethical views of other cultures. They are not engaged in the normative project of trying to understand what moral code is best. In other words, they are out to understand **what is regarded** as good and bad by some culture or society. But they are not concerned with **what is** good and bad. Inquiry into what is good and bad is normative inquiry, and as soon as we engage in this sort of inquiry, we are doing normative ethics. We will take up normative ethical theories of right action in the next chapter. In this chapter we will be concerned with more fundamental issues about the nature of normative ethical claims.

This chapter will be organized around these two questions:

- Are there any moral truths?
- If so, what makes them true, what is the basis for these normative truths?

Suppose the answer to the first of these questions is no, there are no moral truths. Will refer to this view as **nihilism**. According to nihilism, there is no value in the world. This would appear to make our lives pointless and absurd. The Existentialists have a good deal to say about this possibility.

Or, perhaps there is value in the world, but it is merely subjective. **Moral subjectivism** is the view that there is only right and wrong relative to me, or you or some other subject. This is a form of individualized **moral relativism**. This view answers our first question in the affirmative and answer the second question concerning what makes moral truths true in terms of my say so for me, your say so for you, etc.

A further possible meta-ethical position is collectivized versions of moral relativism. It might be that good and bad are matters of social convention, like laws or standards of etiquette. Perhaps the most popular view of this sort is **cultural moral relativism (CMR)**, the view that what is right is right only relative to cultural groups and is determined by the values, standards and traditions of a culture.

The assorted varieties of moral relativism deny that there are any universal moral standards and make morality a matter of the say-so or authority of people, either individually or in groups. One further view that takes morality to be a matter of say so would be **Divine Command Theory (DCT)**. This is the view that what is right is right because God says so. In a way this makes morality objective because God's commands will apply to people universally. However, like varieties of relativism, DCT makes morality a matter of say so.

But might there be moral truths that aren't just a matter of say so. Moral relativism and DCT both make morality a matter of someone's say so. We'll call this view **Moral Realism**. We can readily point to things happening in the world to explain what makes our descriptive claims true. Consider the simple descriptive claim, "Marsha's house is green." We need only point out Marsha's house in normal daylight to see what makes this claim true. But what about claims about how things ought to be? I can't just point to something in the world that makes a normative ethical claim like "people shouldn't torture innocent puppies just for fun" true. Or can I? The puppies suffering might be a morally relevant fact in the world. At any rate, the moral realist owes us some answer to questions like these. Moral realism takes there to be moral truths and takes these to be made true by something other than the say so of people (or God). We should then want to know what that is.

We will take up each of these positions in turn in this chapter. In doing so, we are engaged in **meta-ethics**, inquiring into the status of morality, whether it is subjective or objective and how so. Our meta-ethical inquiry will set the stage for our final two chapters on the normative ethics of actions and social justice, the normative ethics of societies.

Nihilism

Nihilism is the view that there is no value in the world, nothing is good or bad, and we are just deluding ourselves in thinking that anything matters at all. On this view all of our ethical judgments are simply mistaken. Note that this is not the view that it's morally fine to do as we please. That too is an ethical judgment and is no more true or reasonable than any other value judgment according to nihilism. We may not be able to decisively refute nihilism, but that would be no argument in support of it. And building a case for nihilism presents some obstacles. In particular, the nihilist owes us an error theory, an account of how our moral experience and judgments so badly mislead us.

We are moral beings. Morality is a central part of our lived experience. Perhaps this is not the case for psychopaths and sociopaths, but this is just why we count these as pathologies. The psychopath and the sociopath are lacking in a central area of human functioning. That these individuals lack moral experience is no more reason to think the rest of us are deluded than to think that they are simply blind to something real the rest of us appreciate. The rest of us do have moral emotions. When we harm others, we feel guilty. And when we are harmed by others, we may feel indignant. We also have moral motivations. We are generally motivated to treat others fairly even when we can see some advantage in putting ourselves first. And we are motivated to condemn cruelty and praise kindness. Most of us fall short in assorted ways now and then, and we often recognize this ourselves whether someone else judges us or not. If all the moral aspects of our lived experience are mere delusion, then the nihilist owes us some account of how this experience deceives us so systematically.

In suggesting that our moral experience is just a delusion, the nihilist is offering a skeptical view of morality. This might remind you of the more general skeptical hypothesis of Descartes' evil deceiver. While we can't rule out the more general skeptical hypothesis that all of our experience is a deceptive illusion, we generally tag this as a philosophical curiosity and continue to confidently believe that our senses provide us with fallible, but basically reliable evidence of an external reality. Worries about Descartes' evil deceiver present no obstacle to scientific inquiry about the physical world and the benefits of this are pretty obvious even if it turns out that we are living in a simulation of some sort.

I'll suggest we treat nihilism in the same vein. Perhaps we can't rule out the possibility that our experiences as moral beings are mere illusions. Yet, all the same, we can better understand our ourselves and benefit in ways we care about by engaging in systematic inquiry into our lived moral experiences. Our lived experience as moral beings presents us with prima facie evidence for thinking that morality is real in some sense. The nihilist owes us some compelling argument for doubting our lived moral experience before we should grant his moral skepticism any more credence than Descartes's evil demon.

Is Morality Subjective?

Setting aside nihilism, we now turn the idea that morality is subjective. We first need to ask again just what we mean by "subjective." There is a broad sense in which your entire lived experience is subjective. It is, after all, your experience and you are a subject. But there are many aspects of our subjective experience that reflect an objective reality. My subjective visual experience informs me that my bike has a flat tire and my subjective tactile experience quickly confirms this. This subjective experience gives me a reason to believe that my bike tire is flat as an objective matter of fact. Setting aside skeptical curiosities like Descartes' evil deceiver, it seems that we can learn quite a bit about what is objectively true on basis of our subjective experience. But because this broad sense of what is subjective includes experiences that track objective reality, it also doesn't get at what is supposed to be interesting about the notion that morality is subjective.

Rather, the idea of subjectivity that we want to capture in connection with moral subjectivism is that morality itself is relative to you. Things we care about and things we prefer seem to work this way. You determine what you prefer or care about and no thing or person can do that for you. What's best relative to your preferences and concerns depends entirely on you and you may prefer things I don't prefer. The subjectivist takes morality to work like what we prefer or care about, not like subjective experience that tracks things that are objectively real or true. So, we'll call moral subjectivism the view that there is no objective goodness, but only what is subjectively good relative to individual subject.

Moral subjectivism is no mere psychological thesis, but an ethical thesis. If I regard the death penalty as morally justifiable in some cases and you don't, that is just a psychological

observation about our differing beliefs and attitudes. The possibility that one of us gets this moral issue wrong remains open. Moral subjectivism, however, is the view that what *is* right is right only relative to me or relative to another and what is right relative to me or another is determined by our respective opinions and attitudes. It is not just individual moral beliefs and attitudes that are relative to individuals according to the subjectivist, but morality itself that is relative.

Moral subjectivism is the limiting case of moral relativism, where morality itself is relativized to the individual. It has some obvious appeal. Everyone wants to think that they are good and subjectivism allows us all to set the standard of goodness for ourselves. Subjectivism appears to be highly egalitarian since it entails that no one's moral opinion is better than anyone else's. And it provides a convenient route to avoiding unpleasant conflicts. The person who seeks to terminate an unpleasant conflict of opinion on some moral matter by saying "what is right relative to you just isn't what is right relative to me" is appealing to something like moral subjectivism. Though adding that we should just agree to disagree would be inconsistent with moral subjectivism. And we'll take this point up as the first of a series of objections to moral subjectivism. The key to appreciating each of these objections is to see them as concerns about what is deductively entailed by moral subjectivism.

1. No Conflicts over Moral Matters

Moral subjectivism doesn't resolve any conflicts of moral opinion, it rather dissolves the appearance of conflict between moral opinions. People may have conflicts about what to do, but there are no conflicts about who is right on moral matters because rightness itself is relative to subjects on this view. Suppose I think the death penalty is just punishment in some cases and you think the death penalty is always unjust. The death penalty itself can't be both just and unjust. But moral subjectivism doesn't say that the death penalty is both just and unjust. It says that the death penalty is *just relative to me* and *unjust relative to you*. There is no contradiction or conflict here. So, there is no disagreement between us to resolve. According to moral subjectivism, there are no grounds for agreeing to disagree because we never actually disagree about the same moral matters. My approval of the death penalty is just about what is right relative to me and your disapproval is about something entirely different, what is right relative to you. Conflicting opinions concerning moral matters can be unpleasant (though they don't have to be) and moral subjectivism provides a convenient escape from conflict. We might be careful what we wish for, though. While moral subjectivism shields us from conflict and criticism, it also rules out resolving conflicts reasonably, with brings us to our next objection.

2. No Reasoning about Morality

Suppose you feel strongly about the injustice of the death penalty and you've thought a good deal about why you think it is wrong. A disappointing corollary of the no conflicts objection is that your issues with the death penalty are just about you. They can't, in principle, provide any reason for a supporter of the death penalty to change their mind. Reasoning with people presupposes a shared reality we can reason about. But this is just what moral subjectivism denies in the realm of morality. Suppose you attempted to argue that killing is always wrong and so the

death penalty is wrong. A supporter of the death penalty could happily grant that you have a good argument, but since wrongness in this argument is only about what is wrong relative to you, your argument presents no objection to the supporter's argument that the death penalty is morally just relative to him. Since rightness and wrongness are relative to individuals according to moral subjectivism, you and the supporter aren't even offering arguments about the same thing. According to moral subjectivism, no reasoning about the moral status of the death penalty is possible because there is nothing objective here for us to reason about.

3. Moral Infallibility

According to Moral Subjectivism, right and wrong relative to you is determined by your opinions and attitudes. This makes each of us morally infallible. The only way a person could do wrong, on this view is by failing to act according to their own opinions and attitudes. Now consider the case of Dylann Roof, white supremacist and mass murderer, who killed nine African Americans at a church meeting in 2015. Roof acted in accordance with his opinions and attitudes. According to moral subjectivism, his actions were right relative to him. Of course, they were wrong relative to just about everyone else, but subjectivism denies that there is any objective moral standard according to which our moral opinions and attitudes are better than Roof's. We humans often have self-righteous tendencies, but securing infallibility through moral subjectivism comes at a high price. It entails that Dylann Roof was just as right in his beliefs as others are in rejecting his white supremacy.

4. Moral Growth Undermined

Thankfully, we are not always self-righteous. We do sometimes question our own moral positions and occasionally change our view about the morality of some practice. A person might, for instance, think that eating fast food hamburgers is morally unproblematic at one time and then become convinced that animals deserve a kind of moral regard that speaks against factory farming. When moral views change in this fashion, people do not merely drop one moral belief in favor of another. Typically, they also think that their prior moral opinion was mistaken. They take themselves to have discovered something new about what is morally right. But subjectivism cannot account for any changes in our moral beliefs as being changes for the better. For a change in your moral beliefs to be a change for the better presupposes a standard of goodness that is independent of your beliefs. This, however, is exactly what subjectivism denies.

5. Anything Goes

A deep concern that underlies several of the foregoing problems is that moral subjectivism renders morality arbitrary. Anything can be right relative to an individual on the subjectivist view. If a person deems torturing innocent kittens just for fun right, then it is right relative to that person. There is no higher moral standard from which kitten torture can be condemned on this view. It is hard to see how kitten torture could be made right relative to anyone just by that person deeming it OK. This is the problem of arbitrariness. Our experiences as moral beings are not so arbitrary. Our moral judgments and attitudes are usually somewhat more systematic and principled. Moral subjectivism seems more appealing when we focus mainly on hot button moral

issues where people's opinions conflict. But most moral issues aren't like this. Almost everyone agrees that it is wrong to cheat, lie, steal, and kill. We even tend to agree on most of the cases where we'd make exceptions to these general moral rules. Moral subjectivism has little to offer in the way of explaining our broad agreement on most moral matters. And this suggests that morality is not so arbitrary.

Moral Relativism

Moral subjectivism is an individualized variety of moral relativism. There kinds of moral relativism vary depending on what groups of people morality is taken to be relative to. Perhaps more popular is cultural moral relativism (CMR), the idea that right and wrong are relative to cultures and depend only on what is deemed right or wrong in a culture. Cultures are the sources of many behavioral norms, including normative standards of dress, food preparation, polite behavior and so forth. And in cultural traditions these tend to get mixed in with matters of morality. Especially when religion is involved. So, in India, for example, it is considered wrong to eat cows. This is a normative standard that holds relative to Indian culture, but not relative to American culture. The cultural difference in normative standards is clear. What remains an open question is whether how this culturally based difference concerns morality. CMR takes the view that all of morality is a matter of cultural say so, just like styles of dress, culinary traditions, and religious customs. Now, can this view be right?

For starters, here is a bad argument for CMR.

1. Different cultures deem different things to be morally right or wrong.
2. Therefore, what is morally right or wrong is relative to cultures (CMR).

This argument isn't valid. We can think of ways for the premise to be true and the conclusion to be false. The fact that American's say its morally fine to eat cows doesn't deductively guarantee that this is fine. The premise here leaves open the possibility that Indians are on to the moral truth about cow eating and Americans just didn't get the memo. To get the conclusion that morality is relative to culture, we will need to add a premise. So here is a valid argument:

1. Different cultures deem different things to be morally right.
2. Whatever a culture deems right is right.
3. Therefore, what is morally right or wrong is relative to cultures (CMR).

Now we have a valid argument. We have unearthed the hidden assumption behind popular reasoning from moral differences between cultures to CMR. Which means we can now ask some critical questions about that assumption. Is it true that whatever a culture deems right is right? Can we think of counterexamples, cases where a culture deems something right that clearly isn't right? We don't have to look far. American culture used to deem slavery right. Cultures around the world have embraced various forms of oppression of minorities and women. Some cultures have practiced human sacrifice. There appear to be plenty of examples of cultures getting

morality wrong. So, we should at the outset have some doubts about the second premise of this argument.

Next, we might worry about what it is for a culture to deem something right. As we are culturally engrained to think equality is a good thing, most of us would probably say that a group deems something right when a solid majority of its members deem it right. But why not take a group to deem something to be right when the strongest or most aggressive members of the group deem it right? This is how things work with gangs and outlaw militias. If right and wrong are merely matters of convention, why should we favor egalitarian democratic say-so over gangland style strongman say-so? Note that it won't do to appeal to values independent of the say-so of groups or their members here, since moral relativism denies the existence of any independent values.

So, we've encountered a few serious problems for an argument in support of CMR. Now let's look at some arguments against CMR. Again, the problems we'll raise stem from what follows deductively from CMR.

1. *Anything Goes (again)*

According to CMR, whatever is right according to values and traditions of a culture is right relative to that culture. On this view, a culture's moral code defines what is right relative to it (again, what *is right*, not just what is considered right). In some cultures, it is considered right for fathers and brothers to kill female family members who have been raped. Such "honor killings" (committed in order to preserve the family's honor) strike us as morally horrible, and for this reason CMR will count honor killings as wrong relative to our culture. But CMR will count honor killing as right relative to those cultures that accept them. It looks like anything goes, on CMR.

According to CMR, all it takes for something to be right relative to a culture is for that culture to deem it right. So, if a culture deems racism, genocide, slavery or pointless sadistic torture right, then it is right relative to that culture. And this strikes many as a *reductio ad absurdum* of CMR, a refutation of CMR on the grounds that it leads to absurd consequences. So, we can formulate this argument against CMR as follows.

1. If CMR is correct, then honor killings are morally right relative to cultures that deem them right.
2. Honor killings are not right (full stop)
3. So, CMR is not correct.

This argument is valid. And the first premise correctly identifies a straightforward deductive consequence of CMR. So, the only way to avoid the conclusion that CMR is a bad theory about the nature of morality is by denying the second premise and maintaining that in fact honor killings are right relative to certain cultures (again, *are right*, not merely are considered right). So how might such a denial of the second premise go?

The relativist might start by asking, “Who are we to judge that honor killings are not right (relative to cultures that deem them right)?” If the best answer we can supply in reply to this question amounts to nothing more than a veiled appeal to the moral standards of our own culture, then the relativist has a point. But why suppose the best we can do in replying to this question is a mere appeal to the standards of our culture? Granted honor killings are wrong relative to the standards of most cultures. But most cultures may adopt moral standards that oppose honor killings for very good reasons. If members of cultures that accept honor killings fail to appreciate those reasons, that might be due to a culturally induced moral blind spot rather than morality itself depending on cultural attitudes. Having some sense of the moral worth of persons may be all that is needed to see the injustice in honor killings and it is far from obvious that this recognition depends in any way on cultural traditions or values. We recognize our own moral worth regardless of cultural traditions and values. Basic human sympathy and compassion is all that is required to recognize the same moral worth in other persons. For a culture to deem honor killings morally just would then seem to require denying the full humanity of women. It seems at least possible for a culture to get things wrong on this point. But . . .

2. *CMR makes Culture Morally Infallible*

According to CMR, what is right relative to a culture is whatever is deemed right relative to a culture. Since the moral standards of a culture define right and wrong relative to a culture, the standards of a culture can't, by definition, get morality wrong. On the face of it, the case of honor killings looks like a straightforward counterexample. Further evidence of cultural fallibility (the possibility of a culture getting morality wrong) can be found in your choice of genocidal episodes, culturally condoned practices of slavery, colonialism, caste systems, etc. We seem to have an abundance of compelling evidence for the fallibility of culturally engrained moral standards. This evidence for cultural fallibility is, like any sort of evidence, itself fallible. But suppose we try to defend CMR by suggesting that what looks like evidence of another culture's moral fallibility only looks like evidence through the distorted moral lens of our own culture. This move shows how our own culture can be fallible in its moral judgments. CMR can't have it both ways. So we have abundant historical and current evidence that cultures are not morally infallible, contrary to the deductive consequences of CMR.

3. *Reasoning about Morality is Undermined on CMR*

A straightforward consequence of the moral infallibility of cultures on CMR is that it will be conceptually impossible for any reason or argument to be offered against the moral standards of a culture. Since what is right relative to a culture is defined by that culture's standards, no rational argument to the contrary is possible. Those who disagree with a culture's moral standards, whether members of that culture or another, have no rational basis for doing so and will simply be wrong by definition. CMR has it that moral standards are based on the authority of the culture. The commands of an authority are merely to be followed, not rationally questioned. There is no room for reasoned argument where things are made so by the say-so of authority.

This stands in sharp contrast to the experience many of us have as moral beings. We do reason about what is right. We sometimes critical question the moral dictates of authority figures, cultural norms, or even our own opinions. But where there is no truth of the matter independent of cultural standards, there is no reasoning to do. We can only obey. CMR entails the unsavory consequence that we are never able to think for ourselves about what is right.

4. *The Moral Reformers Dilemma*

The impossibility of moral growth entailed by CMR can be understood in terms of the moral reformer's dilemma. We recognize a few remarkable individuals as moral reformers, people who, we think, improved the moral condition of their culture in some way. Common examples include the Buddha, Jesus, Ghandi, and Martin Luther King Jr. While the relativist can allow that these individuals changed the moral views of their societies, none can be said to have changed their societies for the better. This is because CMR recognizes no standard of moral goodness independent of what is accepted in a culture and identifying any changes in a culture's moral standards as changes for the better presupposes a standard of goodness independent of cultural say so. The relativist is committed to taking the most overt and violent forms of racism to be right relative to pre-civil rights American society and wrong relative to post-civil rights American society. But since standards of goodness are determined by the prevalent views in a society, there is no standard of goodness to appeal to in evaluating the change our society underwent in the civil rights movement as a change for the better. According to CMR, anyone who takes Martin Luther King to have improved American society by nudging it in the direction of equality is just mistaken about the nature of morality.

5. *CMR Undermines Tolerance*

In each of the objections we have considered to moral subjectivism and CMR, we have reasoned deductively from a clearly formulated relativistic view about the nature of morality to morally problematic consequences. That is, we've first gotten clear about the view we are considering, and then we have asked what would follow from that. In each case, the consequences have been unsavory. We'll conclude our critical evaluation of CMR with one more such case, the case of tolerance. Here we encounter a sad irony. Fans of cultural relativism often endorse relativism in part because it seems to support tolerance and respect for people of diverse cultures. CMR does take the differing moral standards of different cultures to be right relative to their respective cultures. It rejects the notion that the moral standards of one society could be objectively correct. This line of thought has led many who value cultural diversity and tolerance to embrace Cultural Moral Relativism. But this is a mistake. CMR does not entail that we should be tolerant of diverse cultures. Again, let's reason deductively from what CMR says.

1. According to CMR, whatever a culture deems right is right relative to that culture.
2. Culture X deems intolerance of other cultures right.
3. Intolerance of other cultures is right relative to Culture X.

The argument here is straightforward and valid. CMR entails that we should be tolerant of diversity if and only if our culture deems tolerance of diversity to be a good thing. If a culture deems intolerance to be good, then, according to CMR, intolerance is good relative to that group. Our culture may deem tolerance and respect for diversity a good thing, but as we've seen, this doesn't give intolerant cultures any reason to reconsider their intolerance. CMR thus turns out to be a deeply conservative view in the sense that it undermines all possible reasons for changing our moral outlook. CMR is a view that gives a dominant racist culture moral standing, and further denies us any reasonable grounds for arguing against the intolerance of a dominant racist culture. People who value tolerance and respect for diverse individuals or groups would do much better to endorse tolerance and respect as objective moral standard. The widespread embrace of CMR among multiculturalists is self-defeating.

[For a further compelling line of argument against moral relativism, see Paul Boghossian's piece, "[The Maze of Moral Relativism](#)." Boghossian argues that moral relativism undermines the normativity of moral beliefs altogether and ultimately collapses into nihilism. If you prefer to listen, here's a [Philosophy Bites podcast](#) in which Boghossian explains his line of argument.]

Relativism and the Social Sciences

Assorted branches of the social sciences aim to better understand and explain the diversity of cultural practices, and this includes trying to understand what people in different cultures believe is right or wrong. In describing culturally based beliefs about what is right, we aren't concerned with what is right, but just with what is considered right in a culture. Here we are characterizing a kind of **descriptive cultural relativism** which is commonly embraced in anthropology, for instance. The question of what *is* morally good or bad remains an open question for ethics.

Suspending judgment is methodologically important for understanding. This is just as true in philosophy as it is in sociology or anthropology. We suspend judgment at the stage of trying to understand a new view. Only once we have a clear understanding can we then turn to critical evaluation. Social sciences like anthropology and cultural and ethnic studies are out to understand cultural practices and perspectives. Suspending judgment is essential to doing this well.

So, guarding against ethnocentrism is important when an anthropologist investigates cultures that are different from her own. But the methodological importance of suspending judgment for the sake of better understanding is not a permanent obstacle to critical evaluation of the moral points of view transmitted through culture. Ethics, unlike sociology or anthropology, is a fundamentally normative discipline. Its goal is to evaluate moral views and try to see which are most reasonable in light of the kinds of ethical evidence and argument we can uncover. Here we benefit from the social sciences and the understanding they produce of the moral perspectives of different cultures. The diverse perspectives on morality to be found in differing cultures bring new evidence to our ethical inquiry, but like all evidence it remains open to critical reevaluation.

Divine Command Theory

Divine command theory (DCT) is the view that morality is just a matter of God's commands. According to DCT, what is right is right simply because God commands it. Note that like moral subjectivism and CMR, DCT takes morality to be based on say-so of some authority. It's just that the authority in the case of DCT is God rather than the individual, culture or society. For this reason, even though DCT makes morality objective in a sense, DCT is open to some of the same objections we have encountered for other authority-based views of morality. DCT view makes ethics easy, so long as we can be sure we know what God commands. If we can somehow be confident about that, ethics requires no critical thinking, just total obedience.

We had a much earlier encounter with DCT in our discussion of Plato's dialogue, *Euthyphro*. In that dialogue Socrates asks if the gods love what is good because it is good or if what is good is good because the gods love it. DCT is the latter option, that what is good is good because God loves it. Plato points towards the central problem with DCT in this question and, again, it's that DCT makes ethics completely arbitrary. In principle, God could command that anything be right. God could command that we torture puppies, commit genocide, and treat children like livestock. According to DCT, if God does command these things, then they are right, end of story. In fact, many people have sincerely taken God to have commanded these things (except for the puppy torture). However, hopefully, the idea that any of these things could be morally right strikes you as absurd. In spite of our occasionally differing ethical opinions, ethics does seem to be systematic and coherent. Right and wrong are not completely arbitrary. It seems at least that there is some reasoned systematicity to our ethical opinions in spite of the differences we sometimes arrive at. If this is right, then we should reject any meta-ethical view that makes ethics completely arbitrary. And this means rejecting the view that right and wrong is simply a matter of God's command.

The religious believer has better meta-ethical options than DCT. When I share Plato's compelling objection to DCT with students, it's not uncommon for someone to object that God would never command us to torture innocent puppies because God is good. I think this is exactly the right response for a believer to offer. But this response is not a defense of DCT. Any believer that makes this move is joining Socrates in rejecting DCT and taking God to command what is good because it is good. If God is essentially good, then what is right is not made right merely by his command. Rather he commands what he commands because of his goodness. When the religious believer takes God's goodness to be what is ethically fundamental, they abandon conventionalist meta-ethics in favor of a kind of theological ethical realism. Of course, the challenge of understanding God's good nature remains.

If ethics is a matter of authority as both DCT and moral relativism would have it, then there is no inquiry to engage in beyond figuring out what the relevant authority says. This would make ethics a singularly boring topic to look into. But we will find quite a few interesting things to say about morality in the next chapter. So, we might take our inquiry into normative ethics in the

next chapter to constitute one further argument against authority approaches to ethics like DCT, CMR and moral subjectivism. Ethics just isn't as dull as these views would have it. Inquiry into morality is apparently possible. And where there is inquiry, there must be some object of inquiry, some truth to be figured out.

Moral Realism

We'll now turn to the idea that morality is objective. In thinking of morality as objective, we just mean that goodness and badness, right and wrong, aren't determined by the say-so or authority of any subjects, not even God, nor the say so of groups of subjects as in CMR. What could it be for morality to be objective in this sense? One possibility is that there is real value in the world. Perhaps happiness is fundamentally good, quite aside from the fact that we generally prefer it. Perhaps people matter objectively and have a kind of moral value that transcends the feelings and opinions of others.

An immediate concern for the idea of real objectively existing value is that we can't see or otherwise empirically detect such value. We have no scientific instruments for registering and measuring the fundamental objective value of happiness, or community, or persons. Of course, we may subjectively value these things, but then we subjectively value assorted different things and this seems to favor moral subjectivism or relativism. But bear in mind that my subjective sense of the importance of people or happiness might just be subjective in the way that my visual experience is. That is as a subjective indicator of independent objective facts.

Moral realism is the view that there are moral truths that are not merely a matter of the say so of some person or group. We will consider arguments for some realist views about morality in the following chapter. First we will conclude this chapter with answers to some objections and misconceptions concerning moral realism.

1. But Morality isn't like Science!

If objective moral value is real, shouldn't we expect to find some scientific confirmation of this, some proof or evidence of its existence? Perhaps, but what sort of proof or evidence we should expect to find remains an open question. It may be helpful to consider the sort of proof or evidence science has been able to provide for the most fundamental features of physical reality. Note first that we don't have any direct observational evidence of the existence of fundamental physical forces like gravity or electromagnetic charge. We know fundamental forces only through their effects. You can observe the effects of gravitation when you fall off a bike, for instance. Less hazardously, but perhaps more depressingly (if you haven't been riding your bike, for instance) you can measure the effect of gravity when you step on the bathroom scale. But nowhere here do you get to observe the force of gravity itself. You only get to observe its manifestations, its effects. It is far from clear how real fundamental moral value is any worse off. We can't observe goodness or badness in the way we can a reading on a bathroom scale or a voltmeter, but perhaps we do have another kind of indirect experience of it. We sometimes have

very powerful experiences of indignation when we are treated badly. And we may have experiences of satisfaction when justice is served, or experiences of joy or gratitude when we see that something good has happened. These are pervasive aspects of our life experience. We do have moral experience and this calls for some explanation. The too common suggestion that these things are “merely subjective” doesn’t explain much. My visual experiences are subjective as well (in the sense of being aspects of a subject, me). Recall that the moral subjectivist takes morality to be merely subjective in the further sense that they don’t track anything happening in the objective realm. Perhaps the moral subjectivist can explain the occasional differences in our moral opinions. But we may be far too impressed with this. Our subjective moral experiences are much more alike on the whole than they are different. Our moral experiences are well enough aligned from one person to the next that we can all appreciate the subtle moral dynamics of stories. No one has any trouble recognizing the various virtues and vices of the characters in a novel or a play, for instance. We are largely on the same page when it comes to identifying fortuitous events and tragedies.

Even in the case of hot button moral issues, thoughtful people can often appreciate the conflict in terms of understandable real moral values. To take one current issue, the moral value of having our government protect us from dangerous criminals crossing our southern border is obvious to all. We also find legitimate moral considerations at work in the sympathetic understanding of the plight of people who seek asylum in our country. There may be a number of narratives and perspectives to consider in approaching complex issues like immigration. Assorted moral values will be relevant to these diverse perspectives. But even here, the situation doesn’t present difficulties for the idea of objective morality. On the face of it, security for citizens matters. And so, does concern for the basic human rights of others. The challenge is how to weigh these real legitimate moral considerations in situations where they are in tension with each other.

The idea of real objective moral value appears to be on par that of real objective fundamental physical forces in science so far. We have not found any deep difference between ethical inquiry into objective moral value and scientific inquiry into the nature of the physical world yet. There may be important dis-analogies between ethics and science to uncover and explore and we may get to some of them ahead. But so far, so good for the idea of objective moral value.

2. Dogmatism and Intolerance

If there were objective moral truths, wouldn’t this lead to people who have this truth, or at least think they do, to be dogmatically confident about their righteousness? People who are dogmatic in their moral views do tend to think their views are objectively correct. So perhaps we should be concerned about this correlation. But is it objectivism that leads to dogmatism, or rather dogmatism that leads to misguided objectivism? For a person who is dogmatically sure of their view, asserting their view as objectively correct is a convenient way of rationalizing their intolerance of other views. But too often, when we do encounter characters of this sort, the intolerant bigot or religious fanatic for instance, we also find pretty good reasons to think their

views are objectively false. Dogmatism is much more highly correlated with being objectively wrong than it is with thinking there is an objective truth of a matter when there isn't one.

Here we find further useful analogy with science. Science is concerned with discovering objective matters of fact. But this is not achieved through dogmatic certainty. Progress in science is more often a matter of overcoming dogmatism through open-minded evidence-based inquiry. Our reasons for thinking scientific results reveal objective truths is generally strengthened by testing these views against alternative hypotheses and critical challenges. Dogmatism is an obstacle to getting at the truth in science simply because it suppresses potentially revealing evidence. Open minded evidence-based inquiry is our best means of learning what is objectively true because it aims to consider all possible explanations and evaluating each in light of conflicting views.

For there to be objective morality is for there to be truths about what is good that don't depend on our beliefs, attitudes or wills. A proper appreciation of this should be humbling since it alerts us to the possibility of getting things wrong. Our confidence in our moral opinions should not be absolutely certain if morality is objective. Dogmatism only serves to shut us off from potentially important evidence and argument. This doesn't mean we should never be confident; it just means that our degree of confidence should be based on how thoroughly we have considered the evidence, reasons, and objections. Even where this justifies high levels of confidence in our moral views, we should never be too confident to consider further evidence and insight from other perspectives. Well-reasoned confidence is rationally incompatible blinding certainty.

3. Won't Objective Morality lead to Intolerance?

This concern may be based on confusing the existence of objective truths with dogmatic certainty as we've just discussed. And as that discussion suggests, acknowledging our fallibility in ascertaining objective truths should lead us towards tolerance of competing views, not away from it. The best way we have of getting at objective truth, in ethics as in science, is to consider the full range of possible views and critically evaluate each. So, getting at true moral beliefs suggest that we should be tolerant of diverse moral beliefs, even given that many of them will get morality wrong.

But then, false moral beliefs are liable to be the springs of bad actions. Must we also tolerate these? That is a separate and more difficult question. Note that there doesn't seem to be anything morally problematic about intolerance of some bad actions. To the contrary, we may have a moral duty to not tolerant murder, fraud and other bad things. But objective morality may also require us to respect human autonomy in ways that leave us vulnerable to actions motivated by false moral beliefs.

4. What about Cultural Diversity

Just as human beings are fallible as individuals, we should acknowledge that human beings are fallible in groups. Cultures can get morality wrong. To take a couple of pretty clear examples,

American culture got it wrong in practicing slavery until about a century and a half ago. More recently, German culture got morality wrong in the holocaust. We've already suggested that cultures that condone honor killings are getting morality wrong. But acknowledging the moral fallibility of cultural traditions doesn't by itself suggest we should be intolerant of other cultures. For starters, where we find a morally significant difference in the practices accepted by two or more cultures, acknowledging fallibility should alert us to the possibility that it is our own culture that has things wrong.

Many of the things billed as moral matters in cultural or religious traditions are not really moral issues at all. There may be objective moral truths, and yet none that speak to how animals should be butchered, whether women should wear veils, how many spouses a person can have, whether sins need to be confessed in a particular manner, and so forth. Many of the things deemed morally obligatory or forbidden in many cultural or religious traditions might turn out to be morally optional according to what is objectively morally correct. Cultural variation along these lines may be benign. Though we might worry further about how transgressors of otherwise morally benign cultural standards are treated.

Next, many variations among cultures may amount to different ways of expressing the same underlying objectively good things. For instance, it may be morally good to express appreciation for hospitality. Slurping one's noodles is a way of expressing this in Japan, but not so much in other places.

Finally, cultural identities and the traditions through which these are expressed might have positive objective moral value in their own right. Cultural traditions and values bind communities of people together and to a common shared past. Most of us find positive value in our varying cultural backgrounds and identities. Objective morality might require tolerance and respect for these, at least when nothing of greater moral importance is at stake. Morality probably doesn't require tolerance and respect of cultural traditions of marauding and pillaging.

5. Objective Moral Truths would be Absolute

We may be concerned that if there are objective moral truths, these would be absolute in the sense of not allowing for exceptions. Perhaps objective moral truths would be like universal laws of nature in that they hold in all places and at all times. So, if "thou shalt not kill" is an objective moral truth, we might worry that this would rule out killing self-defense.

This concern is hasty in linking the idea objective truth to truth that is absolutely universal and exceptionless. For moral truth to be objective is just for it to hold independent say so, authority, or cultural tradition. So here is a logically coherent, if silly view: there is just one objective moral truth and it holds that that Jimmy shall not flush his dead goldfish down the toilet. This view is silly because it fails to explain any of our moral intuitions or evidence. But it does provide a clear example of how there could be moral truths that don't involve exceptionless general rules.

Next, note that many general rules still allow for exceptions. US Tax laws are general rules that apply to all income earners in the US and they allow deductions for IRA contributions, except if you have earned more than a certain amount. In this case the exception is explicit and written into the rules. But this needn't be the case. In fact, the most laws of science are *ceteris paribus* laws, laws that hold "all things being equal." So, for instance, Newton's law of universal gravitation does not accurately describe the descent of a paper airplane. Paper airplanes don't violate the laws of physics. When we try to understand the objective truths of physics, we do so with the understanding that the forces our fundamental laws describe often operate in interaction with other forces and the results will not always align neatly with the generalizations suggested by individual laws. Similarly, when we identify an exception to a general moral rule like "don't kill" we expect an explanation of this exception in terms of some factor that is itself morally compelling. So, it is a mistake to infer absoluteness from objectivity in ethics, just as it is in science and other realms.

Conclusion

Unless you have had some prior philosophical education, you have probably assumed that morality is either subjective, a matter of society's say-so, or a matter of God's say-so. People of religious faith tend to go for Divine command theory and non-believers tend to merely substitute society's say-so for Gods. The idea that morality is independent of these things and so can be the subject of critical inquiry is rather unfashionable outside of philosophical circles. At least that seems to be the bias in contemporary American culture. This chapter has followed fairly standard practice in introductory ethics texts in alerting readers to the hazards of popular opinions about the nature of morality as grounded in authority or say-so, and in defending moral realism against some common misconceptions. This is hardly the end of the meta-ethical story. There remain several views about the specific nature of morality to explore. But the popular favorites, Divine command theory and moral relativism, have been nearly universally rejected by philosophers since Plato for the sorts of reasons we have covered in this chapter. Hopefully you find these reasons compelling. Morality grounded in authority or say-so doesn't afford much opportunity for critical inquiry. The work of this chapter has been to clear away such intellectual dead ends and prepare the way for inquiry into normative ethics. Critical inquiry into substantive theories of what is right and wrong, good or bad, presupposes and object of inquiry that is independent of God's say so or our say-so, either individually or collectively. Seeing how such inquiry can proceed may give you further cause for taking moral realism seriously. The evidence we have to work with starts with our experience as moral beings. While this experience is ours, that doesn't imply that it is subjective in the sense of being relative to us as individuals. Our moral sense makes more sense, and provides the opportunity for further inquiry, when our moral experience is understood as a guide to things larger than ourselves.

Review Questions

1. What is the aim of normative ethical inquiry?
2. What is nihilism?
3. What is moral relativism?
4. Explain moral subjectivism as an individualized form of moral relativism.
5. What is Divine command theory?
6. What is moral realism?
7. How does metaethics differ from normative ethics?
8. Explain the challenge nihilism faces as a skeptical hypothesis.
9. How does moral subjectivism and other varieties of moral relativism undermine reasoning about morality?
10. Moral subjectivism makes each person infallible judges of morality relative to themselves. Why might we regard this as a problem?
11. How does moral subjectivism along with other varieties of moral relativism undermine the idea of moral growth?
12. Explain the concern about arbitrariness for moral subjectivism and other varieties of moral relativism.
13. What is attractive about cultural moral relativism (CMR)?
14. Why should we worry about how what is right relative to a culture gets determined on CMR?
15. Give some examples that raise concerns about arbitrariness for CMR?
16. Explain the Moral Reformers Dilemma.
17. Why does CMR fail to support tolerance?
18. How does talk of cultural relativism in the social sciences differ from CMR as an ethical thesis?
19. What is the problem with Divine command theory and what better alternative is open to religious believers?
20. In what ways is experience based inquiry into morality similar to scientific inquiry?
21. How does experience based inquiry into morality differ from scientific inquiry?
22. How can objective morality accommodate cultural diversity? When should it not?
23. How does objective morality support tolerance?
24. How is dogmatism and intolerance problematic if morality is objective?

11. Right Action

Our focus in this chapter will be normative ethics. Normative ethical principles aren't intended to describe how people think or behave. Normative ethics is concerned how we should be motivated and how we should act. Our project here is to think critically about which normative ethical principles do the best job of explaining our assorted moral intuitions about the broadest range of possible cases. We will start with Utilitarianism, a view of right action based on the idea that happiness has fundamental value. We'll then examine Kant's ethics of respect for persons. On this view persons have intrinsic moral worth, and ethics is concerned with what respecting for persons requires of us.

Both Utilitarianism and Kant's ethics of respect for persons can be understood as aiming to formulate action-guiding normative ethical principles. Later in the chapter we will consider approaches to normative ethics that are not so concerned with identifying exceptionless "laws" of right action. Our understanding of right action doesn't have to be expressible in terms of strict rules. Feminist ethics finds value in caring relationships. But taking relationships to be good doesn't directly lead to specific rules for action as Utilitarianism might. Environmental ethicists have advanced various proposals for expanding the realm of moral relevance to include other species or systems of life as a whole. This is not to deny that people matter morally, but many environmental ethicists deny that people are all that matter. Accounting for the value of non-persons in addition to persons is likely to frustrate attempts to characterize right action in terms of simple formulas or "moral laws."

At the end of this chapter, we will consider a pluralistic approach to understanding ethical motivation and action. The suggestion here will be that a substantive realist approach to normative ethics doesn't require reducing all ethical value to one fundamental kind. Such a pluralistic account of ethical value undermines the quest for simple exceptionless or absolute moral principles. But it also suggests that substantive realist normative ethics doesn't require these either.

Utilitarianism

Utilitarianism is based on the idea that happiness is good. Utilitarian thinkers have traditionally understood happiness in terms of pleasure and the absence of pain. Utilitarianism's best-known advocate, John Stuart Mill, characterizes Utilitarianism as the view that "an action is right insofar as it tends to produce pleasure and the absence of pain." If happiness, conceived of as pleasure and the absence of pain, is the one thing that has value, then this criterion of right action should seem to follow straightforwardly.

In any given scenario, every possible course of action will have a utility. The utility of an action is the net total of pleasure caused by the action minus any pain caused by that action. In

calculating the utility of an action, we are to consider all of the effects of the action, both long run and short run. Given the utilities of all available courses of action, Utilitarianism says that the correct course of action is the one that has the greatest utility. So, an action is right if it produces the greatest net total of pleasure over pain of any available alternative action. Note that sometimes no possible course of action will produce more pleasure than pain. This is not a problem for Utilitarianism as we've formulated it. Utilitarianism will simply require us to pursue the lesser evil. The action with the highest utility can still have negative utility.

Utilitarianism places no privileged status on the happiness of the actor. It's happiness that matters, not just your happiness. So, Utilitarianism can call for great personal sacrifice. The happiness of my child over the course of his lifetime might require great personal sacrifice on my part over the course of his first few decades. Utilitarianism says the sacrifice should be made given that the utility at stake for my child is greater than the utility at stake in my child-rearing sacrifices.

Likewise, Utilitarianism places no privileged status on the immediate, as opposed to the long term, effects of the action. An action's utility is the net amount of pleasure or pain that is experienced as a result of the action over the long run. So, while it might maximize a small child's pleasure in the short run to be given ice cream whenever he wants it, the long run utility of this might not be so good given the habits formed and the health consequences of an over-indulged sweet tooth.

There is an obvious concern to address at this point. We often don't know what the long-run consequences of our actions will be, and even in the short run we are often uncertain about just how much pleasure and pain will be caused for the various parties affected. So, we might not be able to calculate the utilities of alternative actions to figure out which action will have the highest utility. These are practical problems for applying utilitarian theory. But while it might be difficult to tell on a case-by-case basis just which course of action will maximize utility, this is not a problem for Utilitarianism as a normative ethical theory. As a normative ethical theory, Utilitarianism is aimed at identifying the standard for right action, not at telling when a particular action meets that standard. Setting the standard for right action and figuring out how to meet that standard are two different projects.

When we speak of utility as pleasure and the absence of pain, we need to take "pleasure" and "pain" in the broadest sense possible. There are social, intellectual, and aesthetic pleasures to consider, as well as sensual pleasures. Recognizing this is important to answering what Mill calls the "doctrine of swine" objection to Utilitarianism. This objection takes Utilitarianism to be unfit for humans because it recognizes no higher purpose to life than the mere pursuit of pleasure. The objector takes people to have more noble ends to pursue than mere pleasure. According to this objection, Utilitarianism is a view of the good that is fit only for swine. Mill responds that it is the person who raises this objection who portrays human nature in a degrading light, not the utilitarian theory of right action. People are capable of pleasures beyond mere sensual

indulgences and the utilitarian theory concerns these as well. Mill then argues that social and intellectual pleasures are of an intrinsically higher quality than sensual pleasure.

We find a more significant objection to Utilitarian moral theory in the following sort of case: Consider Bob, who goes to the doctor for a checkup. His doctor finds that Bob is in perfect health. And his doctor also finds that Bob is biologically compatible with six other patients she has who are all dying of various sorts of organ failure. Let's assume that if Bob lives out his days he will live a typically good life, one that is pleasant to Bob and also brings happiness to his friends and family. And we will assume that Bob will not discover a cure for AIDS or bring about world peace. And let us make similar assumptions about the six people suffering from organ failure, that if they recover, they will go on to live lives of typically high utility. According to utilitarianism, it looks like the right thing for Bob's doctor to do is to kill Bob and harvest his organs for the benefit of the six patients who will otherwise die. But intuitively, utilitarianism gets the wrong result in this sort of case. This case seems to provide a clear counterexample to simple utilitarianism. This looks like a bit of evidence that calls for a change in theory. But perhaps that change can be a modification of utilitarian thinking rather than a complete rejection of it.

One move open to the utilitarian is to evaluate rules for acting rather than individual actions one by one. A version of **Rule Utilitarianism** might say that the right action is the action that follows the rule which will produce the highest utility in comparison to other possible rules. A rule that tells doctors to kill their patients when others require their organs would not have very high utility in general. People would avoid their doctors and illness would go untreated were such a rule in effect. Rather, the rule that doctors should do no harm to their patients would have much higher utility in general. So, the move to Rule Utilitarianism seems to avoid the difficulty we found with utilitarian thinking about individual actions. Or at least it seems to when we consider just these two rules.

But here is a rule that would have even higher utility than the rule that doctors should never harm their patients: doctors should never harm their patients except when doing so would maximize utility. Now suppose that doctors ordinarily refrain from harming their patients and as a result people trust their doctors. But in Bob's case, his doctor realizes that she can maximize utility by killing Bob and distributing his organs. She can arrange to kill Bob in a way one will ever discover, so her harming Bob in this case will not undermine people's faith in the medical system. The possibility of rules with "except when utility is maximized" clauses renders Rule Utilitarianism vulnerable to the same kinds of counterexamples we found for Act Utilitarianism. In effect, Rule Utilitarianism collapses back into Act Utilitarianism.

In order to deal with the original problem of Bob and his vital organs, the advocate of Rule Utilitarianism must find a principled way to exclude certain sorts of utility maximizing rules. I won't pursue this matter on behalf of the Rule Utilitarian. Rather, I want to consider further just how simple utilitarianism focused on actions goes wrong in Bob's case. Utilitarianism evaluates

the goodness of actions in terms of their consequences. For this reason, Utilitarianism is often referred to as a *consequentialist* theory. Utilitarian considerations of good consequences seem to leave out something that is ethically important. Specifically, in this case, it leaves out a proper regard for Bob as person with a will of his own. What makes Bob's case a problem case is something other than consequences, namely, his status as a person and the sort of regard this merits. This problem case for utilitarian moral theory seems to point towards the need for a theory based on the value of things other than an action's consequences. Such non-consequentialist ethical theory is called **deontological** ethical theory. The best-known deontological theory is the ethics of respect for persons. And this will be our next topic.

Here is a link to John Stuart Mill's essay *Utilitarianism*:

<http://www.gutenberg.org/files/11224/11224-h/11224-h.htm>

Respect for Persons: Kant's Moral Theory

Like Utilitarianism, Immanuel Kant's moral theory is grounded in a theory of intrinsic value. But where the utilitarian takes happiness, conceived of as pleasure and the absence of pain to be what has intrinsic value, Kant takes the only thing to have moral worth for its own sake to be the capacity for good will we find in persons. Persons, conceived of as autonomous rational moral agents, are beings that have intrinsic moral worth and hence beings that deserve moral respect.

The opening passage of Immanuel Kant's *Groundwork for a Metaphysic of Morals* proclaims that "it is impossible to conceive of anything in the world, or indeed beyond it, that can be understood as good without qualification except for a good will." This is a clear and elegant statement of the theory of value that serves as the basis for Kant's ethical theory of respect for persons. The one thing that has intrinsic value, for Kant, is the autonomous good will of a person.

That said, Kant does not understand the expression "good will" in the everyday sense. In everyday discourse we might speak of someone being a person of good will if they want to do good things. We take the philanthropist's desire to give to the less fortunate to be an example of good will in this everyday sense. On Kant's view, the person of good will wills good things, but out of a sense of moral duty, not just inclination. Naturally generous philanthropists do not demonstrate their good will through their giving according to Kant, but selfish people do show their good will when they give to the poor out of a recognition of their moral duty to do so even though they'd really rather not.

So, it is the ability to recognize and act according to a moral duty that gives people dignity and makes them worthy of moral regard. The good will is the capacity to will and act out of duty or respect for moral law. As persons, we have a free will in our capacity to weigh our desires against each other and against the rational constraints of morality and determine our own will. On Kant's view, our free will, our moral autonomy, is our capacity to act according to duty as

opposed to being a slave to our desires or inclinations. So free will, in the sense that is associated with moral responsibility, doesn't mean being free to do as you please without consequence. Rather, freedom comes with moral responsibility for the intentions we act on. Having an autonomous good will with the capacity to act from moral duty is central to being a person in the moral sense and it is the basis, the metaphysical grounding, for an ethics of respect for persons. Now what it means to respect a person merits some further analysis.

Kant calls his fundamental moral principle the Categorical Imperative. An imperative is a command. The notion of a Categorical Imperative can be understood in contrast to that of a hypothetical imperative. A hypothetical imperative tells you what to do in order to achieve some goal. For instance, "if you want to get a good grade in calculus, work the assignments regularly." This claim tells you what to do in order to get a good grade in calculus. But it doesn't tell you what to do if you don't care about getting a good grade. What is distinctive about a Categorical Imperative is that it tells you how to act regardless of what end or goal you might desire. Kant holds that if there is a fundamental law of morality, it is a Categorical Imperative. Taking the fundamental principle of morality to be a Categorical Imperative implies that moral reasons override other sorts of reasons. You might, for instance, think you have a self-interested reason to cheat on exam. But if morality is grounded in a Categorical Imperative, then your moral reason against cheating overrides your self-interested reason for cheating. If we think considerations of moral obligation trump self-interested considerations, Kant's idea that the fundamental law of morality is a Categorical Imperative accounts for this nicely.

Here are two formulations of Kant's Categorical Imperative:

CIa: Always treat persons (including yourself) as ends in themselves, never merely as a means to an end.

CIb: Act only on that maxim that you can consistently will to be a universal law.

Kant takes these formulations to be different ways of expressing the same underlying principle of respect for persons. They certainly don't appear to be synonymous. But we might take them to express the same thing in the sense that each formulation would guide one to act in the same way.

The formulation (CIa), tells us to treat individuals as ends in themselves. That is, persons should be treated as beings that have intrinsic value. To say that persons have intrinsic value is to say that they have value independent of their usefulness for this or that purpose. (CIa) does not say that you can never use a person for your own purposes. But it tells us we should never use a person *merely* as a means to our own ends. We treat people as a means to our own ends in ways that are not morally problematic quite often. When I go to the post office, I treat the clerk as a means to my end of sending a letter. But I do not treat that person *merely* as a means to an end. I pursue my end of sending a letter through my interaction with the clerk only with the understanding that the clerk is acting autonomously in serving me. My interaction with the clerk

is morally acceptable so long as the clerk is serving me voluntarily for his own reasons. By contrast, we use people merely as a means to an end if we force them to do our will, or if we deceive them into doing our will. Coercion and deception are paradigm violations of the Categorical Imperative. In coercing or deceiving another person, we disrupt his or her autonomy and his or her will. This is what the Categorical Imperative forbids. Respecting persons requires refraining from violating their autonomy.

Now let's consider the second formulation (CIb). This version, known as the formula of the universal law, tells us to "act only on that maxim that you could consistently will to be a universal law." The maxim of our action is the subjective principle that determines our will. It includes our purpose. The general form of a maxim might be phrased like this:

I will act in way A in order to achieve purpose P

We act for our own reasons. Different intentions might lead to similar actions. I might put on a dress shirt for the purpose of making a professional impression in the classroom. My son might perform the same action for the purpose of satisfying his mother's wishes for Christmas dinner. We can identify different maxims in terms of these different purposes or intentions. For Kant, intentions matter. He evaluates the moral status of actions not according to the action itself or according to its consequences, but according to the maxim which includes the act and the purpose intended. What gets evaluated as morally right or wrong is the maxim, not just the act.

According to the formula of the universal law, what makes an action morally acceptable is that its maxim is universalizable. That is, a morally permissible action is an action that is motivated by an intention that we can rationally will that others act on similarly. A morally prohibited action is one where we can't rationally will that our maxim is universally followed. Deception and coercion are both paradigm cases of acting wrongly according to Kant. In both cases, our maxim involves violating the autonomy of another rational being and this is something that we, as rationally autonomous beings ourselves, could not consistently will to be a universal law. According to Kant, there is a contradiction involved in a rational autonomous being willing coercion or deception. Universalizing maxims that involve coercion or deception would amount to willing the violation of one's own rational autonomy. Acting out of moral duty is a matter of acting only on maxims that we can rationally will others act on as well. The person of good will recognizes the humanity of others by not making any special exception for herself even when her interests or inclination would be served by doing so.

There is no higher moral authority than the rational autonomous person, according to Kant. Morality is not a matter of following rules laid down by some higher authority. It is rather a matter of writing rules for ourselves that are compatible with the rational autonomous nature we share with other persons. We show respect for others through restraining our own will in ways that demonstrate our recognition of them as moral equals.

Primary Source Reading:

Kant's *Groundwork for a Metaphysic of Morals* can be found here:

<https://www.earlymoderntexts.com/authors/kant>

Ethical Pluralism

In ethical theory, we can understand pluralism as the view that there is a plurality of fundamentally good things. Traditionally, ethicists have tried to analyze right and wrong action in terms of a single fundamental underlying kind of value. We can call this kind of approach *ethical monism*. For utilitarians that single value is happiness, for Kantian respect for persons theorists, it is the value of the person. Ethical Pluralism allows that there may be multiple kinds of fundamental and irreducible value in the world. Happiness and respect for persons might be among these, but there may be others yet. But it doesn't take morality to be a matter of picking or choosing among basic moral values.

Let's think a bit about how ethical pluralism differs from moral relativism. Recall that according to Moral Relativism, what makes something right relative to a group is just that it is deemed to be right by that group. Views that take morality to be matter of authority, whether of the individual, a culture or God, all suffer the same basic defect. They render right and wrong entirely arbitrary. If someone or some group gets to decide what's right and wrong, then anything can be right or wrong. According to CMR, whatever a culture deems to be morally right is right relative to it. So, if our culture says that homophobia, sexism, and racism are fine, then they are what is right relative to our culture and that's the end of it. If some people don't like it, that's just too bad. Moral Relativism denies them any objective standpoint from which to complain or any possibility of providing reasons for changing things. Complaints about the oppressiveness of the dominant group amount to nothing more than the whining of losers. The group that dominates is perfectly well within its rights to dominate on relativist thinking. This hardly sounds like a plausible account of social justice. But it is straightforwardly entailed by Moral Relativism and that's one reason why Moral Relativism is an awful ethical theory. This much is just a bit of review from the last chapter. But bear this in mind for the purpose of recognizing how Ethical Pluralism avoids this defect. For according to Ethical Pluralism, the fundamental ethical values are real. The importance of happiness comes with the existence of pleasure and pain. The value of respect for persons comes with the existence of persons. This doesn't depend on the whim or say-so of any authority.

Suppose morality doesn't depend on the say-so of cultures, God, or any other individual or group. On this view goodness is "out there" in the realm of things to be discovered. It needn't be "way out there," like goodness in some cosmic sense or goodness for the universe at large. We're just interested in goodness for human beings and this might have lots to do with our nature as persons. So let's set aside the relativist's claim that goodness is decided by us and ask what else goodness for humans might be. In doing so, we take goodness to be an appropriate object for inquiry, not merely a matter of custom, something somebody gets to decide, or a tool for tyranny. We have some evidence to guide us in this inquiry and it includes all of our varying perspectives

on what is good (the more the better). But just as in the sciences, our evidence is fallible and needs to be tested, both against other evidence and the explanatory power of theory.

From the 17th through the 19th centuries, the Holy Grail in ethical theorizing was to find a single, rationally defensible criterion of right action. This quest was dominated by utilitarians like Bentham and Mill, and respect for persons theorists like Kant. Both theoretical approaches are value *monist*, that is, they take there to be just one thing that has value fundamentally. For the utilitarian it is happiness that matters and the goal is to formulate a single law of morally right action that aims at maximizing happiness. For Kant it is the good will, or the dignity of the person that matters, and the goal is to establish a single moral law that properly captures what it means to respect the value of the person.

The utilitarian might start with the idea that an action is right if it produces the greatest amount of happiness of any available action. But this clearly conflicts with respect for persons as we saw above in the case of Bob and his vital organs. There are various moves a utilitarian might make to try to address this case, but there are more subtle cases yet where Utilitarianism seems to conflict with respect for persons. So, it looks like we can't coherently sign on to both a utilitarian and a Kantian criterion of right action since they will conflict in interesting ways. Utilitarian standards of right action tend to be logically incompatible with standards like Kant's Categorical Imperative. If what we are looking for is a single criterion of right action that is based on a single kind of ultimate ethical value, it looks like we have to pick a single winner among competing monist ethical theories. But perhaps this sets the wrong kind of goal for ethical theory.

The idea that there might be a single universal and absolute criterion of morally right action strikes many who value cultural diversity as highly problematic. But lest we abandon monist approaches to ethical theory too quickly, we should note that the standards of right action offered by both the utilitarian and the Kantian are highly abstract and for this reason they are quite compatible with a rich range of diversity in more specific derivative guidelines for action. In fact, lots of cultural diversity can be explained in terms of more broadly shared underlying moral values. Eating the dead may be seen as a way of honoring them in one culture, but be considered a sacrilege in another culture. Both of these diverse practices, however, can be seen as ways of expressing respect for persons. The difference between cultures in this case is not really a difference of fundamental moral values, but a difference in how these are to be expressed. Similarly, we consider infanticide morally wrong while other cultures facing more difficult environmental pressures may practice it routinely. What may seem like conflicting moral standards at this more specific derivative level might instead be understood as differing ways of maximizing happiness that are appropriate for the starkly different circumstances that the respective cultures must deal with. So, absolutist, universalizing, monist ethical theories turn out to be considerably more accommodating of cultural diversity than we might have thought at first. Still, they may not be flexible enough.

It might be that some cultures value respect for persons over happiness while others value happiness at the expense of respect for persons and others yet value community or kinship relations more than happiness or respect for individual persons. That is, we might find conflicts in the most basic or fundamental moral values upheld by diverse cultures. How can ethical theory account for this without begging questions against one set of cultural values or another?

Recall that the ethical monist is out to discover a single rationally defensible moral truth that is grounded in a single kind of moral value. In discussing monist ethical theories, I insisted that you can't be both a utilitarian and a Kantian respect-for-persons theorist. This is because these theories offer logically incompatible principles of morally right action. There will be actions (like harvesting the healthy patient's organs in the simple versions) that one theory will deem to be right and the other will deem to be wrong. So, you can't coherently hold both a utilitarian principle of right action and a Kantian principle of right action to be true. If the principles disagree on even a few cases, they can't both be true. But let's set principles aside for a moment. I'm not suggesting we be unprincipled; I just want us to focus on the underlying moral values without worrying about truths that might be based on them. There is nothing logically incoherent about taking happiness and respect for persons to both be good in fundamental ways. And there may be other plausible candidates for fundamental goodness. Happiness and respect were just the ones that got most of the attention in the 18th and 19th centuries. Since then, feminist philosophers have argued that we should recognize a fundamental kind of value in caring relationships. Environmental ethicists have argued that we should recognize a fundamental kind of value in the natural world. Hindus and Buddhists have long suggested that there is a kind of fundamental value in consciousness.

Perhaps this short list is long enough. Or perhaps it is already too long. A moral value is only fundamental if it can't be explained and supported in terms of some other fundamental value. So, if caring relationships matter just because they bring happiness to human lives, then we already have this kind of value covered when we recognize happiness as a kind of fundamental value. But it is not at all clear that happiness fully explains the value of caring relationships. There are issues to explore here and feminist philosophers are just starting to map out this terrain. In any case, kinds of fundamental value might be rare, but still plural.

So, what should ethical theory say about cultures that differ in the fundamental values that shape their customs and codes? Monist approaches to ethical theory would insist that we pick winners in this kind of situation. But should we? Certainly, in some cases we should. The fundamental values of Nazi culture were racist through and through. Good ethical theory should not be accommodating this kind of cultural diversity at all. Recall that our most compelling argument against Moral Relativism was that it is committed to accepting that racism is right relative to racist societies and our condemnation of racism has no more moral force than their endorsement of it.

But what about cases like Confucian cultures that give kinship relationships a higher priority than respect for persons? The more individualistic cultures of the West would favor respect for persons. Must we pick a winner here? Monist ethical theories would insist. But pluralism about ethical value offers us a few other options. The ethical pluralist can say that both cultures are structured around worthy fundamental values and neither unjustly favors one kind of fundamental value at the expense of another. Or a pluralist might allow that some ways of prioritizing worthy fundamental ethical values really are better than others, but hold that there is no strict rational formula for working out which is best. Because we have a plurality of worthy fundamental ethical values and these are not reducible to each other or anything more basic, rational methods of inquiry might not be up to settling the matter and the best we can hope for is good judgment. But however we settle these issues, pluralism about fundamental ethical value opens some new avenues for explaining a broader range of cultural diversity without legitimizing the sorts of obvious injustice moral relativism would condone.

There are many issues to address yet in exploring Ethical Pluralism and I won't get to them all here. But some loom too large to ignore. In particular, you might be worried that over the past few paragraphs I merely assumed that the fundamental values of Confucian cultures are worthy ethical values but the fundamental values of Nazi culture aren't. How do we figure out which fundamental values are worthy and merit a place in our ethical theorizing and which don't? Monist ethical thinkers like Kant and Mill faced the same issue, they were just limiting themselves to identifying one kind of value. If I'm given a fundamental value, say respect for persons, then I can argue for more derivative values, being honest for instance, on the grounds that these are required for respecting persons. But when it comes to fundamental values, this strategy for justifying value is no longer open. I've come to the end of the explanatory and justificatory line. So, what now? What are my grounds for taking some fundamental values to be worthy ethical values but not others? The evidence in ethics is not like the evidence in physics. But then the evidence in physics is not really like the evidence in anthropology. Still, I think we do have evidence in ethics. The evidence in ethics consists of our ethical intuitions. We do have a moral sense about things.

Our ethical intuitions do differ around lots of issues, but that's not an argument for skepticism or relativism. People disagree about how to understand scientific evidence, too. The evidence of our senses can be misleading and even systematically distorted. We certainly don't just sense that the earth spins and travels around the sun. What we sense seems quite contrary to the truth of this matter. Likewise, the evidence provided by our ethical intuitions is fallible and even has the potential for misleading us systematically. Things are no different here than they are in any branch of inquiry. Our job as inquirers in ethics is to account for the evidence of our various ethical intuitions as best we can by formulating theories that help make sense out of them. As we try to systematize our ethical intuitions, we will encounter problem areas where some intuitions conflict with our best theories and explanations. Since our intuitions are fallible, such conflicts don't automatically mean our theories are just wrong. There might be creative ways to reconcile

such evidence with our best theories, or we might find that the evidence is defective or distorted in some way, or we might find grounds to alter or refine our theories in light of the evidence. There are at least these different paths our inquiry might take. Likewise, each of these paths is open when the evidence of the senses seems to conflict with our scientific theories. Inquiry in ethics is pretty much like other kinds of inquiry. Our reasoning engages us in a continual negotiation between our experience and how to best understand it. Our experience shapes our theoretical understanding and our theoretical understanding shapes our experience in turn in an organic process of intellectual growth. Reason doesn't dictate any outcomes; it merely provides the system of currency in which this negotiation towards deeper understanding takes place.

So, let's illustrate how this negotiation works with the case of the Nazis. Why reject their fundamentally racist values? There are probably lots of good reasons, but here's one: The value of respect for persons accounts for a very broad range of ethical intuitions about how we should treat people and there is no way to reconcile general respect for persons with Nazi racism. So much the worse for Nazi racist values, they don't merit any place in our ethical theory. The ethical intuitions of Nazis should be rejected as systematically distorted.

The last issue I'll take up here has to do with oneness. Just why is oneness so special? Why would philosophers like Kant and Mill think it so important to have just one kind of fundamental ethical value? One powerful appeal of oneness is that it allows for a high degree of precision and rigor. Bentham even hoped that we would one day have a calculus of utility that would allow us to rigorously prove which actions will maximize utility and therefore be right. The powerful appeal of oneness here is that it allows us to completely replace human judgment with rational calculation. We have yet to outgrow this intellectual lust for reduction. Many of us still want to see the sciences as in some way reducible to just one, physics. But philosophers of science have been raising a steady stream of questions about our reductionistic inclinations over the past few decades. And even physics itself appears to be stuck with a kind of force pluralism. The fundamental forces, according to our best theory, include nuclear, gravitational, and electrical forces. We have specific theories that explain the behavior of things if we abstract away from other forces and focus only on gravity. And we have specific theories that explain the electrical behavior of things, but only when we ignore other forces that might be at play. Similarly, some version of Utilitarianism might give us the ethical truth about the value of happiness at least when no other important ethical values are relevant. And some interpretation of Kant might give us ethical principles that get at the truth so long as we abstract away from ethical values other than the moral dignity of persons. Plurality in both ethics and physics denies us the satisfaction of a single specific formula that accounts for absolutely everything. But that shouldn't bother us too much. I rather doubt that this kind of intellectual satisfaction is really worthy of human beings.

Review and Discussion Questions

1. What, according to Utilitarianism, has fundamental value?
2. What is the utility of an action?
3. How does Mill's Utilitarianism understand happiness?
4. What is it for an action to be right according to Act Utilitarianism?
5. Describe a problem case for Utilitarianism.
6. Explain Rule Utilitarianism and the risk it faces of collapsing back in to act Utilitarianism
7. What has fundamental value according to Kant?
8. What is the difference between a hypothetical imperative and a Categorical Imperative?
9. Explain a version of Kant's Categorical Imperative.
10. What does it mean to refer to Utilitarianism and Kant's respect-for-persons theory as monist theories?
11. Explain Ethical Pluralism. How does it differ from Moral Relativism?

12. Social Justice

Social justice is just the idea of goodness as applied to social groups. When asked what it means for a society to be just, most of us will think of things like freedom and equality. But things haven't always been thus. Valuing freedom is a pretty recent innovation. We have already noted John Locke as an early advocate of liberal political thinking in the 17th century. Older conceptions of justice were neither egalitarian nor freedom loving. Here we'll consider Plato's.

Plato develops his conception of justice in the *Republic*. Here Plato develops a view of the ideal state as modeled on that of the ideal person. The state is understood as the person writ large. The idea of justice, for Plato, was as much a virtue of the individual person as of the state. Justice was seen as a kind of meta-virtue. The just person is the person who has all the other virtues and has them in the appropriate integrated balance. People have various capacities and abilities and we have various virtues that correspond to those abilities. We can be courageous in facing threats, temperate in managing our appetites, diligent in carrying out our projects, and wise in deliberating about what to do and how. To be a just person is for the various abilities relevant to the various virtues to be playing their proper role. When we turn to the justice of communities, we find different individuals playing the various roles. We want the virtue of wisdom in the ruling class, the virtue of courage in the military class, and the virtues of temperance and diligence in the business class. The just community, in Plato's view, is the community where the various elements stick to their proper roles and cultivate the virtues appropriate to those roles.

Though Athens was a democracy, Plato was no fan of democracy. In his dialogues he has Socrates repeatedly lampooning democracy as rule by the least qualified. This is because the leaders in a democracy are not chosen by the wise, but by the majority, and the majority is often easily manipulated by bad actors. As a result, Plato endorsed a kind of elitism, the rule by experts or "philosopher kings." His idea of justice is one where the various functions of society are carried out by those who have the virtues appropriate to the specific role. Plato's ideal state is a meritocracy. While Plato places no particular value on equality or freedom for individuals, he does support equal opportunity. The point of equal opportunity in this meritocratic system was not to be fair to individuals, though. The goal was simply to identify and cultivate talent wherever it is found.

However, we might feel about the inegalitarian view of justice Plato develops in the *Republic*, he raises an important problem that every political system faces. Specifically, how do we reliably fill positions of power with people who are competent and will conduct themselves in the interest of the public. Plato's answer to the competence issue was to select leaders through a rigorous meritocratic education system. To discourage leaders from abusing their power to serve personal ends rather than the good of the state, Plato also would have his philosopher kings be wards of the state for life, owning no personal property and even severing all family ties to avoid the corrupting tendencies of self-interest.

It's only in the last few centuries that ideals of equal individual rights and freedoms begin to gain traction. We'll turn to these now.

Freedom and Equality

We should note at the outset that freedom and equality are both highly ambiguous notions. We can be equal or unequal in a wide variety of different ways. Socialism, traditionally understood as public ownership of the means of production, emphasizes equality of wealth and resources in ways that are liable to frustrate some kinds of freedom. In more liberal traditions, those that emphasize liberty, equality is incorporated in terms of equal liberties, equal treatment before the law, equal opportunity, equal access to public goods, and so forth. Talk of freedom can also refer to assorted different things. Freedom can be thought of in negative terms as in being free from the dominance of others or in positive terms as in being free to do what we like. And there are many kinds of freedom. Economic freedom is one thing, freedom of conscience is another. Then there is freedom of expression, freedom of association, freedom of movement, and so forth. Clarifying our political thought in liberal traditions requires being specific about what we mean by talk of freedom and equality. Not everyone who claims to love freedom and equality loves the same thing.

Here we will focus on two giants of liberal political philosophy, John Locke and John Rawls. What makes a political philosophy a liberal political philosophy is just that it takes individual liberty, in one form or another, to be a fundamental virtue of the just state or society. So, liberal political thought stands in contrast to both communism on the left and fascism or nationalism on the right. Liberalism rejects aristocracy, authoritarianism, totalitarianism, oligarchy, and plutocracy (I'll leave those fancy words for you to look up). At least until recently, what we have known as "liberal" and "conservative" political thought have both been squarely in tradition of liberal political philosophy. Both place a high priority on liberty, even when liberals and conservatives conceptualize liberty a bit differently. You will find John Locke's thinking to be more in line with contemporary conservatism and John Rawls' thought to be more reflective of contemporary political liberals. We'll begin with John Locke

John Locke (1632-1704)

John Locke's *First Treatise on Government* was an extended argument against the European system of aristocracy and the alleged divine birth right of rulers. In a society that had only known government by the rule of kings, this raises an obvious question. If human society is not legitimately organized by the authority of a ruling class, then how is it to be organized? Locke addresses this issue in his *Second Treatise on Government* which can be found here: [EMT - John Locke \(earlymoderntexts.com\)](http://www.earlymoderntexts.com)

Locke's political theory starts with thinking about what morality demands in the absence of government. According to Locke, in the state of nature (or in the absence of government) people exist in a state of perfect freedom. This partly means that people are free to pursue their own

happiness and well-being. But this perfect freedom is not a license to do whatever one likes or to treat others as one likes. Rather, Locke would understand the liberty we have a natural right to as freedom from domination and coercion.

The state of Nature has a law of Nature to govern it, which obliges every one, and reason, which is that law, teaches all mankind who will but consult it, that being all equal and independent, no one ought to harm another in his life (*The Second Treatise of Government*, Chpt. 2 Sect. 6)

By the moral law of nature, one is not justified in assaulting others except as retribution for an injustice. Likewise, one is not justified in taking another's property except as redress for that person taking or destroying one's own property. But this state of nature inevitably leads to a state of chaos because people are not very good arbiters of justice in their own case. They are prone to inflate the wrongs committed against themselves and seek too much in the way of redress or retribution. And, of course, having a right not to be dominated by those who are stronger or more aggressive doesn't necessarily stop them from violating that right.

Government is justified as the most effective way of securing the natural rights of individuals. In joining civil society, we voluntarily turn our right to protect and enforce our individual rights over to the state. The legitimate function of the state is to secure the equality liberty that people have a natural right to. This view places rather strict limits on the legitimate functions of government. The point of government is just to secure our liberty and its function should therefore be limited to that. Where a government goes beyond this liberty securing role, Locke worries that it becomes an illiberal form of tyranny.

Just what are the rights and liberties government serves to protect? Self-ownership is central to the natural rights equally enjoyed by all. In fact, the idea of self-ownership captures much of how Locke understand individual liberty. This clearly speaks against slavery and other forms of domination or oppression. If a person own's herself by natural law, then she can't also be owned by another. Property rights are then justified as an extension of self-ownership. Locke sees all of nature as initially held in common. When a person "mixes her labor with the stuff of the earth," say, by planting a tree or fashioning a tool from a branch, she acquires a right to the fruits of her labor as an extension of her right of self-ownership. Here, Locke offers a compelling philosophical justification for property rights.

Locke also recognizes limits to the extent of property rights. Specifically, A person does not have a right to more property than they can make practical use of. So, if the apple trees I plant produce more apples than I can harvest and preserve, I have no grounds for complaint when a passerby picks a few for himself. Above and beyond what one can make use of, the fruits of one's labor return to the commons and are to be freely available to others. This limitation on property rights is harder to understand once we introduce a money economy where there are no limits to how much wealth I can set aside for future use and there are many more ways in which I

might put that wealth to use. All the same, it would be hard to argue that anyone can make practical use of say, more than 50 million dollars. Beyond a certain point, wealth is no longer about extending personal liberty, but it may well be about power that restricts the liberty of others. So, a state concerned with protecting the liberties of all equally might have an interest in preventing extreme concentrations of wealth. It should be noted that fans of Locke on the political right are rarely fans of such limitations.

A further limitation on property rights according to Locke holds that the accumulation of private property constitutes no injustice to others “at least where there is enough and as good, left in the commons for others.” Locke takes the natural world and all the resources in it to be a commonwealth. That is, the earth, the waters, skies, and the various systems they contain are taken to be commonly owned by all. I draw from the resources of nature for raw materials when I create something I can then claim as property. As long as there is “as enough and as good” left for others, my accumulation of private property doesn’t limit the liberty of anyone else.

Locke lived in a time when natural resources appeared to be endlessly bountiful and any motivated person who wasn’t happy with the available distribution of property could hop a ship to the new world and homestead a piece of land. Where natural resources can be regarded as practically unlimited, my neighbor’s great wealth doesn’t place any restriction on me investing my energy in creating wealth of my own. But if natural resources are limited and my neighbor has claimed much of what is available in the creation of his private property, then my opportunities are limited to that degree. We can no longer sustain the illusion that natural resources are unlimited. And as we bump up against those limits, Locke’s “enough and as good” proviso becomes much more significant and a problem known as the tragedy of the commons deserves some careful attention.

Locke flourished in colonial times and his legacy bears significant historical stains. On the basis of his theory of property rights, Locke denied that Native Americans held property rights to the land they lived on. He took the vast natural areas in America that were occupied by nomadic tribes to be “wasted,” unclaimed as property because the land was “unimproved” by Native Americans mixing labor with natural resources. He also took land in America to exceed what Native Americans were making practical use of and therefore regarded it as open to settlement and improvement by colonists. It has been argued that Locke failed to apply his own theory of property rights appropriately in failing to recognize Native American labor on the land as establishing property rights. But part of the issue here is a significant problem with Locke’s individualism in his thinking about property rights. There is no room in his theory of property for land that is held and worked on collectively by a community. This is a legitimate concern about Locke’s philosophical views, one that continues to play out in contemporary policies like provisions in NAFTA that undermined collective ownership of land in Mexican villages (the resulting displacement of Mexicans from commonly held land drove a wave of undocumented migration to the US).

A less legitimate complaint against Locke and his liberal philosophy concerns his connections to the slave trade. As a young man under king Charles II, Locke played a bureaucratic role in formulating slavery policies in the colonies and was later granted stock in a slave trading company. Later in life it was his moral reservations about the colonial policies of Charles II that lead Locke to write his political treatises where slavery is antithetical to his core principle of equal liberty grounded in self-ownership. It took two centuries for liberal thinking, established originally by Locke himself in opposition to monarchy, to lead to the abolishment of slavery. A brief overview of this history can be found here: [Does Locke's entanglement with slavery undermine his philosophy? | Aeon Essays](#)

A quick Interlude on Libertarianism

The notion that there is an injustice in funding a social safety net for the less well off with taxes on the more affluent has its roots in a Lockean conception of property rights as natural rights that are closely tied to human liberty. On Locke's view, when we mix our labor with the stuff of the earth, the fruit of our labor is ours by natural right. It is an extension of our natural right to our own selves. Thus, property rights, on Locke's view, are closely tied to human liberty. The contemporary philosopher Robert Nozick extends Locke's line of thought concerning property rights in his entitlement conception of social justice. Nozick's *Anarchy, State and Utopia* gives the strongest argument around for the view we now call Libertarianism. On Nozick's view, any distribution of property and wealth, no matter how unequal, is socially just so long as it was arrived at by just means. Acquiring wealth by one's labor and then building on that through fair trades (those not involving coercion or deception) will be fair.

Taxation beyond what is necessary to keep property rights (and hence human liberty) secure will be an injustice. In fact, it will be a variety of theft. Something along the lines of the views of Locke and Nozick has inspired a good deal of the anti-tax, small government sentiment in U.S. politics over the past several decades. Liberty is seen as closely tied to property rights. To the degree that the government taxes citizens, it takes their property and thereby limits their freedom. It is worth noting that Nozick's view goes well beyond Locke in stripping limits from property rights. While libertarian thought offers principled grounds for accepting extreme inequality, it offers very little in the way of addressing the social instability that can result from this. It is also worth noting that while libertarianism accepts extreme inequality and rejects social safety nets for the less fortunate, it also provides a strong case for reparations to Native Americans and the descendants of slaves. Inequalities that result from legitimate acquisitions and transfers of property are endorsed by Libertarianism. But a significant amount of the wealth in this country is the result of the illegitimate plunder of Native American land and slave labor. According to libertarianism, the current holders of this wealth have no property right in it, but Native Americans and the descendants of slaves do. Ta Nehisi Coates notes this in prefacing his famous essay, [The Case for Reparations](#), with a quote from Locke.

An Answerable Concern for Locke: The Tragedy of the Commons

Garritt Hardin is well known for his clear articulation of the Tragedy of the Commons in the late sixties. Hardin was mainly concerned about human population growth, but this is just one instance of a much broader kind of problem. A tragedy of the commons is any case where some commonly held resource gets exhausted to the point where it has little value left to offer. Such a tragedy is bound to occur eventually whenever a commonly held resource is finite and freely utilized by self-interested agents.

Hardin introduces the notion of the tragedy of the commons with a tale about the fate of shepherds who share a pasture in common. Each shepherd notes that if he runs one more animal on the commonly held pasture, he will get the full benefit of that animal's value when he takes it to market, but since the pasture is held in common, he will bear only a fraction of the cost of raising the animal. As a result, each shepherd finds it to be in their self-interest to run an additional animal on the pasture, and then another and another until the commonly held pasture is depleted to the point where it is of no use to anyone. This dynamic is at work in a broad range of issues including fisheries, fresh water supplies, air pollution, and climate change.

Once we have a clear understanding of the logic of the commons, it is equally clear that there are only a limited number of ways to avoid a tragedy of the commons. Again, a tragedy of the commons is the inevitable result whenever we have a finite commonly held resource that is freely utilized by self-interested agents. The only way to avoid a tragedy of the commons is to prevent one or another of the three conditions that give rise to one. Perhaps we cannot expect individuals to consistently refrain from acting on their own interests. But there remain the possibilities of expanding the commons in some way, regulating access to the commons or otherwise reducing use of the commons.

Fish hatcheries provide a good example of how we might expand the commons. A degree of overfishing can be sustained if we also breed fish to replenish the diminished supply. But in many cases, notably including climate change, strategies for expanding the commons aren't sufficient for avoiding a tragedy of the commons. Given this, we have only one route to avoiding a tragedy of the commons, and that is regulating the use of the commons.

There are more and less obtrusive ways to regulate or reduce use of the commons. Requiring pollution controls on cars has been a very unobtrusive way to regulate pollution and avoid an air quality tragedy of the commons. Fees at campgrounds and catch limits with fishing licenses also regulate use of commons effectively and without imposing much of a burden. Imposing taxes on the use of the commons or providing incentives for not using the commons are useful strategies. So is seeking technological alternatives to the use of a commons.

The most challenging tragedy of the commons situation we have confronted so far is climate change. In the case of climate change, the Earth's atmosphere is used as a commons for absorbing CO₂ emissions that result from burning fossil fuels. For several decades now, attempts

to reduce CO2 emissions through variations on each of the strategies just mentioned have been proposed, but little progress has been made. We are beginning to suffer the consequences of our failure to adequately address this tragedy of the commons. Things will no doubt get much worse as governments, companies, and citizens continue to seek creative ways to bring down CO2 emissions.

Once we have a clear understanding of the logic of the commons, it is pretty obvious that some form of regulating the use of commonly held resources is often called for. We are very prone to think of government regulation as an imposition on our liberty. But the destruction of commonly held resources can impose a much greater imposition on our liberty. So, regulation to avoid a tragedy of the commons is quite in line with Locke's view that the legitimate role of government is to secure out liberty. This should help us understand why conservatives as well as liberals were proponents of environmental regulation through the 60s and the 70s. A better understanding of how tragedies of the commons happen and how they can be avoided should help us cooperate more effectively. It would be a serious mistake to see Lockean love of liberty as an obstacle to effective action. The tragedy of climate change is no friend of liberty.

Here is a link to Garritt Hardin's article, "The Tragedy of the Commons" [E:\twelve-one\tragedy.PDF \(garretthardinsociety.org\)](E:\twelve-one\tragedy.PDF (garretthardinsociety.org))

An Objection to Locke: The System

Now we will consider an objection to Locke's political philosophy, one that will set the stage for our discussion of John Rawls. The political thought of Locke is highly individualist. One aspect of this is that Locke takes the rights and liberties of individuals to be the only thing that matters in political philosophy. Beyond this, Locke would limit the legitimate role of government to securing those individual rights and liberties. Is this enough to secure justice?

Even when our rights and liberties are secured, our lives are significantly affected for better or worse by what I'll refer as "the system." So, what is this system? The system we have in a society is constituted by various subsystems like the market economy, our tax system, our education, health care and environmental management systems. As Rawls sees it, government has important work to do in establishing and upholding systems that are fair. Here Rawls goes beyond Locke. Rawls political philosophy still prioritizes individual rights and liberties. But Rawls sees something more at play. Justice in a society also depends on having a fair system. So what is it for a system to be fair? This is the question Rawls sets out to answer. But before we get to that, let's explore the system some and how it operates.

Anymore, we don't create wealth from our own labor and ingenuity in a social vacuum as Locke suggests. With the possible exception of the vegetables I grow in my garden, none of my wealth is entirely the product of mixing my labor with the stuff of the earth. Rather, nearly all of our productive activity is carried out in the context of a complex fabric of social structures and interrelations buoyed by a substantial technological infrastructure. Enjoying the fruits of my

labor nearly always requires doing business with someone else and what I get out of this depends as much on the favorable social environment and technological infrastructure that makes doing business possible as it does on the efforts I bring to the table. In light of this, the view of property rights offered by Locke is unrealistically individualistic in that falsely assumes property and wealth is the product of individual effort alone.

Having a functioning well-ordered community is a necessary condition for succeeding in every line of business (even gangsters depend on the system they exploit). The businessman who has profited from a fair exchange with his customers depends on the underlying system that makes it possible for him to do his business in the first place. His success may require a healthy and well-educated workforce, stability in the economic system, a citizenry that is well informed enough to politically sustain just social institutions, a citizenry that respects the law, a customer base that is doing well enough themselves to afford his product and so forth. It will also depend on physical infrastructure we rely on the government to provide or at least regulate. Roads and bridges to sewer systems are a few examples. Utility companies are often private businesses, but these require government oversight and regulation since they constitute natural monopolies. All of these things and others are hidden ingredients in the wealth created through the businessman's activities. Given this insight into the kind of sophisticated market economy we have, taxation for the maintenance and upkeep of our various social and infrastructure systems is not theft, but fair compensation for the benefits we derive from participating in the system we are productive members of. Limiting the role of government to the protection of rights and liberties turns a blind eye to need for government involvement in the upkeep of the various systems we all depend on.

The Locke/Nozick approach to social justice where the legitimate activity of government is limited to securing our natural rights and liberties is roughly analogous to field biologists aiming to secure the well-being of squirrels without giving any regard to the health of the ecosystem that sustains them.

Given some sense of the importance of the various systems we rely on government to uphold and maintain we can turn our attention to Rawls and the concept of fairness.

Rawls: Justice as Fairness

Where all other things are equal, we are liable to think of fairness in terms of equal treatment. The fair way to divide a cookie between three equally hungry boys is into three equal sized parts. But in the real world, things are seldom equal.

Consider the case of Jones and his three sons. Jones has done well in life and build up a modest fortune of 1.5 million dollars. But now his life is coming to an end as he enters the late stages of terminal cancer. So, he now faces with the question of how to divide his modest fortune among his three sons. Again, all other things being equal, the fair thing to do would be to give each a half million. But things aren't equal. Jones' eldest son, John, is a brilliant young man. He is near the end of a PhD in computer science at Stanford and he is being recruited for a high paying

Silicon Valley job with Google. Jones second son, James, is a star athlete on a full scholarship at UCLA where he is majoring in business. In addition to being talented on the basketball court, James is highly popular, well regarded for his forthright and easygoing manner, trusted and liked by peers and superiors alike. The third son, Joe, well, he's a nice guy. He has no particular talents. He has tried hard in school, but hasn't done especially well. He's a bit awkward socially, the sort of guy that just isn't going to find a date for the prom. In addition, he has a fairly expensive lifetime medical condition, he's type 1 diabetic. What would be the fair way for Jones to divvy up his modest fortune among his three sons?

Jones can't just pull a clear answer out of thin air, but he is clever and has a friend who does research in neuroscience that can help. His friend has developed a drug for highly selective memory loss that temporarily mutes a person's sense of personal identity. Jones administers the drug to one of his sons (it doesn't really matter which one) and instructs that son to decide how his fortune is to be divided among the three. He offers the following instructions: "you are one of my three sons and you are fully informed about the life circumstances of all three. Now, with just your own self-interest in mind, you are to decide how my estate is to be divided between you and your brothers." The son who decides how the fortune is divided must do so without knowing which of the brothers he is. The effects of the drug render the son who decides incapable of favoring the interests of one son over another, and this gives us some grounds for thinking his decision will be reasonably fair. Let me know how you think the drugged son will decide.

What we've described here is similar to the method Rawls recommends for selecting fair principles of social justice. Rawls is aiming at a conception of justice as fairness in the sense that social systems won't advantage any particular kind of person at the expense of others. Rawls' proposes that we can get onto the ideal of justice as fairness by means of a thought experiment that involves reasoning from what he calls "the original position." In the original position, we imagine that we are perfectly rational agents with full information about the consequences of the various possible social arrangements. We are then given the task of designing the principles of justice that will structure our society and we are expected to do so with an eye to what will be in our own best interest. But there's a catch. In reasoning from the original, we operate behind a veil of ignorance about our own personal circumstances and characteristics. So, in the original position, behind the veil of ignorance, I must think about what set of social institutions will work out best for me without knowing whether I am weak or strong, healthy or diseased, clever or dull, beautiful or ugly, black or white, born to a wealthy family or a poor one and so forth. If I am rational and self-interested, I will want to set things up so that I can enjoy the benefits if I have characteristics that are highly valued in my society and I put them to good use. But at the same time, I will want to hedge my bets to assure that I still have a good life if I am not so lucky or if my best efforts fail.

Of course, the original position thought experiment is just that, a thought experiment. No one could actually place themselves behind the veil of ignorance, nor reason perfectly rationally about all the possible social arrangements that might result from her choice of principles. Still

Rawls has devised a way to think about what is fair, and we can apply this to approximate relatively impartial judgments about what a fair society would look like.

On the basis of the original position thought experiment, Rawls argues for two principles of justice as fairness:

The Equal Liberty Principle: Each person is to be granted the greatest degree of liberty consistent with similar liberty for everyone.

The Difference Principle: Social practices that produce inequalities among individuals are just only if they work out to everyone's advantage and the positions that come with greater reward are open to all.

The Equal Liberty Principle has a longer history. The idea that everyone should be granted the greatest degree of liberty consistent with similar liberty for others is defended at length in John Stuart Mill's essay [*On Liberty*](#). In fact, we could take some variation on this principle as the core tenet of Liberalism as a political theory. This principle doesn't tell us that people should be free to do as they please no matter what. At some points, my being free to do something is liable to interfere with your being free to do something. For instance, my being free to host parties with live bands into the early hours of the morning might interfere with my neighbor's being free to get a decent night's sleep. In the interest of maximizing equal liberty for all, we would be justified in restricting people from activities that would interfere with the liberty of others. This has many familiar applications. Neighborhood zoning regulations are one example. A good deal of environmental regulation illustrates this idea. Maximizing liberty for all equally requires that we restrict businesses from being free to pollute where doing so would adversely affect the health of others.

Rawls thinks the equal liberty principle will be selected by rational agents reasoning from the original position because no rational agent in this position would choose to be less free than necessary nor grant some (possibly someone else) greater liberty than others (possibly herself).

The Equal Liberty Principle is only concerned with equality of liberty. But we can be equal or unequal in many other ways. In fact, being equally free is liable to lead to other sorts of inequalities. If we are all free to plant apple trees as we see fit, we will probably wind up with an unequal distribution of apples simply because some of us will plant more trees and do a better job of tending them. So long as this is merely the result of people exercising their equal liberties, there is nothing unfair about this. If I'd wanted more apples, I could have spent more time growing apple trees and less time playing chess.

Given equality of opportunity, the Difference Principle holds that a system that produces inequalities is fair so long as it works out to the benefit of all. This will strike some as puzzling. How could we have inequalities that benefit all? It might seem that liberty and the fact that we want different things can account for some of this. I can't claim unfairness when my neighbor

has more apples because I'd rather play chess than harvest apples. The inequality here is one that serves the interests of both of us, my neighbor's interest in having lots of apples and mine in playing chess. But the difference here isn't the result of any system, fair or unfair. It's just the result of our individual choices. What we are concerned with in connection with the difference principle are inequalities that result from our social systems, not just the free choices of individuals.

So, we want some clarity on how there could be inequalities resulting from a system that work out to the benefit of all. Here it will help to get familiar with the idea of a zero-sum game. A zero-sum game, as the term suggests, is a game where the sum of the winnings of the winners and the losses of the losers add up to zero. A friendly neighborhood poker game is a zero-sum game. The combined winnings of the winners will be exactly equal to the combined losses of the losers. Likewise, slicing a pie is a zero-sum game. Ordinarily, the only way I get a bigger piece is if others get smaller pieces where the difference between how much bigger my piece is and how much smaller other pieces are is zero. But not every game is a zero-sum game. To take a fanciful example, suppose the pies I bake are magic pies that are favorable to me in a specific way. When I cut myself a bigger piece, the size of the whole pie doubles. Now, when I cut myself a bigger piece, I get more than other people who want some pie, and the others get more than they would have if I'd taken a smaller slice. So, this is a positive non-zero-sum game, but we can easily describe a negative non-zero-sum game. Suppose my method for dividing a cookie between three hungry boys is to just toss it on the table and have them grab at it. They will probably wind up with unequal size chunks of cookie, but the sum of the winnings of the winners and the losses of the losers will add up to less than zero because a good part of the cookie will wind up as crumbs on the floor. So, a social system that operates as a positive non-zero-sum game can, under the right circumstances, both systematically produce inequalities and yet leave the less advantaged better off than the less advantaged would be under a different system. Such a system will satisfy the difference principle.

A market economy can provide a good example of a positive non-zero-sum game that would satisfy Rawls difference principle. This will depend on how the market economy functions in the context of the broader system of social policies and institutions. But let's first consider Adam Smith's notion of the invisible hand. Adam Smith is well known as the founder of economics, but he was also a notable moral philosopher who did important work on the emotions. We'll be concerned, though, with one of his key economic insights, one that is often touted as making the case for free markets. First, we need to describe a scenario. Let's imagine that we are all part of a small low-tech economy where one of our technologies is making matches. The standard method for making matches to light fires is to sliver off individual sticks with a knife and dip them one by one in the solutions that make them ignite easily when struck. In this economy we have a free market, meaning that we have property rights and the freedom to set out own prices and negotiate our own deals with customers. I, being clever and self-interested, devise a way to dramatically reduce the amount of labor I have to invest in each match by building a jig that

allows me to split a small log into a hundred little sticks at once and then hold all hundred as I dip them together in the various solutions. Because I have dramatically lowered my labor cost, I can take my matches to market, charge less for them than my competitors and dramatically increase my market share while still making a tidy profit. I get rich as a result, probably driving several of my competitors out of the match making business. But my competitors, also being clever and self-interested, will find other ways to contribute to the village economy. Perhaps one starts growing flowers and becomes the first florist and another discovers yeast and invents bread. And folks in the village can afford to pay for flowers and bread now that they are paying so much less for matches. In this happy little story, competition in a market system generated an inequality when I got rich thanks to my clever innovation in match making. And the community as a whole benefits because now we are investing much less of our collective energy into making matches. Instead, we have cheaper matches plus bread and flowers. We are all better off as a result. In this story, the private vice of greed leads me, as if by an invisible hand, act in a way that benefits the community.

We should note that contrary to myth, Adam Smith notion of the invisible hand was not intended as an endorsement of greed or *laissez-faire* capitalism. In a poorly regulated market economy, greed is just as liable to lead to predatory behavior as productivity boosting innovations. But all the same, under the right conditions, the profit motive can bring out the best in people and result in wealth creation that benefits all. We can look to the historical example of East and West Germany prior to the fall of the Berlin Wall. Communist East Germany had a Soviet style command economy that dictated a high degree of equality and sharply restricted property rights. West Germany had a mixed economy with a healthy market based private sector. The result was less equality in West Germany. But still, the least well off were substantially better off than the least well off (pretty much everybody) in East Germany. What this indicates is that while, on the one hand, the market economy with its institutions of private property inevitably generates inequalities, it is so much more effective at generating wealth that the least well off can be better off in spite of the inequalities.

But this is likely to happen only under the right conditions. West Germany had then, as Germany has now, a mixed economy where a vibrant private sector market economy is supplemented with a substantial system of social benefits. The profit motive is harnessed and the well to do are quite well off, but taxes are higher and those who don't fare well in the market economy still get decent health care, modest but safe and comfortable places to live, access to education etc. Compare this America where the private sector market economy dominates and the least well off are much worse off than the least well off in communist East Germany. Again, it is ultimately the system as a whole that gets evaluated according to the difference principle.

Hopefully it is clear now how Rawls' view of justice as fairness would endorse taxing the well off in order to provide for things like education, health care and a social safety net. His is a more expansive view of the role of government than Locke or Nozick would support. But to what degree will taxing the rich be justified on Rawls' view? Clearly taxing the successful members of

society to the point where they are no better off than those who are largely unproductive will not meet the Rawlsian ideal under the difference principle. If people are not rewarded for hard work and innovation, then it's liable not to happen and everyone suffers as a result. Too much taxation of the well-off will be unjust on Rawls' view precisely because it doesn't work out to benefit of least well off (or anyone else). Rawls would deem communism and socialism as unjust under the difference principle for just this reason (not to mention the ways these systems undermine equal liberty). Rather, Rawls would aim for that sweet spot where hard work and innovation are well rewarded, so everyone has a reason to do their best, and yet those who fail for whatever reason are not left by the wayside but still have opportunity and enjoy some security and quality of life.

Rawls represents the "liberal" side of liberal political philosophy. In some corners of popular political discourse, the term "liberal" has been reduced to a term of derision and its meaning has been so badly distorted as to make it unrecognizable. Liberals are routinely described as socialists or communists by a fair number of political pundits. Hopefully some acquaintance with Rawls will help you can recognize the lie in this rhetoric. Liberals, as the name should suggest, support liberty. Communists do not. Socialism traditionally refers to public ownership of the means of production, as opposed to having a market economy with property rights. Mainstream liberalism as represented by Rawls political philosophy, endorses a market economy, just one with guard rails and safety nets to correct for some of the excesses of *laissez faire* capitalism.

Illiberal Political Philosophy

When I initially wrote this text, my aim in this chapter was to explain the philosophical underpinnings of the two dominant mainstream traditions in American politics. As we've discussed, both what we popularly known as liberalism and conservatism are positions within a broader tradition of liberal political thought. We can understand a political philosophy as part of that broader liberal tradition when it gives priority to the equal rights and liberties of individuals as a political value. Illiberal political philosophies will just be those that reject the rights and liberties of individuals of paramount importance. I originally prefaced this chapter with a discussion of Plato to provide some historical perspective, some appreciation for how recent an innovation it is to treat individual rights and liberties as important in political thought. Over the past few years, though, the broad tradition of liberal political thought that has guided this country since its inception, in both its "liberal" and "conservative" variants, has come under threat. So, a few cautionary words about illiberal ways of thought are in order.

Political observers have recently heard lots of talk about authoritarianism, populism, nationalism and other assorted "isms." I'm not going to take up a detailed analysis of these here, but I do want to address an underlying current common among them. Authoritarianism, as the term suggests, prioritizes the will of an authority figure. But for an authority figure to gain and sustain power, he must have the support of a sizable chunk of the population. So, would-be authoritarians will need to appeal to the concerns of ordinary people and make themselves popular and this is what populism is. So, populism can lead to authoritarianism, though that

depends on the concerns of the people. If ordinary people care most about individual rights and liberties and have some ability to defend themselves against the rhetorical trickery of a demagogue (a leader who appeals to people through emotion and prejudice rather than rational argument), then populism won't provoke a turn away from individual rights and liberties. But people have concerns beyond individual rights and liberties and these can eclipse the tenets of liberalism.

We'd be hard pressed to explain how popular opinion could turn against the broad tradition of liberalism if varieties of illiberalism had nothing to offer people. Mass movements in support of nationalism (which prioritizes national interest over individual rights and liberties), offer the powerful appeal of a shared identity and the social cohesion of a common cause. Indeed, one of the classic criticisms of liberal political thought is that it fails to provide shared ideals that can be the basis of a sense of shared identity, purpose and community. Liberty alone is thin gruel for those seeking a sense of meaning and purpose in life in the public or political sphere. And prioritizing liberty as a political value requires maintaining neutrality on a broad range of other values and conceptions of the good life. Pushing a specific further set of values in the realm of politics as the basis of community and shared identity is bound to marginalize and threaten adherents of other ways of life. But the whole point of liberal values like freedom of conscience is to avoid this.

Like nationalism and some kinds of populism, Communism also held the appeal of a shared identity based on commonly held values. In direct opposition to liberal traditions, both communism and nationalism prioritize the good of a collective over concern for citizens as individuals. Fairly recent history is rife with examples of how collectivist thinking, both on the left in the form of communism and on the right in the form of nationalism, have licensed extreme brutality. I won't pursue historical examples or details here. But it should come as no surprise that prioritizing collectives over individuals is liable to be pretty hard on individuals.

What I do want to say in the way of caution concerning collectivist ideologies of all stripes is mainly that collectives don't suffer. The very idea of a collective is an abstraction. Collectives have no existence beyond the individual members that make it up. And so, it is hard to see how a collective can have any value of its own. My intuition here is roughly Kantian. People have intrinsic moral worth, not nations. Advancing the interests of people has nothing to do with pitting the interests of one nation against those of other nations. America is not a human being. But we might consider making America humane again.

Review and Discussion Questions

1. Explain Plato's conception of justice. How does it differ from contemporary liberal views of justice?
2. Why is Plato no fan of democracy?

3. In what way does Locke see us as having a natural right to liberty?
4. What is the function of government according to Locke and how is it justified?
5. How does Locke justify property rights?
6. Explain the tragedy of the commons and how it provides a rationale for regulating use of commonly held resources.
7. How does the creation of wealth and property in our society differ from the idealized individualistic conception Locke offers?
8. Explain Rawls' two principles of social justice as fairness.
9. Explain Rawls original position thought experiment and how it gives us a way of thinking about fairness?
10. What is a zero-sum game and how does the idea of a non-zero-sum game help us understand Rawls' difference principle?