

Knowledge For Humans

KNOWLEDGE FOR HUMANS

CHARLIE HUENEMANN



Knowledge For Humans by Charlie Huenemann is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/), except where otherwise noted.

CONTENTS

Introduction	1
--------------	---

Part I. 1. Knowledge: The Basics

1. Knowledge	5
2. Truth	11
3. What Was Francis Bacon Talking About?	15
4. Questions to Consider	19
5. Further Reading	21

Part II. 2. Skepticism: A Dialogue

6. Questions to Consider	47
7. Further Reading	49

Part III. 3. The Grand Deception Doubt (Or, the GDD)

- | | |
|--------------------------|----|
| 8. Questions to Consider | 57 |
| 9. Further Reading | 59 |

Part IV. 4. First Answer to the Skeptic

- | | |
|---|----|
| 10. Knowledge of Appearances, or
Phenomenalism | 65 |
| 11. Problems with Phenomenalism | 81 |
| 12. Conclusion | 88 |
| 13. Questions to Consider | 89 |
| 14. Further Reading | 91 |

Part V. 5. Second Answer to the Skeptic

- | | |
|---------------------------------|-----|
| 15. Internalism and Externalism | 95 |
| 16. G. E. Moore's Hands | 102 |

17. Epistemology Naturalized	105
18. Back to Bacon	113
19. Questions to Consider	116
20. Further Reading	118

Part VI. 6. Scientific Knowledge

21. Logic, Math, and Science	121
22. How Do We Know the Laws of Nature Are True?	129
23. But How Do We Know the Laws of Nature Are True?	136
24. Social Conditions of Scientific Knowledge	143
25. Questions to Consider	146
26. Further Reading	147

Part VII. 7. Social Conditions of Knowledge

27. Observations Are Theory-Laden	153
28. Pernicious Background Theories	158

29. Morality of Knowledge	163
30. The Open Society	168
31. Questions to Consider	175
32. Further Reading	176

Part VIII. 8. Knowing Our Weaknesses

33. Seven Fallacies of Highly-Human Thinkers	187
34. Reflection on the Fallacies	198
35. Questions to Consider	204
36. Further Reading	205

Part IX. 9. How to Argue with Other People

37. Arguments as Occasions for Learning	213
38. Being Fair, and Even Generous	217
39. Friendly Adversaries	221
40. Helping the Discussion to Continue	226

41. Questions to Consider	229
42. Further Reading	230

Part X. 10. Bayesianism and What Is Likely

43. David Hume and Miracles	233
44. Bayesianism	243
45. Questions to Consider	255
46. Further Reading	256

Part XI. 11. Epistemology and the Internet

47. Information in Historical Context	263
48. Algorithms	268
49. Knowing through the Internet	274
50. Some Advice	280
51. Questions to Consider	288
52. Further Reading	289

Part XII. 12. Conspiracy Theories

53. The "Seed Belief" Model	293
54. Monocausal Explanation, and <i>Cui Bono</i> Inferences	297
55. Where Seed Beliefs Come From	302
56. Real Conspiracies	308
57. Questions to Consider	313
58. Further Reading	315
Appendix	317

Introduction

Is it possible to exaggerate the importance of knowledge? We can of course disagree about whether it is important to know the history of astrology, or the grammar of dead languages, or the mating habits of warblers, but only a stubborn blockhead would insist that it is not important to know. One might even be said to forfeit one's humanity by denying the importance of knowledge.

Given its extraordinary importance, it follows that it would be good to know what knowledge is, how to gain it, how to be sure one has it, and its effects upon individuals, political states, and societies, and so on. For many of these questions we should turn to philosophers. But unfortunately quite a lot of these questions have so fascinated academic philosophers that their discussions of them have spiraled off into inaccessible regions of forbidding jargon. Moreover, many of these discussions are bound to a single discipline, as if there is only one set of questions one should ask about knowledge. It is hard for a curious human being to know where to go to get started in understanding knowledge in some more expansive fashion.

Hence Knowledge for Humans, meaning knowledge for humans who are intelligent and curious, but have not yet been

shunted into specialized regions of abstract scholarship. The idea behind this text is to offer some introductory philosophical discussions about knowledge combined with some attention to science, history, media, politics, and psychology. It is meant to pull together different aspects of knowledge into a package that a philosophically curious reader might find interesting.

I wish to thank my students for reading through the text and offering feedback, to my friend and colleague Professor Richard Greene for doing the same, and to the Open Educational Resources team at Utah State University who helped to put the text into an accessible form. Any errors in the text are due to my own ignorance, appropriately enough.

PART I

1. KNOWLEDGE: THE BASICS

“Knowledge itself is power.”

— Francis Bacon (1597)

Most of the time we want knowledge. Knowing means understanding what is true and perhaps being able to shape events toward our own ends. Ignorant people are generally not esteemed (though sometimes we might envy the person who is in a bad situation but does not know enough to be troubled by it; that’s when we say, “Ignorance is bliss”). We seek knowledge both for the advantages it gives us and even for its own sake. In a great many cases, if not in every case, it is simply *better* to know than not to know. Philosophers sometimes have said that the pursuit of knowledge is essential to human nature. “All humans by nature desire to know” is how Aristotle begins his *Metaphysics*.

Because knowledge is so important to us, we fight over it. One group claims that X is true, another group insists that X is false, and a fight breaks out at least in words, if not in fisticuffs, over whether X is true—which is to say whether it should count as known. Authoritative institutions often seek to control knowledge both in terms of who can have it and what should count as *known*. We often turn to science as an institution to decide whether claims are known or not, which gives “science” – meaning *scientists*, at universities and in research labs – a huge amount of power. Each year the world spends a trillion dollars on research and development in science which means there is a lot of money we are willing to spend on figuring out what is known.

So, what is knowledge? What is it to know something? Philosophers usually start to answer big questions like this by making distinctions and making the discussion more precise.

1.

KNOWLEDGE

First, let's ask what sorts of things are "known" in the sense we are interested. In a claim of the form "Sluggo knows X," what sort of entity goes into the "X" spot? We could fill in the blank with phrases like "how to ride a bike," "how to speak Esperanto," or "how to dance like a maniac." These are abilities Sluggo may or may not have, and we might call this sort of knowledge *know-how*. Philosophers typically are not very interested in know-how since we have not yet come across many interesting questions or problems to raise about it. Instead, we could put into the "X" spot phrases like "that Bangladesh and India share a border," or "that Jupiter is larger than Mars," or "that Francis Bacon is credited with having said that knowledge is power." For obvious reasons, this is called *knowledge-that*, and what is known, or what goes in the "X" spot, are claims that may be true or false which philosophers call "propositions." Philosophers are very interested in propositional knowledge.

So, let's continue to focus on propositional knowledge. We

will want to know what features a proposition has to have in order for us to rightly say that someone *knows* it. We have already seen one back when we were observing that people sometimes fight over knowledge or *truth*. A claim to *know X* is, among other things, a claim that *X* is *true*. No one would ever claim to know something they believe is false. (Well, they might if they were lying. But if they were not trying to lie, and they were well aware of what they are saying, and they were not under the influence of some strange drug or head injury, they would never say, “I think *X* is false and that I know it.”) Furthermore, we would not claim that someone else *knew* something if we thought that that something were false. I would not say that Sluggo *knows* the world is flat since the world isn’t flat. Sluggo might *believe* it is, but he does not *know* it because it isn’t.

This is a tricky point, so we should spend a little more time on it. Some time ago, people thought the sun revolved around the earth. They had good reason for believing this since it certainly looks like the sun revolves around the earth at the rate of once per day, and we do not feel the earth to be in motion. But would we say these people *knew* it? They believed it, yes, and they had their reasons, yes, and they *thought* they knew it, yes. But did they know it? We might say they “knew” it, but it would be important to keep those scare-quotes attached

because they did not *really* know it. They did not know it because the claim is false, and they were wrong.

(Or so it seems to me. Perhaps my use of the word “know” has been corrupted by reading too much philosophy. But at the very least there is one widespread meaning of the word “know” which implies that what is known is in fact true, and in what follows, that is the sense of “know” we will be using.)

This discussion also suggests a further feature of propositions we claim to know: we must also *believe* them, or in other words, *believe they are true*. There may be some cases of people really knowing something deep down in their bones but not admitting it to themselves. Perhaps these are cases when someone knows something and, in a certain sense, does not believe it. But let’s mark these cases as special exceptions to the more general rule that we believe the things we know.

From what we have seen so far, knowledge is *believing true propositions*. But there is a further feature we need to add. Knowledge usually involves some amount of evidence or reasoning or reliable report. It is not a lucky guess, like correctly predicting a coin toss. We can call this general feature

“justification,” meaning that if we really do know some proposition, we could provide some reasons or evidence for the truth of that proposition.

And so we arrive at a time-honored definition of knowledge: *To know a proposition is to have a justified belief in that proposition and for that proposition to be true.* If I believe it and have some good reason for believing it, and if it is true, then I know it. For obvious reasons, this is called the “justified true belief” definition of knowledge, and in fact it is so common that it is sometimes just referred to by its initials: **JTB**.

But philosophers love to test definitions against their own imaginative powers, and so challenges have been posed to JTB. Here is a case where the letter of the definition is met, but intuitively, it does not seem to be a case of knowledge. Suppose that Molly walks into her bedroom one night and flips on the light switch, but the light does not come on. She tries the switch a few more times, but still the light does not come on. She considers that the light bulb has been in use for a long time and that it is about time for it to have burned out. Just to be sure, she checks to make sure the bulb is screwed securely into its socket. It is. “Ah,” she thinks. “The light bulb must have burned out.” Let us also suppose that, in fact, she is right:

The light bulb has burned out. But what she doesn't know is that Doug is in the basement, and he has turned off the circuit breaker for her bedroom. So the light bulb would not have come on even if it had not been burned out. Does Molly's belief that her lightbulb has burned out count as knowledge?

She believes the lightbulb has burned out, and she is right about that. She also has at least some justification for her belief since she has tried the switch a few times. She remembers how old the bulb is and correctly infers that it would be likely for the bulb to have burned out at this time. She confirms that the bulb is screwed in securely. So she meets the JTB conditions. But in a very important sense, Molly has simply *lucked out* this time. She *accidentally* got to the right conclusion, and her belief does not really count as knowledge.

To try to repair the JTB account, we might further require that the truth of the proposition is *not accidental* to the person's justification for their belief. In Molly's case, as we said, it was lucky for her that the bulb was in fact burned out; its being burned out did not actively contribute to the evidence she was gathering for her belief since all of that evidence would have been the same anyway even if the bulb was in good working order due to Doug's interference. Another way to make this

point is to say that, in a case of knowledge, the truth of the thing being known *helps to explain* why the person believes what they do.

With this in mind, we might offer the **JTB+** definition of knowledge: S knows P if and only if S believes P, S's belief has justification, P is true, and the truth of P helps to explain why S comes to believe P.

Of course, much of this is still unclear. How much justification does one need? Will any old justification be sufficient, or must the justification be of a certain kind? What does “helps to explain” mean? Is it okay to help a little, or must it be a lot? And even if we manage to make these matters clear, can we be sure that there aren't some further clever challenges to JTB+ arising from the fertile imaginations of philosophers?

But we will leave these worries and unclariities to the side, at least for now, and be content to have given at least a set of features central to cases of knowledge. However, we should explore this “truth” business a little further.

2.

TRUTH

At this point someone might raise an objection:

Objection: You have claimed that knowledge implies truth—in other words, in order to know something, that something must, in fact, be true. But this makes the definition entirely useless. If we have to know what's true in order to figure out what we know, then why bother with any definition for "knowledge" in the first place? Why not just rest content with the truth?

This is a good question. It forces us to become clear about what we want in a definition. Sometimes we want to use definitions to help us sort things into categories. Consider, for example, the standards set forth by the American Kennel Association for figuring out when this or that animal is a member of this or that breed. In this case, we want the

definition to act as a sorting mechanism to help us decide whether this or that thing should be called whatever it is we are defining. But other times, we simply want a definition to tell us in a more general way how the target concept relates to other concepts. This increases our understanding of the concepts, though may not decisively settle any disputes.

The JTB+ definition of knowledge is definitely the second sort of definition. We have seen that knowledge is related to belief, truth, justification, and explanation, though this has not brought us any closer to being able to assemble some sort of “litmus test” we can use to determine which of our beliefs should count as knowledge. The effort to assemble such a test has a long history in philosophy; it is the effort to refute skepticism. We will discuss that at length later on.

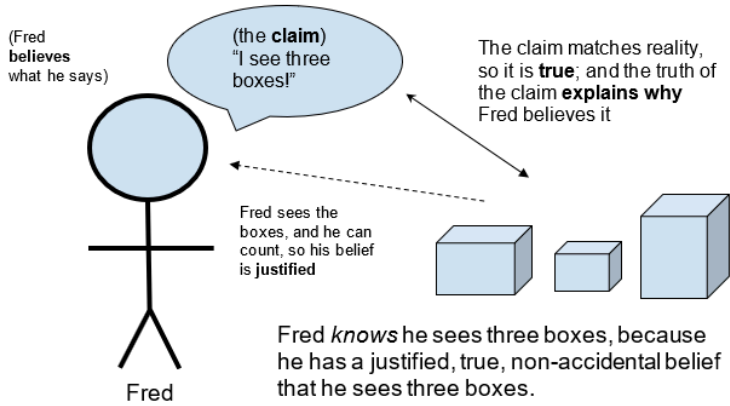
But let us pause over *truth* a bit longer. What is *truth*? This is the sort of question philosophy is famous for, and one might expect a very impressive and mysterious answer like “Truth is beauty” or “Truth is what releases us from ourselves” or something else. These are interesting claims to reflect on. But in fact, philosophers typically rely on a simple and straightforward meaning of truth. First, we need to ask again what *sorts* of things are true. Are we talking about people,

concepts, neutrons, or what? Once again, it is propositions or sentences that we say are true or false. So what makes a proposition true? Here is the simple answer: a proposition is true if and only if it describes how things really are. When a proposition matches reality, the proposition is true. That's it.

Objection: Once again, this seems like a useless answer. How do we know what reality is? And without knowing what reality is, how can we determine which propositions are true?

These are great questions, and they are at the foundation of epistemology. If we want knowledge, we want to know what's true, or what reality is, and how we should go about discovering what reality is, and what we should do when we are not completely sure what's true. That's what this book is about.

Here is everything we have said so far summarized in cartoon form:



Media Attributions

- Figure 1.1 © Charlie Huenemann is licensed under a [CC BY-SA \(Attribution ShareAlike\)](#) license

3.

WHAT WAS FRANCIS BACON TALKING ABOUT?

The quote at the start of this chapter says that knowledge itself is power. That sounds mightily impressive. What does it mean?

One thing it might mean is that if you have knowledge, you will have power. Maybe it won't be a lot of power, if what you know is trivial. But if you know an important secret, it may give you some power. Certainly a lot of the power humans have is due to their scientific knowledge. If knowledge implies truth, then having knowledge means knowing what reality is. Once you know that, you will be much more successful in your endeavors than you would be if you didn't know how reality is. Reality punishes ignorance with failure.

But claiming that knowledge *itself* is power might also mean something else. It might mean that knowing how reality is—just the knowledge of it, never mind doing anything with that knowledge—somehow improves our condition as beings who are able to think. It is good to know, and not just because knowledge will make us more successful, but because, well, it is just good to know. Many things are simply *intrinsically* good (meaning they are good in and of themselves). It is good to love and to appreciate poetry and to have fun and to have friends and to *know*. These things make us better human beings because they are part of living good human lives. Aristotle called knowledge a *virtue* which just means that it is a good thing for a human to have because the human is human. A virtue is a power in the sense of being a capacity for a specific kind of being. Humans are the sort of being that is capable of knowing; it is a power we have. Indeed, one might say it is a human superpower since we have not yet met other sorts of beings who are better at it than we are.

There is a third thing the claim might mean. Because of the first two meanings (the ones about success and virtue), human societies value knowledge. Generally, as someone gains more knowledge, they ascend in social status. The two are not always tied, of course. Some smart people lack social status, and some people with high social status don't know very much. But, on the whole, there is a connection between the two that has

existed over the long course of human history. Having higher social status means having more power. And so it follows that knowledge is power since having it gives someone greater social authority and privileges.

These three common benefits of knowledge—success, virtue, and social authority—are important to keep in mind as we explore knowledge philosophically. Some philosophers have treated knowledge as if it were in a vacuum, as if the “power dimension” of knowledge could be safely ignored. But such an approach leads to an impoverished understanding of the character of knowledge and its importance. We need to consider also *what is getting done* when people are making claims to knowledge or accusing others of not having knowledge.

We might consider an analogy. Imagine a society in which people generally made judgments about whether certain action were “pure.” The people in this society who perform *pure* actions are highly respected, receive medals, lead parades, and so on. Those who perform *impure* actions are shunned and sometimes imprisoned. We would want to know a lot more about what “pure” means and whether there was anything to it—meaning whether there really was such a

quality as “pure” or whether it was simply imagined by people and perhaps motivated by a desire to raise some people up and push others down. This is a good question to have in mind as we explore knowledge. Is there really anything to it? Or is it a social device for putting people into favored or disfavored groups? Of course we think that knowledge is necessarily connected to *truth*. But at the same time, we know how hard it is to be sure a claim, especially an important claim, is true.

Given the power dimension of knowledge, it would be marvelous if we had some clear, impartial method for determining what is true and what is known.

4.

QUESTIONS TO CONSIDER

1. Suppose Marcelo says that he knows the moon is made of green cheese. You ask for his justification, and he says that he works in a cheese shop and so he knows cheese when he sees it. Thus, he knows that the moon is made of green cheese. You want to show Marcelo that he's wrong, but in looking things up, you find a credible article in *Nature* claiming that, in fact, there is evidence of trace elements of green cheese in lunar soil. Now what do you say to Marcelo?
2. Petra adds 59,086 and 63,212 and gets 122,398. She is obsessive about such things, so she checks six times and gets the same answer each time. Why is she wrong to say she *knows* that 59,086 plus 63,212 equals 122,398? Once you have answered this question, consider this one: how many times did you check the calculation? How do you *know* that Petra is wrong?
3. Imagine encountering a society that has built rocket

ships, atomic clocks, and X-ray machines. But when you talk with their scientists, they believe they are harnessing the spirits of demons locked inside in the earth, and their science textbooks look completely different from our science textbooks. Put the following responses in order from “least plausible” to “most plausible”:

1. These people believe false things, and they have unknowingly built some very impressive devices.
2. These people believe things which seem false, but if we studied their beliefs more carefully, we would find that they have the same science we have, but expressed in wildly different terms.
3. What these people believe is just as true as what we believe, and their science is radically different from ours; the two sciences are not merely the same knowledge expressed in different terms (so #2 is false).

Explain your ranking.

5.

FURTHER READING

Online encyclopedias such as the *Internet Encyclopedia of Philosophy*, the *Stanford Encyclopedia of Philosophy*, and *Wikipedia* offer overviews of philosophical problems and theories regarding knowledge. There are also many introductory texts, including:

Charles Landesman, *An Introduction to Epistemology*
(Wiley-Blackwell, 1996)

Robert Martin, *Epistemology: A Beginner's Guide*
(Oneworld, 2010)

Jennifer Nagel, *Knowledge: A Very Short Introduction*
(Oxford, 2014)

The story of Molly and Doug and the burned-out light bulb is known as a “Gettier Problem,” named for Edmund Gettier and an article he published in 1963. The Gettier Problem led to a great number of papers providing new JTB+ accounts

or raising new problems to the new accounts. To get a sense for this extensive literature, you can simply search for “Gettier Problem” on the internet.

A very good general account of Gettier problems, and how they just won’t go away, is Linda Zagzebski, “The Inescapability of Gettier Problems,” *The Philosophical Quarterly*, vol. 44, no. 174 (Jan., 1994), pp. 65-73.

PART II

2. SKEPTICISM: A DIALOGUE

“Skepticism relieved two terrible diseases that afflicted mankind: anxiety and dogmatism.”

— Sextus Empiricus (c. 200 CE)

Patient: Help me, doctor! I know too much!

Doctor: Goodness me, what a curious complaint! But don't worry. We'll get it sorted out. Here, have a half-full glass of water, or what at least has all the properties of normal water. I have put a straw in it. Isn't it curious how it appears to be broken at the water line?

Patient: (*slurps*) Ah, thank you.

Doctor: Now why is it you think you know too much?

Patient: I'm afraid I shall have to be immodest, if you are to understand my problem. I am insatiably curious and am always reading about science, history, politics, economics, literature, philosophy, and, well, everything. I have a fantastic mind: I understand things very quickly and make astute connections among the things I learn. And I rarely forget anything I have read.

Doctor: Wow!

Patient: And if I am honest, I must say I have the most logical mind I've ever encountered. I draw inferences like nobody's business! I have been busy over the last few years writing a long book which I am titling *The Encyclopedia of All Human Knowledge*. And, well, not to brag, but just to tell you how things are, I am certainly the person to write it, for there are very few things that belong in an encyclopedia that I don't already know.

Doctor: I must say, you certainly are a *rara avis*, which means—

Patient: “Rare bird,” coming from Juvenal’s sixth satire in which a good wife is said to be as rare as a black swan.

Doctor: I didn’t know that! You do indeed know a lot. But such broad knowledge would be a wonderful thing to possess. Anything you can think to ask yourself, you can answer. Anytime you feel a bit of wonder, you can offer an explanation. With such knowledge you can answer anything! Why on earth do you regard it as a problem?

Patient: I am bored out of my head! I have written half of the encyclopedia, but I can’t possibly go on, since I already know everything I’m going to say. What’s the point? You must understand, doctor, that learning new things is my greatest delight. But now I know too much, and there’s nothing for me to learn! Oh, how I wish there were something for me to learn!
(*breaks down in tears*)

Doctor: There, there! Here, dry your tears on this handkerchief, which is either white or blue, depending on how you look at it. I see your problem, and though I have not run

into it before, I think I know how to treat it. What you need, my friend, is a healthy dose of *skepticism*.

Patient: Skepticism? But you should be warned, doctor, that I not only know everything, but I also know the explanations for everything and all of the justifications. I don't think skepticism can possibly help!

Doctor: Well, you may be surprised. In my experience there is always some room for doubt. What's more, I think you will find that if you apply skeptical doubts of the right kind, in the appropriate measure, and at the right time, you will be able to transcend your great knowledge and reach a very fine form of happiness.

Patient: I guess it's worth a shot. I'm miserable now as things are.

Doctor: Good! Now since you are such an intelligent person, I will teach you seven skeptical doubts which I am sure you will be able to apply as needed. I call them the "seven modes."

Patient: Sounds good. But I doubt that this will help!

Doctor: See, you are making progress already! The **first** mode has to do with *history*. You say you are writing an encyclopedia?

Patient: Yes, the most comprehensive and accurate encyclopedia ever written, if I may say so myself!

Doctor: Will your encyclopedia include a history of encyclopedias?

Patient: Well, yes, I suppose it must.

Doctor: So then you will include the great, long tradition of encyclopedia writing. You will include Gregor Reisch's *Margarita Philosophica* of 1496, and Paul Scalich's 1559 encyclopedia, and of course the wonderful *Cyclopaedia*

Ephraim Chambers published in 1728 which led to the French *Encyclopédie* which led to the *Encyclopedia Britannica*. But your encyclopedia will improve upon all of these, will it not?

Patient: Immeasurably!

Doctor: But the funny thing is that each of the authors of these older works also thought they were having the final say and that they were improving immeasurably upon the encyclopedias of the past. Why do you think you are better off in this regard?

Patient: Well, they were all smart people for their time, but I know more than they did then.

Doctor: But would they not have said the very same thing?

Patient: I guess so.

Doctor: And they were wrong by our lights. But what should we conclude from this evidence? Doesn't history suggest that writers of encyclopedia always take themselves to have knowledge and (at least so far) have always been wrong about that?

Patient: That does seem to be the case.

Doctor: Good! So that is the first mode. The **second** mode has to do with *culture*. All of those encyclopedias I mentioned, the ones by Reisch and Scalich and Chambers, were written by western Europeans, were they not?

Patient: Yes. Reisch was born in Württemberg in 1467 and went to—

Doctor: Hold on now, the point's been made! They were Europeans. But it would be very surprising, would it not, if there were not also encyclopedias written in many other lands such as China and India and Egypt.

Patient: Yes! Why, in China, the 1726 *Chinese Collection of Pictures and Writing*—

Doctor: Yes, yes, exactly! But these other encyclopedias no doubt had very different entries and very different explanations owing to the different cultures and languages and religions and philosophies and systems of science. And yet all of these non-European writers of encyclopedias also took themselves to be giving the final say and improving immeasurably upon every other effort.

Patient: Yes, I suppose so.

Doctor: And, within their own cultural contexts, they had as much reason to take themselves to have superior knowledge as you do now in your culture. Am I wrong?

Patient: No, you are right. They thought they knew everything—though their “everything” was not the same

“everything” that the European writers took themselves to know.

Doctor: So that is something to consider: whether any person bound to a single culture can really be said to know everything. Or even *anything*! For wouldn't anything you claim to know be understood quite differently by someone from a very different culture?

Patient: Yes, I guess that would be true. Every culture has a different way of understanding things, and to the extent they don't square up with one another, it's hard to be confident about what one culture claims to know.

Doctor: Onward, then, to the **third** mode which concerns your own *brain*.

Patient: And a very fine brain it is!

Doctor: One of the very best, I agree! Consider the

unimaginably long history that has brought about your brain. The long process of evolution from bacteria to vertebrates including fishes and reptiles and mammals and, eventually, humans, took millions and millions of years. And it was no smooth process, as you know. There were plenty of failed starts and extinctions along the way, and plenty of accidents, both lucky and unlucky, that led to things being the way they are today.

Patient: Yes, I know this well! I have several insightful entries on evolutionary history.

Doctor: I'm sure you do. And so I am also sure you realize that our brains did not develop for the sole purpose of gaining knowledge. Our cognitive set up had to be *just good enough* to allow our ancestors to reproduce before dying. And that doesn't necessarily require an organ that is excellent at knowing. It only requires an organ that isn't a total disaster, at least in the department of making babies.

Patient: I hadn't heard it put that way, but you're right. Evolution just selects for reproductive fitness, not necessarily for *epistemic* fitness.

Doctor: Exactly. So that means that, for our knowledge, we are relying on an organ that wasn't exactly evolved for being good at knowing. We would be skeptical of the use of any device that wasn't designed for the purpose we're using it for, wouldn't we? So shouldn't we be skeptical of our own brains when it comes to knowing things?

Patient: I guess so ... hey, wait a minute! It's *our brains* that are telling us this! Our brains figured out evolution by natural selection—well, Charles Darwin's brain figured it out, at any rate—and our brains have deduced these skeptical consequences. So don't we have to trust our brains in order to not trust our brains?!

Doctor: Now you're getting that hang of it! It's hard to know what to believe, isn't it?

Patient: I feel dizzy.

Doctor: Here, have another sip of this stuff that certainly seems to be water. Now let's turn to the **fourth** mode which is related to what we were just talking about. I call it the *animals and aliens* mode. Human brains evolved in ways that have ended up helping human beings lead human lives, right? But the same is true for ape brains, dolphin brains, elephant brains, even frog brains—each species evolved brains that suit the lives those species live.

Patient: Yes, it's the same story across the board.

Doctor: And there is so much interesting variation! Think about bats with their echolocation, or how pit vipers can see heat. And did you know that, while we have only three sorts of photoreceptors in our eyes, mantis shrimp have as many as sixteen different kinds?

Patient: Indeed! They can see deep ultraviolet light, as well as far red light, and polarized light.

Doctor: Just think about the sorts of encyclopedias these

animals would write! What they experience is so different from what humans experience, but their experience is just as valid, wouldn't you agree?

Patient: Sure. In many cases, the animals experience much more than we do.

Doctor: And I presume you are familiar with all of the arguments suggesting that there almost certainly is life elsewhere in the universe.

Patient: Of course! There must be, given how enormous the universe is. It is statistically quite certain that life has evolved on countless other planets.

Doctor: Some of those aliens have probably written encyclopedias as well! And in all likelihood, those aliens would have evolved in very different ways from any life on earth. So what they would claim to know would be even more different than what our own bats and vipers and mantis shrimp would claim to know.

Patient: Yes, it would be impossible for us to conceive what they would think they knew.

Doctor: But there is no reason to think any less of them for that, is there?

Patient: Of course not! Okay, okay, I see the point: animals and aliens would have knowledge that is radically different from human knowledge, and every bit as valid, if not more so, which means human knowledge cannot be the be-all and end-all.

Doctor: I could not have said it better myself! Are you ready for the **fifth** mode?

Patient: It is hard to believe there's even more doubts to consider!

Doctor: You can be sure of it! Set aside the last two doubts about our brains and the brains of animals and aliens. Suppose we could assure ourselves that our brains are in fact excellent trackers of truth. Still, a person as educated as you must surely realize that there is hardly anything to put into an encyclopedia that isn't contradicted by some other expert somewhere?

Patient: Don't I know it! You know, as a hobby I routinely correct entries on *Wikipedia*, and boy have I gotten into fights—sometimes even over the most trivial things!

Doctor: It seems like there is no bit of knowledge that isn't contested by someone. But often, that someone has some reasons for raising their objections, especially if the matter has big consequences. No one cares much about the exact temperature outside the Poughkeepsie courthouse at noon, but start talking about raising the minimum wage, and *boo boy!*

Patient: Yes sir!

Doctor: So, it would seem, on a great many matters, *experts*

disagree. And we know from the first mode, the one based on history, that experts can be very wrong. So if experts, who dedicate much of their careers to understanding important things, can disagree with one another, and may well be shown to be wrong later on, then what hope for getting things right do the rest of us have?

Patient: Not much. I mean, I know how hard it is to be an expert since I am one on nearly everything. But there are always people disagreeing with me, and sometimes, if I'm honest, I'm not sure they're wrong!

Doctor: Interesting, isn't it, that experts, who represent the highest knowledge a human can have on a topic, end up serving as a reason for being skeptical, even about that very topic!

Patient: Well, only because experts so often disagree! It makes it hard for a nonexpert to know what to believe or whom to trust.

Doctor: Exactly. So let us consider next the **sixth** mode which is based on *logic*.

Patient: Aha! I am an expert logician!

Doctor: So I gather! And where did you learn logic?

Patient: I attended a very fine school in Des Moines.

Doctor: Des Moines? And where is that?

Patient: In Iowa!

Doctor: And where is that?

Patient: In the United States, of course.

Doctor: And where is that?

Patient: Is this a trick? It's in North America. And before you ask, that's on planet Earth, in the solar system, in the Milky Way galaxy!

Doctor: And where—

Patient: In the Virgo Cluster, which is in the Laniakea Supercluster, which is in the universe!

Doctor: Good! And where is the universe?

Patient: What sort of silly question is that?! No one can answer that!

Doctor: Ah, so you don't know where you learned logic.

Patient: Uh ... *what???*

Doctor: I'm just having some fun with you, of course, calling to your attention the idea of a *regress*, which is a long chain of questions either without an end or with a big question mark at the end. A regress suggests that if you can't answer the long chain of questions, you don't really have knowledge of what's at the tail end of it. Here's how it's relevant to your case. Didn't you say earlier that you not only knew everything, but also all of the explanations and justifications for everything?

Patient: Yes, I did say that! Though now I'm becoming less sure ...

Doctor: So, if I asked you about something you know—let's say, about some event in the War of 1812—you could tell me about the event, and you would be able to give me evidence for believing what you do about that event?

Patient: Yes, I would! I am very scrupulous.

Doctor: And so then, you could give me reasons to believe that evidence? Evidence for the evidence, so to speak?

Patient: I think so ... well, to some extent...

Doctor: And evidence for the evidence for the evidence?

Patient: Hold on! Explanations have to stop somewhere!

Doctor: But where? With things for which you have no evidence?

Patient: Well, yes, I mean, I *guess* so...

Doctor: But this means everything you believe is ultimately based on things for which you have no evidence!

Patient: Just because there's no evidence for something doesn't mean I shouldn't believe it!

Doctor: Come again?

Patient: Well, just because I have no reason to believe them doesn't mean I have no reason to believe ... uh, wait a minute. I'm not sure what I'm saying

Doctor: Excellent! See? Your swelling of knowledge is already going down!

Patient: I feel everything slipping away ...

Doctor: Isn't it a nice feeling? All those things you thought

you knew and had to keep track of and worry about are all slipping away. No knowledge, no worries! This is the perfect point at which to tell you about the **seventh** mode, my favorite, which I call the *forever open alternative*.

Patient: What's that?

Doctor: All the doubts we have been through have had something in common: they have all indicated that there could well be something we hadn't thought of. Maybe it's a future discovery, or the discovery by another culture, or by animals or aliens, or something our brain isn't picking up, or some further question we hadn't thought to ask. The forever open alternative is the general possibility that there is something important we haven't considered yet. Who knows why! The world is such a big, complicated place, and we are so tiny and live such short lives that it would be incredible if there weren't extra complications we have not considered!

Patient: That seems unavoidable. It's always possible that things are more complicated in ways we have not even imagined!

Doctor: Yes, and maybe we can't even see the ways in which they are more complicated! The forever open alternative is simply the recognition that we may not know as much as we think we do. It's the essence of skepticism, really. There's always an open possibility that we just don't really know what we think we know.

Patient: Doctor, I think you have cured me! I feel like I can diminish my confidence in everything I used to believe!

Doctor: Glad to hear it! You will learn through experience which mode to use in each case, since some will be more effective than others, depending on what you are doubting. But if you apply yourself, you find you are able to place yourself in a blissful state of *not knowing*. You will balance yourself between each claim and its denial, neither affirming nor rejecting either one, but floating above them all. And that, my friend, is a happy state of mind!

6.

QUESTIONS TO CONSIDER

1. Consider the following five alleged facts. How would you go about doubting them? Feel free to make use of the doctor's seven modes if they are useful.
 1. The aroma of freshly brewed coffee is much nicer than the aroma of an angry skunk.
 2. It is impossible for a human being to levitate using only their own powers.
 3. Two plus seven equals nine.
 4. Canada is larger than India.
 5. The total entropy of an isolated system can never decrease over time and is constant if and only if all processes are reversible. (This is called "the second law of thermodynamics.")
2. It is sometimes said that the Academic Skeptics believed that nothing can be known, but the Pyrrhonian Skeptics doubted even this. Which group do you think is more skeptical? Justify your answer.

3. It seems very odd to say, “I know X, but I’m not sure of it.” (Try out some examples.) But did anything in the JTB+ account of knowledge suggest that knowledge requires *being sure* of something? Why does it seem odd to us to say we are not sure of the things we think we know? Or does it?

7.

FURTHER READING

Sextus Empiricus, *Outlines of Pyrrhonism*. Sextus was a physician in ancient times. He recommended skepticism as a way of attaining peace of mind (*ataraxia*), and his book was meant as a guide. A useful summary of his views can be found also on *Wikipedia*.

Peter Unger, *Ignorance: A Case for Skepticism* (Oxford, 1979) provides more contemporary arguments in favor of skepticism.

PART III

3. THE GRAND DECEPTION DOUBT (OR, THE GDD)

Consider this possibility from the philosopher René Descartes (1596-1650):

“Nevertheless, the belief that there is a God who is all powerful, and who created me, such as I am, has, for a long time, obtained steady possession of my mind. How, then, do I know that he has not arranged that there should be neither earth, nor sky, nor any extended thing, nor figure, nor magnitude, nor place, providing at the same time, however, for the rise in me of the perceptions of all these objects, and the persuasion that these do not exist otherwise than as I perceive them? And further, as I sometimes think that others are in error respecting matters of which they believe themselves to possess a perfect knowledge, how do I know that I am not also deceived each time I add together two and three, or number the sides of

a square, or form some judgment still more simple, if more simple indeed can be imagined?”¹

We might reformulate Descartes’s doubt into a set of steps as follows:

1. Everything I believe about the world is based on what has come to me either through my senses (what I have seen, read, heard, or experienced), or through my own power of thought.
2. It seems to me there could be a being, like God, who has power over what I sense and even has power over my own power of thought.
3. If God wanted to, he could make it *seem* like I am perceiving something when *really* I’m not, and God could make me think something *must* be true when *really* it isn’t.
4. I cannot know for certain that God isn’t deceiving me in these ways.

1. Rene Descartes, “First Meditation,” in *Meditations on First Philosophy* (translated by John Veitch, 1902).

5. So I really cannot be certain about anything I believe about the world.

Call this “the Grand Deception Doubt,” or the **GDD**, for short. If you prefer not to bring God into arguments like this, consider replacing God with a very powerful demon or some mad scientist who has captured you, removed your brain, and put it into a jar with a bunch of wires connecting it to a supercomputer which gives the scientist complete control over what you think you perceive and what you think you think.

Is there any way to refute the **GDD**? No, not really. But let’s see why. To refute it, we would need some way to “defuse” it and show that, in fact, we could not be experiencing such a massive and thorough deception. Our first move might be to point out that it would just be too impractical for some god, or some demon, or some mad scientist with a lot of advanced technology, to accomplish such a deception. Who could have the power to do such a thing? Why would they want to do it? What advantage would they have in deceiving us so thoroughly? But the answer to all of these questions is the same: “Who knows?” That is to say, we do not know what motivates gods and demons and mad scientists, and for all we do know, they might have very good reasons for deceiving us. Maybe they are competing in a cosmic contest. Maybe the

GDD is cheaper than creating an actual universe. Maybe they are teaching us a lesson. Maybe they are malicious. Maybe, maybe, maybe. *For all we know*, the skeptic will say, there might be a powerful demon with both means, motivation, and opportunity to deceive us. *So long as we cannot rule that out*, we cannot be certain of the truth of what we experience.

Next effort. We might point out that this is crazy talk, and normal people don't go around worrying about the possibility of a **GDD**. But the skeptic will respond that the fact that "normal people" don't worry about something does not in any way show that the thing isn't true. And—by the way—the skeptic will point out that the very existence of these alleged "normal" people is also cast into doubt by the **GDD**. They could be part of the grand illusion, like non-playable characters in a video game.

OK, let us try a more philosophical approach. We have experiences. The experiences must come from somewhere. We could believe that they come from what they seem to come from—namely, a real world with other people and animals and trees and buildings and so on. Or we could believe they come from some radically different source—a god or demon or mad scientist. It is more rational to adopt the simpler and

most straightforward explanation which is that experiences are coming from what they seem to be coming from. Therefore, it is rational to reject the **GDD**.

This is a much better sort of reply. For one thing, it uses words like “rational” and “therefore,” so it seems very philosophical! But is it enough of a reply? No, it isn’t. It makes two highly questionable assumptions. The first is that our ordinary view, that our experiences come from a real world, is simpler than the view that our experience comes from some other source. But why should the ordinary view count as “simpler”? It’s more *common* (or seems to be), but “common” does not mean the same as “simpler.” The fact that the view is more common could be entirely accidental. If we were all taught that experience comes directly from an eggplant in the heavens, then that belief would be more common, but we should not on that account take it to be more likely *true*.

The second questionable assumption is that it is more rational to believe the simpler and most straightforward explanation. But why should we believe this principle is true? Is it because we know the world to follow rational rules and principles? Why must this be true? I might believe that the world would be a more rational place if people shared their surplus wealth

with others who needed it, but it does not follow that the world is as I believe it rationally should be. Why should this case be any different?

At this point someone might throw up their arms and insist, “It doesn’t matter! You can raise all the skeptical scenarios you like, but in the end we are all going to keep believing the ordinary view of things!” Most skeptics would agree. People will continue to believe what people will continue to believe. The skeptic’s point is merely that people do not have *good reason* for believing what they believe, and if people claim they do, they are mistaken.

This puts us in a somewhat awkward position. Is it true that we don’t have good reason for believing what we believe? Is it no more rational to believe the claims of scientists and historians than to believe that all our experience comes from an eggplant in the heavens?

8.

QUESTIONS TO CONSIDER

1. Suppose some event indicated that you really were living in a GDD situation. Suppose, for example, a pop-up screen entered your visual field in real life and notified you that you are living in a simulation. How would your life change—if at all? Explain.
2. Our senses give us a misleading picture of the world, and news media oversimplify and exaggerate events. At any moment some high number of our beliefs are false (though we don't know which ones). But these facts do not ordinarily trouble us. Does the **GDD** add anything new that is especially troubling? Or should we be more troubled than we typically are about ordinary life?
3. “The **GDD** refutes itself. For if we were in a **GDD** situation, then our ability to conceive of demons, mad scientists, or even God is unconnected with whatever reality there is. So we cannot trust our ideas about such things, and thus, we cannot trust our idea that the **GDD**

is even possible.” Discuss.

9.

FURTHER READING

Descartes's *Meditations on First Philosophy* is a classic expression of the **GDD**, and Descartes's systematic response to it is ingenious if not entirely compelling. Many versions of the work can be found on the internet.

It is also interesting to work through a modern expression of the **GDD** in the form of the "Simulation Argument" put forward by Nick Bostrom (available online at <https://www.simulation-argument.com/simulation.html>).

Bostrom is not urging skepticism so much as trying to establish that probably we *really are* living in a simulation. The basic idea is that if it is possible for simulated beings to be conscious, and if advanced civilizations are likely to run simulations of the universe, then (given the probable populations of such simulated beings) it is more likely than not we are living in a simulation.

PART IV

4. FIRST ANSWER TO THE SKEPTIC

There are so many good reasons to be an extreme skeptic. There are the seven modes offered by the skeptical doctor in the earlier dialogue, and the Grand Deception Doubt (**GDD**), and just the routine experience of being wrong, which happens to everyone, and which could happen to anyone with regard to everything, it seems. And yet hardly anyone is really this skeptical. Why is this? Do we have good reasons for not being skeptics? Or are we just proud and stubborn and unwilling to think we might be wrong about everything?

After expressing the **GDD**, René Descartes went on to try to refute it. He wanted to show how to organize our knowledge so that everything is based upon secure principles and we would be free from skeptical doubts. He did this by establishing one thing he could never be wrong about: namely, that *he existed* as a thinking thing. “I think, therefore I am,” he said. He was not sure whether he had a body, or whether there were other people in the world, but he knew that he could not be wrong about his own existence. This insight of Descartes’s

is called “the cogito” since in Latin he would have said “*Cogito, ergo sum.*”

What made him so sure? Ultimately, he found he could not make sense of any way of being wrong about his own existence. You may try to prove this to yourself. Try to imagine some situation in which you do not exist but are somehow deceived into believing that you do exist. It is impossible, for you would have to exist in some fashion in order to be deceived about anything! If you are not convinced, then try this experiment: grab something that doesn’t exist, and then try to trick it into believing that it does exist. Any success?

But if this first insight is secure, it is hard to know what to do next. Descartes goes on to argue for God’s existence and that God is not malevolent and that God would not deceive us about things that seem to us to be clearly and distinctly true, and before you know it, he has built an entire system of knowledge upon his cogito foundation. Readers generally do not find his arguments compelling. But they like the cogito. It seems like the one definitive claim philosophers have been able to prove!

In this chapter, we will explore another way someone might try to answer the skeptic, and in the next chapter we will explore a second way. The first way requires us to be content with a system of knowledge that is based upon what *appears* to be true.

10.

KNOWLEDGE OF APPEARANCES, OR PHENOMENALISM

The skeptics who offer the **GDD** claim that, for all we know, everything we ever sense could be an illusion. One way to reply to these skeptics is to call their bluff. *So what* if everything is an illusion? We still have to deal with appearances. We still have to get up in the morning, do our exercises, go to work, deal with people, solve problems, pay taxes, and die. Does it really matter whether what we are experiencing is only a very thorough illusion? What practical difference would this make in our lives?

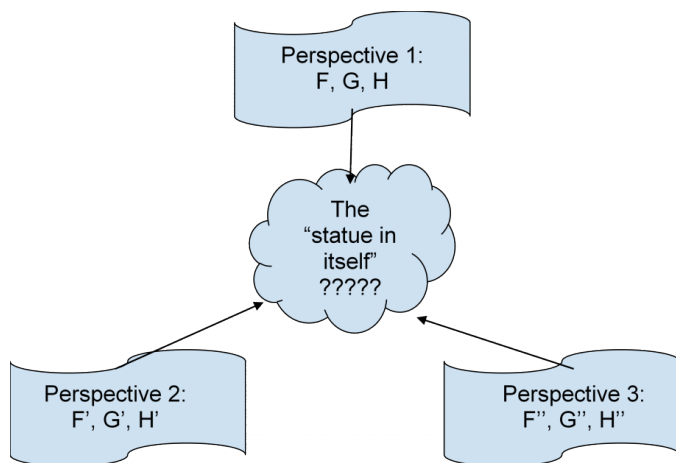
This response to the skeptic maintains that we can get all the knowledge we shall ever need through a careful study of appearances, and we should never have to insist that these appearances come from some “external” world (meaning, a world external to the appearances). The task, then, is to show that we really can get everything we need from mere

appearances. The efforts to show this are sometimes called “empirical idealism” or “phenomenalism,” and such efforts have been made by George Berkeley (1685-1753), Ernst Mach (1838-1916), Bertrand Russell (1872-1970), and Rudolf Carnap (1891-1970), among others. We will sketch a similar effort in this section.

We can begin this project by thinking through what we would have to do if we wanted to create a very compelling virtual reality machine. With such a machine, you can strap on goggles or a helmet and be treated with the sights and sounds of another world—a Martian landscape, Middle Earth, ancient Rome, or whatever you like. But we want our VR machine to be much more thorough. We want not only sights and sounds, but smells and textures and tastes. We want to feel the ground beneath our feet, and we want to experience the difficulty of walking up steep hills. We want the world to be wide open so that no matter where we decide to turn or run or jump or crouch, the machine will respond with the appropriate set of images and features for us to experience. *We want the real in our virtual to seem as real as our real in the real.* (By the way, isn’t that an interesting sentence?)

Let’s think through what this requires by taking up a specific

example. Let's say I am in the VR of ancient Rome, and I am going to walk up to a statue and examine it from all angles. Obviously, the statue that I am looking at does not exist in the real world. In the VR, I am presented with images that show how a real statue would look from all possible angles. I walk around the sculpture, and I am presented with a series of images of the sculpture, smoothly transitioning from one to the next. I lay down on the ground, and I see an image of the statue from below. But, again, there is no statue—not really. All that exists is a very extensive set of images produced by the VR program. I experience three sets of images similar to one another but also different which together “suggest” that there is an object “out there” that corresponds to the sets of images I experience.



Of course, in a VR situation, there is no “statue in itself” except as a collection of images included in a number of “perspectives.” We pretend that there is a “statue in itself,” though we only see the images from one perspective or another. Image F is very similar to image F’, which is very similar to F”, but they are all subtly different from one another in such a way to suggest the illusion that we are moving around a stable, fixed object (“the statue”) which does not really exist. Setting this up requires a lot of thought and a lot of coordination, though there is no reason to think it cannot be done. It just requires a lot of planning in how the images are constructed and the order in which they are viewed.

To make the sort of VR we are imagining, this sort of construction would have to take place for each and every virtual object and for all of the angles from which the virtual “object” could be seen. Similar sorts of programming would have to be done for every sound and smell and taste. Imagine having to program just how a virtual orange would taste after biting a virtual lemon as opposed to how it would taste just after brushing your virtual teeth! Imagine having to program just how the smell of some virtual old socks would change as you became used to the smell.

“Hold on,” I imagine someone thinking. “What do you mean by my ‘virtual’ teeth? Can’t I use my real teeth?” No, you can’t. For in the VR we are imagining, we will have to present you with images of your (virtual) feet if you look down at them, images of your (virtual) eyes if you look into a (virtual) mirror. We shall even have to present you with images of your (virtual) brain if you decide to do brain surgery on yourself. Basically, all of the sensory images you experience—indeed, *all sensations whatsoever*, whether of yourself or other things—will have to be generated from the VR programming. Plugging into this VR system means leaving the so-called real world behind and experiencing only the sights, sounds, tastes, textures, smells, and other features of the virtual world. And that will include the sensations of your own body.

I keep pointing out that the ‘objects’ in the virtual world don’t really exist. That is to say, they do not exist in the *normal* way we take ordinary objects in our ordinary reality to exist. But there may be another sense in which they do exist: they exist as complicated structures within the VR program. For example, I might say that the statue’s left elbow *just is* the set of F, F’, and F” that relate to one another in the way described above. The same, of course, for the head and for the left knee. What groups them together as images of the left elbow, or the left knee, or the head, is that the images share *a certain structure*, or have great similarities to one another, and the images change

smoothly from one to another according to general rules that hold for any observer who is “walking around the statue” in the same way we are.

Hopefully, we have spent enough time thinking through this example to have a solid sense of what is being proposed under the names “empirical idealism” or “phenomenalism.” These philosophical views assert that our real world is in fact like the virtual reality we have been thinking about. There are no material objects existing “out there” for us to experience. There are only the sensations we experience, organized in the way a very advanced VR would be programmed so as to generate the appearance of a world of objects. Really, when we refer to the Eiffel Tower or the pyramids of Egypt, we are referring to all of the perceptions one might experience when one undergoes the actions that initiate certain sequences of perceptions we associate with “experiencing the Eiffel Tower” or “experiencing the pyramids.” *Only experiences exist.* And this includes the experiences of our own bodies for the reasons given above. We experience our bodies through our senses, or through images in mirrors or recordings, but these also are just experiences and not different in kind from the statues and towers and pyramids we can experience.

Normally (when we are not doing philosophy), we believe something like this: there exists a world of objects which cause us to have perceptions of those objects. A phenomenalist drops out the world of objects and inserts in its place some other cause, or perhaps just a question mark, and keeps everything else—all our perceptions—exactly as it was before.

Objection: What is the point of working though all of these construction details if, in the end, all our experience remains unchanged?

But remember why we started this section in the first place. We wanted to find some way to answer the skeptic. The skeptic was asking how we could be sure we were in a real world instead of some sort of virtual reality (caused by demons, mad scientists, God, etc.). *We called the skeptic's bluff.* That is, our effort has been to show that our world could very well be virtual in just the way the skeptic seems worried about, and it really would not matter for all of our intents and purposes. The **GDD** need not bother us.

In calling the skeptic's bluff, we gain some advantages. First, all our knowledge is based on things we immediately experience—the sights, sounds, smells, tastes, textures, and perceptions of daily life. We could be wrong about what those

experiences are experiences *of*, but we cannot be wrong about the fact that *we have* those experiences. And if we base the rest of our knowledge just on those experiences, we will have placed all our knowledge on a secure foundation, and we can tell the skeptic to run along and pester somebody else.

There are some interesting further questions we might consider:

1. If only experiences exist, how am I ever wrong about anything?
2. How can I be sure other people (meaning, other conscious minds) exist?
3. What causes my experiences? Who or what organizes my perceptions?

We will briefly consider each of these questions in turn. But first I will offer a hint to help you in thinking through the answers to them for yourself. For the phenomenalist, all knowledge is rooted in experience. So the answer to any question will be a question about *what our experience is* when we *ordinarily* try to answer those questions. If you want to know how a phenomenalist answers a question, ask yourself what experiences you have when you try to answer those questions in a more commonsensical context.

How am I ever wrong about anything? Well, what happens ordinarily when we have the unfortunate experience of being wrong about something? Perhaps I think that $5 \times 7 = 42$; someone takes me slowly through the calculation and shows me otherwise. Perhaps I think India is larger than Canada; someone grabs a globe and asks me to study it more closely. Perhaps I think Woodrow Wilson is president; someone shows me a recent newspaper that suggests otherwise. Generally, when we learn we are wrong, we can describe *the experience* of being shown wrong, and once we have the experience of being shown wrong, we have all the phenomenalist requires to answer the question.

This means that truth means something slightly different for the phenomenalist than it does for non-phenomenalists. Non-phenomenalists think that we are wrong when what we think or say does not match how things really are (where “really are” means outside of one’s experience). But for the phenomenalist, being wrong means saying or thinking something that does not match what the rest of experience shows. When someone is wrong, their belief is *not consistent with* a relevant set of experiences (like doing a calculation carefully, studying a globe, or reading a newspaper). This might be seen as a point in the phenomenalist’s favor, actually, since that *does* seem to be what we do when we try to determine whether we are right

or wrong: we look to other experiences and see whether our beliefs are consistent with them.

How can I be sure other conscious minds exist? Well, once again, consider what experiences suggest to us that other conscious minds exist. To the extent that I am sure of the existence of other minds, it is because of what other people seem to say or do. On the basis of my interactions with other people, I come to believe they have their own thoughts and feelings, and that they have experiences much like my own. If someone presses the objection, asserting that, for all I know, other people could be characters generated by a computer program, or non-playable characters, I may have to admit that this could be the case. I really cannot be sure at a very fundamental level. But who can? This again seems to be a point in the phenomenalist's favor. We should be skeptical of any answer to a skeptic that promises too much, or promises knowledge that we don't ordinarily take ourselves to have.

What causes my experience? This question requires special treatment if we think that whatever causes experience has to be something that does not show up in our experience. If we think this, then obviously we cannot examine our ordinary experience to show us how we ordinarily go about answering


the question. At first glance, it appears that a phenomenalist really has no compelling answer to offer. Our experience could be coming from a world of spatiotemporal, material objects, or it could be coming from a demon, or a mad scientist, or God, or some psychedelic eggplant with telepathic powers. There is no way of knowing since, for the phenomenalist, all our knowledge is based upon experience, and obviously, we do not have any experience of anything outside of experience.

But on further thought, we may notice that the phenomenalist does not *need* to answer this question. The basis of knowledge, for the phenomenalist, is experience, or *data*, or “what is given.” What we know we know because it is either present in experience or constructed from experience. If someone asks us about things that are totally outside of experience, we may legitimately respond, “You are asking about things of which no knowledge is possible.” And no one should expect us to do the impossible! So the question “What causes our experience?” is something like a meaningless question, like asking whether triangles are married or if the color green is for sale.


So phenomenalism has a lot going for it as an answer to the skeptic. One might even praise it as the most responsible approach to knowledge someone could take, as it forces us to

make sure everything we think we know is connected in clear ways to what we experience and cautions us against believing in anything beyond what experience shows. Yet in my experience, very few people wish to be phenomenologists. In a class of thirty students, I may have one or two who are interested in adopting the view. Strangely, it is because people really like the idea of there being a material world even if it lies outside of any possible experience. For most people, it is not enough to believe in the experience of touching a statue and feeling its hardness; *there must also really be a statue* in addition to whatever I can possibly experience of the statue. I call this “strange” because it is strange that people should feel so confident about an object that they themselves will admit—at least, after a bit of philosophical discussion—is an object that they can never possibly experience (“the statue in itself”).

Let us explore the strangeness of this confidence a bit further. Ordinarily, before we begin to philosophize, we think that there are objects in the world, and these objects somehow cause our experiences of those objects. If we are asked why we believe the objects cause our experience, our general feeling is something along the lines of, “What, are you crazy? These experiences must come from somewhere! My experiences are experiences of objects.” This attitude is what we may call “common sense realism.”



But now suppose I extend common sense realism a bit further. I say, “Well, surely you don’t think that is the end of the matter, do you? For these objects must also come from somewhere just as the experiences must come from somewhere. And just as experiences come from objects, I tell you now that objects must come from *schmobjects*.”



According to this new view, which we may call “uncommon sense realism,” it is not enough to merely claim that objects cause our perceptions. We must explain what causes those objects. The cause of those objects is what we should call *schmobjects*, I say. *Schmobjects* cause objects in exactly the way that objects cause our perceptions of objects.

If you ask me why I believe in *schmobjects* in addition to objects, I shall reply, “What, are you crazy? These objects must come from somewhere! Objects are the effects of *schmobjects*.” You can hardly accuse me of being less rational than the person

who made a similar claim about experiences coming from objects.

But wait, there's more. *Schmobjects* can't just exist by themselves, can they? They must be caused by *flobjects*. And *flobjects* are caused by *globjects*. And *globjects* are ...

And off we go in an endless regress to crazyland. We should cut off this endless regress somewhere, don't you think? And wouldn't it be reasonable to cut it off at the point where our experience ends and not go on to assert the existence of things of which we have no experience? That would leave us with phenomenalism, wouldn't it?

But strangely, most people seem to like to assert the existence of objects and don't feel arbitrary at all in denying the existence of *schmobjects*. What can I say? People are weird.

Media Attributions

- Figure 4.1 © Charlie Huenemann is licensed under a [CC BY-SA \(Attribution ShareAlike\)](#) license
- Figure 4.2 © Charlie Huenemann is licensed under a [CC](#)

[BY-SA \(Attribution ShareAlike\)](#) license

- Figure 4.3 © Charlie Huenemann is licensed under a [CC BY-SA \(Attribution ShareAlike\)](#) license

11.

PROBLEMS WITH PHENOMENALISM

Many people reject phenomenalism for some strange and unconvincing reason. But there are also good lesser-known reasons for rejecting phenomenalism. I will sketch two of these reasons: the problem of necessary truths, and the problem of sense data.

The problem of necessary truths is the problem of how phenomenalism can account for our knowledge of truths that strike us as necessary truths. This is a bit tricky to explain. Some truths seem like just accidental truths about the world. Giraffes have long necks. The moa is, sadly, extinct. And here is some good news: there are currently more than 5,000 black rhinos, which is more than twice as many as there were in 1995. All of these truths, it seems, could have been otherwise. Giraffes could have evolved to have shorter necks, the moa could have survived, and black rhinos could be even more (or less) populous than they currently are. Call these truths that

could have been otherwise *contingent truths* (“contingent” means that they depend on other facts for their truth).

Other truths in our experience seem *more necessary* and less contingent. Consider the claim that the earth is gravitationally attracted to the sun or that for every action there is an equal and opposite reaction. Are these claims true in exactly the same way that it happens to be true that giraffes have long necks and the moa is extinct and there are more black rhinos than before? Normally we think of *laws of nature* as more necessary and less changeable than other claims about the world. They describe *what must happen*, and not just *what happens to have resulted*. But what in our experience could make these truths *more necessary*?

The mere fact that we do not come across exceptions to the laws of nature does not show they are more necessary. We also never come across short-necked giraffes or existent moas or herds of a million black rhinos, but that does not show that there *could* not be such things. Indeed, it is hard to imagine the necessity of a truth being somehow *apparent* in our experience. Would it glow somehow or have a warning sticker on it telling us that this portion of our experience is non-negotiable?

A phenomenalist might try to answer this question by simply insisting that there are laws that govern our experiences, and these laws have some sort of greater authority or are somehow more fixed and less changeable than the contingent truths we experience. Perhaps the laws of nature are due to the program in the simulation or due to rules set down by God or a mad scientist. But if a phenomenalist insists on this, then they must appeal to facts that lie *outside our experience*: facts about the program, or God, or the mad scientist. There is nothing *internal* to our experience to show that the laws of nature are more necessary than any other accidental generalization we come across (like “all giraffes have long necks”). This suggests that one cannot be a “pure” phenomenalist and make sense of laws of nature. They have to *reach outside* their experiences in order to explain something they have found to be true *inside* their experience.

Alternatively, a phenomenalist might bravely deny that any general truths are more necessary than others. They may insist that “all giraffes have long necks” is just as necessary, or just as contingent, as “for every action there is an equal and opposite reaction.” But the cost of making this denial is that it will ruin our efforts to explain natural phenomena. Normally, we try to explain things in nature on the basis of a relatively small set of

fixed laws (like the laws discovered in physics and chemistry). Doing so, we think, gives us a deeper understanding of how nature works since it shows us why particular facts *must* happen in the way they do. If we deny that there are any laws of nature and instead say there are just truths about how things happen to be, it becomes ludicrously easy to explain anything. If someone asks us why apples fall to the ground or why giraffes have long necks or what makes gold more dense than copper, we can simply declare, “That’s how things are!” and call it a day. It would not make sense to dig any deeper to try to find the fixed features of nature that explain what we experience.


So, one problem with phenomenalism is that it is not clear how it will allow us to distinguish contingent truths from necessary truths. A second problem is that phenomenalism requires us to believe in some strange things that exist in a middle place between objects and us, namely *perceptions* or experiences. Other philosophers have called the perceptions “sense data.” They are the things I have put inside clouds in the diagrams. We have been presuming all along—indeed, even the **GDD** itself makes this assumption—that there is a difference between objects and our perceptions of objects, and that those perceptions could, in principle, exist without those objects causing them. But is this a good assumption? How confident should we be that sense data can exist on their own?

One must admit that sense data are weird things. There is no scientific evidence for their existence—they do not show up in brain scans, or under microscopes. But they are supposed to be the most obvious things in the world that even an extreme skeptic cannot doubt. Do they vanish into nothingness when they are not being experienced, and pop back into existence when we have experience? When I see a giraffe, then close my eyes, and then look again, have I experienced two different sense data or have I experienced the same sense datum twice? How can I be sure whether it is one or two sense data—or three or more? If you stand where I stood and look at the same thing, do you now enjoy the same sense datum I enjoyed a minute ago? Or are the sense data two identical copies? There are not good reasons for favoring one answer to these questions than any other. For being the most obvious things in the world, sense data are not very obvious after all!

But is there any way we can understand our experience without using sense data? Yes, there is. Suppose we do away with sense data as intermediary objects existing between objects and us. We say instead that we are experiencing the objects *directly*. But of course your experience may not be exactly the same as my experience even though we are experiencing the very same object. But this difference, we shall

say, is not due to you having one sense datum and my having another; it is instead due to you perceiving the object *in one way* and my perceiving the object *in a different way*. “Different ways of perceiving” are not sense data; that is to say, there are not different *things* standing between us and the objects we perceive. Rather, you and I are seeing the same thing, but *in different ways*. It is a difference between adverbs rather than a difference between nouns.

This might sound like we are merely playing with words, but in fact, this simple change in how we talk about our experiences makes phenomenalism impossible. The whole idea of phenomenalism is that we can separate our perceptions from the objects alleged to cause those perceptions and make do with just the perceptions themselves. But if there is no way to separate perceptions from those objects—if perceptions *just are* those objects, perceived in a certain way or from a particular point of view—then the perceptions cannot be separated from the objects, and phenomenalism does not even get off the ground. There is no gap between the observer and what is observed.



This way of understanding our perceptions avoids postulating sense data which (as we have seen) are weird things to postulate. It does postulate *different ways* of experiencing things, but that is well understood. You see an object from some position, under certain lighting conditions, while wearing sunglasses, and so on. This does not mean that the object you are experiencing is different from the object anyone else is experiencing; it means you are experiencing the same object as other people are experiencing, but under different conditions. This means the objects really do exist after all.

Media Attributions

- Figure 4.4 © Charlie Huenemann is licensed under a [CC BY-SA \(Attribution ShareAlike\)](#) license

CONCLUSION

In this chapter we have seen one way to respond to an extreme skeptic. That way is to deny that we need to have knowledge of objects outside our own experience and to insist that we can make do with just the experiences themselves. The effort to base all our knowledge on actual experiences is known as *phenomenalism*, or empirical idealism. Many people do not accept phenomenalism, though their reasons for rejecting it are sometimes strange. But we have also seen that a phenomenalist does face two problems. One is the problem of distinguishing necessary truths from contingent truths. It is not evident how experiences alone can provide such a distinction. The other is the problem of making sense of experiences, or sense data, as objects. It is not obvious when two experiences should count as the same experience, or whether the same experience can pop up at different times, and these are basic questions that should be answerable about any object said to be real.

13.

QUESTIONS TO CONSIDER

1. Imagine two people arguing over whether Plato was taller than Aristotle. They explore old texts, compare them, and go round and round arguing about it. Now imagine two people arguing over whether Don Quixote was taller than Hamlet. They explore old texts, compare them, and go round and round arguing about it. Someone might say that these two arguments are different because *there is a fact* about whether Plato was taller than Aristotle, but there is no such fact about Hamlet and Don Quixote. What would a phenomenalist say?
2. Can two people experience the same hallucination? Is there a difference between saying they experience *the same hallucination* and saying they experience *identical hallucinations*? What about two people having the same idea or the same concept?
3. Suppose some characters in a video game became

conscious and you can talk to them. Suppose further that you explain to them the **GDD**, and they become skeptical of their knowledge of their world. You read this chapter to them—or send them a copy of it—and as a result some of them become phenomenologists. Does this make any sense? Can they become phenomenologists if the objects they thought they were experiencing were only virtual to begin with?

14.

FURTHER READING

Bertrand Russell wrote an essay explaining exactly how phenomenalism works from a logical point of view in lecture three of his book *Our Knowledge of the External World* (1915), which is available in several places on the internet including Wikisource. If that is not detailed enough for you, you might turn to Ernst Mach's book, *An Analysis of Sensations* (1897), or Rudolf Carnap's *The Logical Construction of the World* (1928) or David Chalmers' *Constructing the World* (2012).

But the most famous and readable phenomenalist of them all was Bishop George Berkeley whose *Treatise Concerning the Principles of Human Knowledge* (1710) lays out a very clear account of phenomenalism with powerful arguments for its appeal.

Ludwig Wittgenstein's *Philosophical Investigations* is a classic

text ranging over many topics, but sections 398-411 throw doubt upon the existence of sense data. It is a difficult text to work through because you have to read it very carefully and slowly and think through each step, but it is a very rewarding exercise.

PART V

5. SECOND ANSWER TO THE SKEPTIC

15.

INTERNALISM AND EXTERNALISM

At this point, it is important to distinguish two different general approaches to questions in epistemology: *internalism* and *externalism*. The difference is between figuring out knowledge “from the inside” (internalism) and figuring out knowledge “from the outside” (externalism). An internalist approach works with the resources each of us has as knowing beings—what we can sense, what we already believe, and what we can reasonably conclude from those things. An externalist approach instead takes a broader view of knowing beings, considering not just what is “inside” them but also their circumstances and their patterns of success in the past. It would not be far off to say that internalism is epistemology from a *first-person perspective* while externalism is epistemology from a *third-person perspective*.



An analogy might be helpful. Suppose you are asked to rate the performance of a certain plague doctor in 14th-century Florence. You might first ask *by what criteria* you are supposed to rate this doctor. Should you rate the performance according to the set standards of 14th-century medicine? If so, then perhaps the doctor did a great job; he rubbed onions and dead pigeons on the bodies of the plague victims and gave them vinegar to drink just as he was supposed to. This would be an *internalist* rating, or a rating based on more local information. But if you are supposed to rate the doctor according to today's medical practices, you will have to conclude that his performance was, well, not so great. This would be an *externalist* rating since the standards come from factors far

beyond the plague doctor's own beliefs and experiences. There is something valuable to be learned from each approach, though they are very different from one another and call upon different sets of facts.

So how does this analogy apply to epistemology? In this way. Sometimes we may find it important to understand the individual's local information: their own experience, reason, and beliefs and the processes by which they come to know what they think they know (*internalism*). Other times we may want to know, in fact, whether those processes *really do* deliver knowledge according to *what the rest of us on the outside* think we know (*externalism*). We might examine the individual's methods for forming beliefs in relation to what we know about their situation and what we know about the reliability of those methods.

Many philosophers have insisted that internalism best captures the meaning of "knowledge," and others have insisted that externalism best captures the meaning of "knowledge." But I see no reason to make any such insistence any more than I see any need to insist that one way of evaluating 14th-century Florentine plague doctors is better than the other. Each way gets something right depending on what we are interested in.

Over the last couple of chapters we have seen how difficult it is to offer good *internalist* answers to the skeptic. What this means is that it is hard for us to find within ourselves reasons or experiences that show that what the skeptic says is false. Phenomenalism is one route a person could take, trying to base all of their beliefs on sense data. Or they could follow Descartes's route and establish God's existence and God's good nature so that they can trust whatever seems to them clearly and distinctly to be true. These answers, as we have seen, only go so far in answering the skeptic.

But what about an *externalist* answer to the skeptic? Here there is much greater promise as an externalist can call upon a broader circumstance that reaches beyond our own experiences as individuals. Most significantly, an externalist can draw upon the circumstance that, as a matter of fact, we are usually *not* being radically deceived. To see how this works, suppose you are sitting in the library wondering if you are in a **GDD** scenario. An externalist comes along and asks what you are doing. You answer that you are undergoing a skeptical crisis, and the externalist says, "Nope, clearly you are not in a **GDD** situation. For you are in the library, sitting in a chair. If that is what it seems to you that you are doing, then you are right, my friend, for you really are doing it. Your beliefs

are true, and your senses are reliable, so you have knowledge, and the **GDD** is refuted.” We can then imagine the externalist sauntering away, whistling a jolly tune.

Now it might seem to you that the externalist is missing the point. What you want to know is whether you *really* are sitting in a chair in the library. But the externalist will say that, *yes, really, you are*; it is a fact. “But how do you know?” you ask. The dialogue continues as follows:

Externalist: I know this because my eyes are working well (I visited the optometrist just yesterday), and I know what chairs and libraries are (I just watched a stimulating educational video on the topic), and I see you sitting here in a chair in the library.

You: Well, you *think* you do! How do you know you’re not being deceived by an evil demon or a mad scientist?

Externalist: What an odd question! Do you see me standing here talking to you?

You: Yes, I *think* I do.

Externalist: Do you see any demons fluttering about or wires coming out of my head or mad scientists around me?

You: No.

Externalist: Do we have good evidence for thinking people are routinely deceived by demons or mad scientists? Is this something you see reported by credible media agencies?

You: No.

Externalist: Do you regard it as a sound epistemic practice to go around believing in stuff that you do not experience and is not reported to you by any credible source?

You: No.

Externalist: Then why on earth are you doubting whether we are here having this conversation?

In this riveting dialogue, the externalist is *changing frameworks* on you. Your doubt about being in the **GDD** scenario was arising from an internalist perspective, and the answer being offered by the externalist is coming (rather predictably) from

an externalist perspective. From that perspective, the **GDD** seems sort of silly.

Maybe you are still not convinced? Well, the aim of this chapter is to show the merits of externalist approaches in epistemology.

Media Attributions

- [Figure 5.1](#) © Cuddly Beast is licensed under a [CC BY-NC \(Attribution NonCommercial\)](#) license

G. E. MOORE'S HANDS

The British philosopher G. E. Moore (1873-1958) famously offered an externalist reply to radical skepticism. Recall that radical skeptics are not sure of the existence of any material objects, not even the existence of their own hands since they might merely be having the experience of having hands while, in reality, not having any hands. To refute such radical skeptics, Moore would sometimes dramatically hold up one hand and assert, “Here is one hand.” Then he would raise his other hand and say, “Here is another.” He would conclude from this vivid demonstration that there are at least two material objects existing in the so-called external world. And from this conclusion, he drew a further conclusion: that the external world exists.

Now it is very tempting to make fun of Moore in providing such a simple argument. But he was a fiercely intelligent person, and he knew exactly what he was doing. What he was doing was calling to everyone’s attention that the starting place of the skeptic—that I know my experience but not what lies

beyond it—is not in fact more obvious than the starting place of the non-skeptic—that, in fact, we typically do know that we have hands (as well as many other things). In some cases, we may not know that we have hands. If we have been in a terrible accident, for example, and the ends of our arms are wrapped in thick wads of bandages, then we may not be sure that we still have hands. But ordinarily we are quite confident that we have hands, and the burden is on the skeptic to offer some positive reason for thinking we are mistaken about this. If we are sure we have hands, then the proof of the existence of the external world is relatively straightforward: “Here is one hand ... and here is another.”

An externalist, following Moore’s line of thought, might well argue that the **GDD** is an abuse of language and the meanings of words (this particular line of thought is associated with Ludwig Wittgenstein). For consider, how do we learn how to use words like “know” and “illusion” and “deceive” and “doubt”? We learn them in very ordinary situations, in classrooms and at home and in theaters and on the street. We learn, for example, that stage magicians deceive us with hidden pockets and trap doors and sleight of hand. We learn about mirages and optical illusions. We learn about liars and cover-ups and conspiracies. In these ordinary situations, there aren’t any deep puzzles about what we know or what it means to be deceived. And we learn when it is appropriate to doubt or how

severe our doubt should be in particular real circumstances. It is not easy—it requires a lot of thought and experiment and so on—but the challenges we face are familiar and common.

Then we walk into a philosophy class, and we are asked to apply the concepts that we learned in ordinary circumstances to circumstances that are unlike anything we have seen before: deceiving gods and demons and mad scientists who exert malicious control over *everything* we experience. But our ordinary concepts are not meant to hold up in such extraordinary circumstances! It is a bit like learning that all numbers greater than zero are either even or odd, and then being asked whether *infinity* is even or odd. It is not just a hard question, it is an *impossible* one. To this extent, the skeptic raising the **GDD** scenario is pushing our concepts well past their breaking points. By *changing frameworks*, the externalist is trying to pull us back into the circumstances where those concepts are meant to apply.

17.

EPISTEMOLOGY NATURALIZED

The American philosopher W. V. Quine (1908-2000) offered an externalist epistemology in which our understanding of knowledge is based on a scientific understanding of our situation. Through science we are learning more and more about the natural world, about human psychology, and about the ways in which humans form beliefs about the world. Our scientific knowledge is not absolutely certain, of course—it may be that next week we learn that our current scientific theories are wrong in many fundamental ways. But contemporary science does represent *the best we have been able to do* so far (or let us assume this is so; more discussion of this will come in the next chapter). Is it not natural to use what we have learned about human psychology and the world in order to understand what it takes for human beings to know something?

This is what it means to *naturalize* epistemology: it is to see

epistemology as continuous with our science of the natural world, including the humans in it. When we ask, “How do we know that we have hands?” we should not seek some ground-shaking answer that will cause skeptics to run for cover. We should take the question seriously in the way that a scientist would. How do we know that we have hands? Well, our nerves are sending signals to our brains that indicate to us what our hands are doing and where they are, light waves are bouncing off our hands and entering our eyes, and signals are sent from our eyes to our brains where they are processed in such a way as to give us the belief that we have hands. Normally, our nerve signals are *extremely reliable* when it comes to telling us such things. If we doubt this, we could run an experiment with many people, some with hands and some without, and determine just how reliable our nerve-signal-processing functions are. In the end, we will find that it is virtually certain that those of us who think we have hands do in fact have hands.

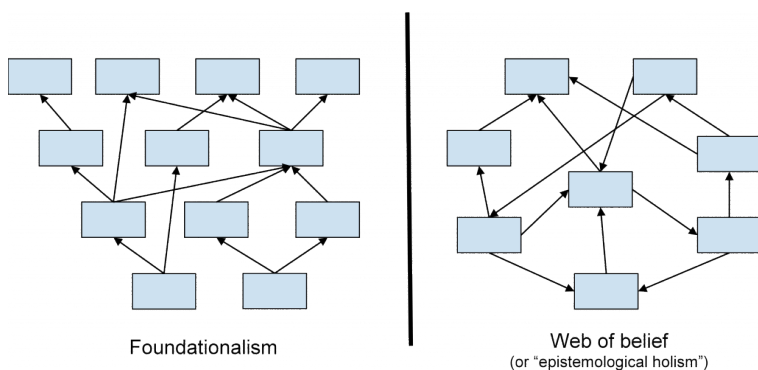
Objection: This so-called justification of our knowledge is circular. (Calling it “circular” is to say that it assumes what it is supposed to prove.) After all, if I am really doubting whether I have hands, then I am also doubting whether there are other people, and whether there is any scientific knowledge, and whether careful so-called scientific experiments really show anything about reality. But the externalist is supposing that we

do have all of this scientific knowledge and then uses that knowledge in order to justify that parts of that big system of knowledge—specifically, the parts that describe our nerve signals and brain functions—are trustworthy. The externalist is assuming our knowledge of the external world in order to justify the claim that we do have knowledge of the external world. How convincing is that supposed to be?

Quine responded to this objection of circularity. His response was that the objection arises from a mistaken view about how knowledge works. The objection supposes that there should be some basic and fundamental things about which we could not possibly be wrong—perhaps the cogito or the basic experiences that the phenomenalist appeals to—and that knowledge, in order to be knowledge, needs to be based upon these basic and fundamental things. We might call this view *foundationalism* since the view is that all knowledge, in order to count as knowledge, must be founded upon basic and fundamental beliefs we cannot possibly be wrong about.

Quine argued that this is a mistaken view of knowledge. He suggested instead that our beliefs about ourselves and about the world “hang together” in a kind of *web of belief*. In a web,

all of the strands and their connections rely on other strands and connections; the strength of the whole web is distributed across all its parts. Similarly, in a web of belief, a belief is supported by other beliefs which are supported by other beliefs which are supported by other beliefs including, perhaps, the first belief we started with. Our beliefs are in this sense *mutually supportive*. This is what it means to say they “hang together.”



According to Quine, we should not expect all of our knowledge to depend upon a few beliefs that are absolutely certain. Rather, we hope that our beliefs support one another in an overall coherent way. We still might regard some beliefs as very central to our web—meaning that many other beliefs depend on them, such as the belief that our senses are not deceiving us. But even these beliefs might be called into question if our other beliefs demand it. Suppose, for example, that we think we see a floating cat, and then our roommates

show us a very clever projector they are using to make the image of a floating cat. Now we have to decide whether to believe our roommates and their explanation or to believe our senses that there really is a floating cat and our roommates are lying to us for whatever strange reason. Our other beliefs—for example, that our roommates geek out over technological tricks, and floating cats are not commonly found within anyone's experience, and holding my hand in front of the projector lens makes the floating cat disappear—eventually persuade us to give up on the floating cat and to believe that it was only a projection. We can imagine different circumstances that would persuade us not to believe our roommates.

But here is another worry we might consider. The person who defends the web-of-belief view, or *epistemological holism* (as it is called), thinks that our beliefs all hang together in some mutually supportive structure. But might not two people each have mutually supportive webs of belief that disagree fundamentally with one another? Suppose one person believes that humans traveled to the Moon in 1969. A second person believes it was all a hoax. Each person has a mutually supportive web of beliefs supporting their belief about humans on the Moon: one person has all the beliefs we would expect including beliefs about film footage and reports from newspapers and NASA and so on, and the other person has beliefs about government conspiracies and cover-up

operations and movie sets made to look like the Moon and so on. We cannot fault either person with inconsistency. But clearly they cannot both have knowledge, can they? So what should we say?

(Can we say that they *both* have knowledge? So then, it is *true* for one person that humans traveled to the Moon and *true* for the other person that it was a hoax? Remember, we are not merely saying that this is what each of them *believe*. We are talking about *knowledge*. So we are saying that what is true may vary from person to person and not just for subjective things like favorite colors and banjo tunes, but for all sorts of things, including moon landings, ocean levels, and the shape of Mt. Fuji. Can we make sense of this? We will explore the idea further in the next chapter.)

Supposing for now that “true-for-you-but-false-for-me” is not an option, what is the epistemological holist to say about the Moon landing case? Quine, and indeed any naturalized epistemologist, would insist that the person who believes in the Moon landing is right, and the other person is wrong. Why? Because, as a matter of fact, *we did send humans to the Moon in 1969*, and all of the reports from the news and NASA are quite accurate. Remember the virtue of externalism: we can

appeal to facts outside an individual's sets of beliefs. We know the facts in this case, and we can trace how the person who believes in the Moon landing came to have their belief, and we can connect that belief to the facts. We can also trace how the conspiracy theorist came to have their belief and connect that belief to facts about paranoia and spurious claims made by other paranoid people. Case closed.

Objection: But wait! Who is to say that Quine's overall web of belief, which tells him that the Moon landing person is right and the other person is wrong, is the right web of belief to have?

Well, Quine would say, *we* are the ones to say—those of us who share a naturalized worldview. Of course, we will all admit that we might be wrong. But until some better web of beliefs comes along, we will keep on with the one we have. If in raising your objection you are expecting Quine to produce some fact as solid as iron that will favor one web of belief over another, then you are still using the mistaken foundationalist view of knowledge. We are always in the middle of working out our beliefs from our current web of beliefs, making adjustments where we need to, and trying to keep everything hanging together. There's nothing more a human can do so far as knowledge goes. And given what we have been able to work out so far, we can be pretty sure that humans landed on the

Moon in 1969, and people claiming otherwise are simply wrong.

Media Attributions

- Figure 5.2 © Charlie Huenemann is licensed under a [CC BY-SA \(Attribution ShareAlike\)](#) license

BACK TO BACON

We should recall Bacon's claim that "Knowledge itself is power." In this chapter we have seen that a second way to answer the skeptic is by changing our view of knowledge from internalism to externalism and from foundationalism to epistemological holism. The result is that we should believe the world pretty much is as we already think it is, and we should make changes in our beliefs only when some tough evidence comes along that forces us to make a change. But consider the effect of such a view for society as a whole. On a great many topics and questions, most people will end up with a dominant view of what is known. There will be some people in the margins who disagree with the dominant view, but those of us holding the dominant view will discount their beliefs as easily as we brushed aside the view of the conspiracy theorist at the end of the previous section. *Externalist epistemological holism* (such a name!) seems to give us *license* to reject the claims to knowledge made by those in the margins. And this should worry us.

Of course, we might say that the same view also gives the groups in the margins license to continue to maintain their webs of belief in the same way that those who believe the dominant view are maintaining theirs (for ease of discussion in what follows, let's simply call these two groups "the Margins" and "the Dominants"). But of course there is a considerable difference in power between the Margins and the Dominants. Who will decide what's taught in public schools? Whose knowledge will inform policy decisions? Who will get the jobs and grants for research and development? The Dominants, of course. Note also that externalist epistemological holism does not offer any reason for thinking that the Dominant view is *better justified or more thoroughly known or more true* than the Margin view. The only justification for the Dominants overruling the Margins is the fact that they *have the power* to do so.

But let us not write off the Margins just yet. If the Margins are able to produce experiments or problems or questions that the Dominants must deal with in some way, there is the possibility for the Margins' view to prevail. One might think here of the way in which Copernican astronomy (sun-centered solar universe) eventually replaced Ptolemaic astronomy (earth-centered universe). The history of this transition is long and complicated, but overall, when seen from high altitude, the Margin's view outperformed the Dominant view by criteria

the Dominants themselves shared (such as the value of accurate astronomical predictions). Similarly, the Darwinians overcame Aristotelian view in biology, and the Einsteinians overcame the Newtonians in physics.

And with these examples coming up, it is time to turn to science.

19.

QUESTIONS TO CONSIDER

1. This chapter has included several technical terms. It may be instructive to write down definitions for them.

internalism

externalism

naturalized epistemology

circularity (in justification)

foundationalism

epistemological holism

2. Consider your definitions for externalism and epistemological holism. Do you think there could be such a thing as *internalist epistemological holism*? Think it through, and describe what such a position would maintain.

3. As it has been written, “Our other beliefs ... eventually persuade us to give up on the floating cat and to believe that it was only a projection. We can imagine different circumstances that would persuade us not to believe our roommates.” Please describe these different circumstances, and show why they would lead us to believe it is more likely that there is a floating cat than that our roommates have deceived us.

20.

FURTHER READING

G. E. Moore, “Proof of an External World” in his *Philosophical Papers* (New York: Collier Books, 1962), pp. 144-148. Moore’s essay can also be found at multiple sites on the internet.

W. V. Quine, “Epistemology Naturalized” in his *Ontological Relativity and Other Essays* (New York: Columbia UP, 1969), pp. 69-90.

“Internalism and Externalism,” an article in the *Internet Encyclopedia of Philosophy*, offers an excellent overview of the strengths and weaknesses of both views and also features a helpful bibliography.

PART VI

6. SCIENTIFIC KNOWLEDGE

Scientific knowledge is a great human achievement. Because of this knowledge, we can successfully navigate around our world and describe its size, shape, and mass. We know our world is a *planet* orbiting a star. We can say how fast light travels and how long it would take for light to travel to the next nearest star, which we know is smaller than our sun but denser. We know how to remove a heart that isn't working and replace it with one that is—and possibly, with one we have manufactured ourselves. We know how to replenish soil so that it continues to sustain crops, how to smash atoms together to re-create energy levels comparable to those at the beginning of the universe, and how to inoculate against smallpox. Of course, the list could go on and on.

And on and on it should go lest we underestimate just how much we know. We know that osmium is the densest stable element, that Triceratops lived about 68 million years ago, that the eruption of Krakatoa ejected six cubic miles of rock into the sky, and that the opossum is North America's only native

marsupial. We know human blood comes in various types, and we know what types of blood may be transferred from human to human without deleterious consequences. We know how to turn lead into gold (yes, *really*, but it costs a lot), how to generate x-rays and block them, how to keep subatomic particles in a superposition, and how to arrange them so as to perform calculations across several possible worlds at once.

That is a lot of knowledge. And in this chapter, when we turn eventually to asking whether scientific “knowledge” is actually knowledge, we would do well to remember this astonishing list that captures barely a sliver of all the knowledge we frequently take for granted. Scientific knowledge merits special attention in epistemology because—on first glance, at the very least—it is the greatest tradition of knowledge-getting in all of human history.

21.

LOGIC, MATH, AND SCIENCE

We should begin by situating scientific knowledge among other types of knowledge. As we have seen, some known propositions are *contingent*, meaning that they express facts that easily could have been otherwise. As it happens, for example, some lucky person in Canyon County won the Idaho Powerball jackpot in August 2017. That could have been otherwise—someone else could have won, or the prize might have gone unclaimed for another month. It snowed last Thursday (again), so weather patterns might have shifted slightly, bringing the snow sooner or later or not at all. Of course, for these events to have been otherwise, the particular causes would have had to have been different, and for those particular causes to have been different, *their causes* would have had to have been different, and so on. But none of these changes seem impossible. Each change is “thinkable” or imaginable on its own.

Other known propositions are *necessary*, meaning that they really could not be otherwise. So in geometry, for example, a cube contained within a sphere has less volume than the sphere. It's hard to get around that fact—there is no way it could be otherwise, given the meanings of the terms we are using. In these causes, it is not a matter of re-engineering causes to bring about different effects. The changes themselves are unthinkable or unimaginable.

(Someone might wonder if this cube/sphere proposition might be considered contingent since it is, after all, contingent that the words “cube” and “sphere” mean what they do in English. Clearly, those words might have meant different things. It is a good question. One reason philosophers like to talk about *propositions* is that a proposition is supposed to be the *meaning* of what is said in whatever language. So, yes, the sentence, “The cube is in the sphere” might have meant many different things or nothing at all; but the *proposition* that *the cube is in the sphere* means precisely one thing: the thing that is also meant when we say, “el cubo está en la esfera” or “a kocka a gömbben van” or “tha an ciùb anns an raon” which all mean “the cube is in the sphere.”)

The known propositions of logic and mathematics are

necessary. How do we know they are necessary? Is it simply a matter of what we can or cannot imagine? This is a *very* good question. A first answer might be that we know these propositions are necessary because if we deny them, then we can derive a contradiction from them. So, for example, *five plus three equals eight* is a true and necessary proposition. If we try to deny it, we end up in the following sort of trouble:

$5 + 3 \neq 8$	(suppose)
$(xxxxx) + (xxx) \neq (xxxxxxxx)$	(by definitions of
“5”, “3”, and “8”)	
$(xxxxxxxx) \neq (xxxxxxxx)$	(by definition of
“+”)	
$8 \neq 8$	(by definition
of “8”)	
ABSURD!	(by definition
of “ \neq ”)	

So we might say that necessary propositions are those whose denials entail contradictions or results that are false in virtue of the meanings of the terms. Perhaps this is a good enough answer. But some philosophers—notably W. V. Quine whom we encountered in the last chapter—have wondered whether “meanings of terms” are fixed in such precise ways as to allow for a clear distinction between necessary and contingent truths. Don’t we learn the meanings of terms in rather informal circumstances, which allow for quite a lot of slippage

and unclarity and vagueness? So, for example, what about the claim that Catholic priests are male? Is that true in virtue of the meanings of the terms, or is it a contingent truth based on decisions made by a particular tradition? Is it obviously “more necessarily” true than the claim that some dogs have tails? We might wonder whether there really are hard-edged “meanings” of terms that allow us to definitively determine whether a given claim is necessarily true or contingently true. This objection is generally known as “Quine’s criticism of the analytic/synthetic distinction” and it is an interesting and important discussion to study, but it is a bit beyond our reach in this introduction.

So, having mentioned that objection, I will now set it aside, and continue as if we have some good way of distinguishing necessary truths from contingent truths. Logical truths (such as “ $P \Leftrightarrow P$ ” or “if $P \Rightarrow Q$ and $Q \Rightarrow R$, then $P \Rightarrow R$ ” or “if $P \vee Q$ and $\sim Q$, then P ”) and mathematical truths (including all those in arithmetic, geometry, algebra, calculus, etc.) are *necessary* truths. We know they are necessary truths because if we try to deny them, we will find that we can derive claims that are false in virtue of the meanings of the terms. Other particular facts about the world such as what happened here or there, how long some particular rhino’s horn is, or who stole the cookies from the cookie jar, and so on, are *contingent*. We know they are contingent because if we deny them, we will find

that we can derive claims that, in fact, are false but not false in virtue of the meanings of the terms. Denying that Slim Jim won the Idaho lottery, for instance, might make it harder for us to explain how he was able to afford a shiny new Cadillac, but it will not lead us to derive claims that are false by virtue of the meanings of “Cadillac,” “lottery,” or “Idaho.”

Now what about known propositions of science? In particular, what about the known propositions we identify as *laws of nature*? These known propositions seem to be somewhere between necessary and contingent. We can deny them without running into contradictions about meanings of terms. So, for example, suppose it is a law of nature that force equals mass times acceleration. Indeed, this was once thought to be a law of nature known as “Newton’s second law of motion.” It was thought to be a rock-solid truth, one perhaps that could not be otherwise. But since then, we have learned that this law not only *could be* otherwise, but *it actually is* otherwise since, to update it to Einstein’s theory of relativity, we need to complicate the equation a bit (i.e., taking into account how fast the observer is moving relative to the speed of light). Einstein made this advance upon Newton without running contrary to any of the meanings of the terms involved: “mass” still meant “mass,” but the relation to force and acceleration turned out to be a little different. So it is evidently possible to deny Newton’s second law of motion without

entailing a contradiction. Moreover, it is possible to deny Einstein's laws without entailing a contradiction. And, indeed, any of the known laws of nature can be denied without entailing a contradiction.

Does that make the laws of nature *contingent*? It does if we hold fast to the claim that contingent propositions are the ones that can be denied without entailing any contradiction. But at the same time, there is something about laws of nature that make them seem similar to necessary truths. Laws of nature are *more fundamental*—"closer to the core of reality," so to speak—than contingent facts about particular things. When scientists discover basic laws of nature, they are getting at deep truths about reality, truths that could be different only if reality itself were different in some fundamental way. This depth of the truth of laws of nature makes them seem similar to truths in logic or mathematics, which also could be different only if reality itself were different in some *really* fundamental way.

Perhaps an example will make this idea clearer. Suppose we inflate a balloon until it bursts. We can imagine all sorts of ways to vary this exciting experiment: we could use thicker or thinner balloons, we could use different sorts of gases, we

could inflate the balloon more or less quickly, we could do it on mountaintops or down in the valley, during the day or night, and so on. These are changes we can easily make. But suppose that instead of making any of these easy changes, we want to keep everything exactly the same but *delay the bursting* for an extra minute. That is to say, we want to use the same sort of balloon, the same gas, the same outside pressure, the same rate of inflation but just delay the bursting by a minute. To do this, we will have to change some *natural fact* that has to do with the strength of the balloon material. We will have to change a deep fact about the nature of the world and, specifically, about how much that sort of material can stretch before ripping. That's *really* hard to do. In fact, for creatures like us, it is *impossible*, for humans cannot alter the laws of nature that govern the limits of materials.

Of course, *in words*, or *conceptually*, we can deny whatever law of nature that is involved in this experiment, and our denial will not entail any contradiction. But we cannot deny or change the law *in fact*. We cannot *really* make it false. There is a sort of necessity to the law of nature that simply is not found in the other particular circumstances, all of which we are able to change by using different balloon materials, a different gas, different altitudes, and so on.

So the denial of a law of nature is impossible, but for some reason other than that the denial of the law entails a contradiction. It would be interesting to continue to pursue this line of thought, but once again, this is a topic that takes us quickly into matters beyond the scope of this introduction. For our purposes, we might simply recognize three types of necessity: *logical necessity*, *mathematical necessity*, and *natural necessity* which is the sort of necessity pertaining to scientific laws of nature. The differences among these kinds of necessity are philosophically interesting, but we won't pursue the topic here.

22.

HOW DO WE KNOW THE LAWS OF NATURE ARE TRUE?

This question might be asked in two tones of voice: “how do we know?” in the sense that *maybe we don’t*, and “how do we know?” in the sense of *really, how do we manage to know that they are true?* Let’s take that second sense first. What sorts of observations and reasoning gives us reason for thinking a *supposed* law of nature is actually true?

We should begin by noting that laws of nature are *general* statements about *regularities* that objects and their features *must* obey. Force is *always required* to cause a mass to change velocity. Pressure is *always* directly proportional to temperature and inversely proportional to volume. Overall entropy *never decreases* in an isolated system. Laws of nature are lawlike in the generalities they describe and in their force. They do not invoke specific particular objects, like the phone in your pocket or the top of Mt. Fuji or Larry from down the

hall. Laws of nature state *general truths* which may apply to all relevant particulars but never focus on any specific particulars.


But, of course, *all we ever see* are specific particulars. And, as noticed back in our discussion of phenomenalism, we see that events happen, but we never see the necessity in their happening. So how do we arrive at knowledge of laws of nature?

One might begin by thinking that we arrive at knowledge of laws of nature by observing a phenomenon repeatedly and seeing what regularities hold. So, for example, a line of frogs comes our way, and the first one hops, and the second one hops, and the third one hops, and so on, and eventually we begin to think that maybe “all frogs hop” is true. That is a general statement, and it seems to be based on observing a string of particulars.

But there is a logical problem. It simply *does not follow* from the claim that one frog hops and a second frog hops and a third one hops that *all* frogs hop. Moreover, the observation that all the frogs observed so far hop does not demonstrate any *necessity* in their hopping. Consider for contrast the situation

in which an old man in front of you in line at a convenience store pulls one penny out of his pocket and then pulls another one out of his pocket and then a third one... As frustrating as the experience may be for you, you should not conclude from it that *all* coins in the pockets of old men in convenience stores *must* be pennies. Your tedious experience does not allow you to boast that you have discovered a new law of nature.

Of course, doing real science is harder than just watching frogs hop or old men count out change. Scientists observe carefully, employ control groups, and run tests to find accidental correlations. But even so the logical fact remains that no number of particular observations *show the truth* of a general and necessary claim. Put another way, a general truth, like a law of nature, is always **underdetermined** by its evidence. This, by the way, is known as **the problem of induction**, and it was made famous (or infamous) by David Hume (1711-1776). To make the point in another way, suppose we make five careful observations. Which of the dotted lines most accurately portrays the law of nature we have discovered?



the problem of induction

They *both* do; each one models our data. The problem though is that there just isn't enough data to tell us what to do with the spaces in-between the data points. This is what is meant by "under-determination." The two different ways of construing what is going on make different predictions about what patterns we will find as we make further observations. Further observations may help us to rule out some possibilities. But no matter how many further observations we make, we will always have many different ways of connecting the dots.

So we do not easily "read off" general statements from the data. What do we do? At this point it would be sensible to admit that our observations of particulars do not *demonstrate* what the laws of nature are. Instead, perhaps they only *give us some*

idea of what sorts of correlations there *might* be in the world, and based upon this idea we frame a *hypothesis* about what the laws of nature are. The hypothesis, we will admit, may be true or false, but it is a good guess based on what we have seen so far.

This suggestion obviously brings us to the famed **scientific method**. The method, in short summary, is this. Begin with some observations; frame a hypothesis; generate a prediction from the hypothesis; devise a test to determine if your prediction is accurate; if it is – well done, keep testing; if it isn't – then start over. We need not go into further detail for our purposes, and the method is probably already familiar to most readers.

Note, however, that the scientific method does *not* deliver *certain* knowledge of the laws of nature. Each known law of nature is a hypothesis that has *not yet* been decisively refuted by experiments. That is the most that can be said of the hypothesis, and it is clearly not the same as saying that some law of nature is known with certainty. But perhaps that is okay. As we have seen, knowledge need not be certain in order to count as knowledge; it need only be a *justified, true belief* (recall **JTB**) whose truth helps to explain why it is believed (**JTB+**).

So long as we are *justified* in our belief of the scientific claim that has been made (which we have not yet discussed), then it will turn out that we know the claim – just so long as it does turn out to be true, and the **JTB+** conditions have been met.

This gives us occasion to reflect on the sentence “It’s *only* a theory.” Many times people use this phrase to remind us that the theory of evolution, for instance, is *only* a theory. It is implied that when something is *only* a theory, it is not known with certainty, or perhaps not even known very well at all. But once we recognize that very little of our knowledge of nature is known with certainty, we should reject the implication behind saying “It’s *only* a theory.” If “theory” just means *not known with certainty*, then basically all of the knowledge of nature we use to build bridges, cure diseases, manufacture cell phones, and boil water is *only* a theory. And in fact, this is not a good way to understand the term “theory” anyway. A common definition of the term is that a theory is “a supposition or a system of ideas intended to explain something, especially a system based on general principles independent of the thing to be explained.” Note that nothing is said about *certainty*. A theory is an explanation-provider; whether any particular theory is plausible or not, or very likely true or not, or even certain or not, is a completely separate matter.

This admission, that we can never know with complete certainty whether our claims about the natural world are true, is called **fallibilism**. Our knowledge is fallible, which means it might turn out not to be knowledge at all. It is a good guess we are working with until we have evidence that it is false.

Media Attributions

- Figure 6.1 © Charlie Huenemann is licensed under a [CC BY-SA \(Attribution ShareAlike\)](#) license

23.

BUT HOW DO WE KNOW THE LAWS OF NATURE ARE TRUE?

Now we can turn to the other tone of voice in which the question can be asked. This other way of asking the question asks whether our alleged knowledge of the laws of nature *really should* count as knowledge. Does science count as *genuine* knowledge?

As we have seen, if alleged knowledge should count as genuine knowledge only when we have the sort of absolute certainty that would shock a skeptic into amazed silence, then we do not have genuine scientific knowledge. But in that case, we have hardly any genuine knowledge whatsoever, and we have seen where that leaves us. So suppose we lower our standards a bit so as to allow many of the items we typically regard ourselves to genuinely know: that we have hands, that our senses typically do not deceive us, that regular and constant patterns in our experience (like apples falling and clothes drying and water

freezing) will continue to hold in the future, that other people exist, and so on. Suppose, in short, we adopt a more ordinary attitude toward our experience. Given our ordinary attitude toward our experience, the one we have when an extreme skeptic is not pestering us, do we have good reason for counting our alleged scientific knowledge as genuine knowledge?

It might seem that the answer is obviously **yes**. After all, every scientific claim is accompanied by a list of observations and experiments that (in principle) anyone can access or perform. These observations and experiments belong in the sphere of our ordinary knowledge of the world. The scientific claims or theories are based upon those observations and experiments through the general method outlined in accounts of “the scientific method.” This makes scientific knowledge *continuous with* our ordinary knowledge of the world and just as genuinely known. The idea here is that we can take any reasonable person who trusts their experience and show them step-by-step how we have reached the scientific conclusions we have reached and why these conclusions are reasonable to believe even if, as we admit, there is always the possibility that we might be wrong about some or all of it.

But here is a line of objection that might be presented. The central idea is one articulated by Thomas S. Kuhn (1922-1996), though it should be pointed out that Kuhn himself never used this idea to criticize scientific knowledge. Still, others influenced by Kuhn have used the idea to this end. The idea is this: scientists are always coming up with the best theories they can given the ideas they have, the evidence at the time, and what they are interested in. When scientists come up with these theories, they are creating a *paradigm* or a worldview that basically says “here are the problems we are interested in, and here are the methods we should use to solve them.” Then teams of scientists normally get to work trying to solve the problems they are interested in. We can call this period of normal scientific activity *normal science*.

Every so often some scientist develops a new paradigm or worldview that is *radically different* from the one being employed in normal science. The paradigm presents a wholly different view of the world, a different set of problems to solve and a different set of methods to use in solving them. Sometimes these new paradigms just fail—no one else gets interested in them, they don’t work very well, and they sputter and die. But sometimes younger scientists get very excited by the new paradigm. They regard it as an exciting and useful new development, and they begin to use it and promote it to others. Over time the new paradigm may take the place of the

old paradigm, and that is when a *paradigm shift* or a *scientific revolution* occurs.

What causes a paradigm shift to occur? What explains the success of a new paradigm? For a long time historians of science believed that a new paradigm is successful when it allows for better predictions or better methods of solving problems. But Kuhn's work in the history of science showed that this is not so, or at least, not always so. Remember, the new paradigm is radically different from the old paradigm. This means that the standards of what counts as "better predictions" or "better methods" also changes. There is no common standard of measure between an old paradigm and a new paradigm. For this reason, Kuhn called paradigms *incommensurable*: there is no meaningful way to compare one to the other.

So a paradigm shift does not happen because the new paradigm is clearly *better* than the old one. Rather, Kuhn argued, a paradigm shift happens because purely human and historical conditions give the new paradigm an advantage over the old one. In the simplest possible case, the old guys defending the old paradigm eventually die, and the younger guys with the new paradigm get their jobs causing the

paradigm shift to occur. In more realistic and complicated cases, there are political and economic and ideological pressures that all come into play and end up favoring the new paradigm over the old one. But these pressures do not guarantee that a new paradigm will be “better” in terms of being better knowledge of the world.

The end result of this—though again, not one that Kuhn himself embraced but one embraced by scholars influenced by him—is that scientific progress is an illusion. There are changes in scientific theories, of course, but the changes are not brought on by objective measurements and experiments. They are brought on by social pressures. Science, then, is sort of like fashion or the evolution of styles in art. Attitudes and styles change, and people with the new styles and attitudes view them as improvements, but really the change does not indicate that the new attitudes and styles are closer to the *truth*.

Objection: This could not possibly be true. After all, do we not have better medicine, better technology, and more thorough explanations of nature than any previous generation? Read again the first two paragraphs of this essay!

But here is a reply: perhaps the advances in medicine and

technology *would have happened anyway* under the terms of the old paradigms, and the new paradigms really had nothing to do with those advances. And so far as “more thorough explanations” go, that judgment is being made from the perspective of the new paradigm. The judgment is biased toward the new view. The old paradigms also had very thorough explanations, though, of course, using different terms and ideas. What reason do we really have for believing our paradigm is better than theirs?

This last question should be taken seriously, and not just rhetorically. *Are* our modern scientific theories better than previous scientific theories? Could previous theories, in principle, make sense of the technological advances that have accompanied modern theories? Rather than simply concluding “well who knows? Maybe!” from the comfort of our armchairs, we might actually try to determine whether, for example, Aristotelian science could allow us to make sense of gene therapy. What new advances would have to be made by succeeding generations of Aristotelian scientists? What further elaborations would have to be made to their theories? In the end, would the revised Aristotelian theories be fundamentally different from our modern theories? Or would they just be the same idea in different words? Can someone come up with distinct neo-Aristotelian alternatives to the many ways we understand the natural world around us without simply

repackaging what we think is true in Aristotelian-sounding language?

To date no one has really taken up this challenge in a thoroughgoing way. Of course, we can never know what twists and turns alternative histories might have made, and there is always the possibility that inventive Aristotelians could have kept their tradition going and perhaps could have led to even more impressive technological achievements. But one might well ask *what evidence we have* for believing this possibility is real and whether it is stronger than the evidence we have for believing contemporary scientific knowledge to be genuine. It is not enough merely to *claim* that a rival paradigm might have enjoyed equal success in controlling and predicting the natural world; one must *show* that it is true. Until that challenge has been met, we do not have good reason to think that the old paradigms are “just as true” as the newer ones.

24.

SOCIAL CONDITIONS OF SCIENTIFIC KNOWLEDGE

But even if the more dramatic claim that there has been no scientific progress does not seem compelling, there is a valuable lesson to be drawn from Kuhn's historical argument. The valuable lesson is that science does not develop in a vacuum. Scientists are human beings, and scientific institutions have connections to funding agencies, economics, politics, and culture at large. At any given time, multiple pressures are affecting how science develops. Some of them are "proper," having only to do with evidence, observation, experiment, and theoretical integrity. Some of them have less to do with a concern for scientific truth and more to do with the human ambitions and prejudices of the scientists or their bosses.

An example is the 50-year struggle to recognize the toxicity of lead in gasoline. Lead was introduced into gasoline in the 1920s to stop engines from making knocking noises. It was

already well known that lead was harmful to living organisms and made people behave erratically, but the scientists employed by fuel companies insisted that the levels of lead in gasoline were safe for human beings. They were not, and the evidence was manifestly clear that the levels of lead were unsafe. Committees were formed and studies were performed, but the results of the studies for several decades was that lead should continue to be put into gasoline and further research should be done. In the 1970s, lead was eventually banned by the newly-formed Environmental Protection Agency. By that time, the average level of lead in people's bodies in the US was well beyond safe levels, children were underperforming in schools, and crime rates were rising as a direct result of lead poisoning. A similar story can be told of the ways in which scientific studies commissioned by oil companies skewed the data to suppress information about global warming for decades.

In such cases claims to scientific knowledge have been shaped far more by economic considerations than by a proper concern for genuine knowledge. Of course, it is also true that we eventually learned of the effects of leaded gasoline and the effects of carbon emissions on our atmosphere precisely through scientific inquiry once it was freed from the distortions of economics and politics. So the cases of science being distorted by social conditions are not enough to discredit science as a whole. But they are enough to cause us to examine

claims to scientific knowledge with some awareness of their social contexts.

And, of course, this point does not apply only to scientists but to all of us. Social conditions shape human knowledge. We will turn to this topic in the next chapter.

25.

QUESTIONS TO CONSIDER

1. Someone might try to solve the problem of induction in the following way: “Over time, we have found that inferring a generalization from a number of particular observations has worked, so we should be able to trust doing the same thing now and in the future.” Why doesn’t this solve the problem of induction?
2. I often hear people claim that science can never know “capital-T Truth.” What on earth does this mean, and why would someone think it is true (or “True”)?
3. “We have excellent historical evidence for believing that the claims of science are not based on objective evidence, but are instead just based on the prejudices of the time and place from which they come.” Comment on this assertion—in particular, consider whether it is somehow self-contradictory (*how?*).

FURTHER READING

There is a massive literature on the philosophy of science and the project of constructing science from observations and logic. An overview of the issues very briefly mentioned in this chapter can be found in Peter Godfrey-Smith's *Theory and Reality* (University of Chicago, 2003) and virtually any textbook on the philosophy of science.

Thomas Kuhn's book *The Structure of Scientific Revolutions* (University of Chicago, 1962) has been enormously influential. There are also multiple overviews of his argument available on the internet, including a useful entry on *Wikipedia*.

A more radical view of scientific progress can be found in Paul Feyerabend's *Against Method* (New Left Books, 1975). As his title suggests, Feyerabend argues that there is no scientific method, and our best bet is to try a very broad range of

approaches to understanding the world and see what happens—“anything goes” is his slogan.

PART VII

7. SOCIAL CONDITIONS OF KNOWLEDGE

Imagine having the opportunity to spend a year studying abroad. Of course, many students do this, and they experience life in nations and cultures around the world. They learn what it is like to live in another culture. They make new friends, adopt the local language, celebrate local holidays, learn what the traditions are, and perhaps learn a new version of what is regarded as “common sense.” Travel broadens the mind by teaching us how much of our mental lives is due to just being in one culture rather than another.

We might also try to imagine what it would be like to “study abroad” in other cultures throughout history. Imagine spending a year in Mesopotamia five thousand years ago or in Ancient Rome or in the Mayan Empire or in Japan in the 12th century. Such experiences would broaden the mind to an even greater extent since the “common sense” of these cultures would be so radically different from anything we know or can

even imagine. The locals would regard us as bizarre, strange-thinking aliens, and we would have to work extremely hard to learn what to say, what to assume, what to eat, and what customs we had to follow. Going to school—if there were schools—would raise another cluster of problems as we would have to catch up on the strange (to us) things our companions already knew, and we would have to get a sense for what the “problem space” of knowledge was. Were there gods or magical forces we need to take into account? Is there a creation story that plays an explanatory role? What sorts of questions can we raise, and what questions would be weird or offensive to ask?

After spending a year in another historical culture so radically different from our own, we might be shocked when we returned home. What once seemed familiar would seem extraordinary. What seemed so obvious would now seem novel and arbitrary. We might try to imagine all the difficulties a friend from the other culture would have as they tried to adapt to our world—what they would find weird, baffling, or ridiculous.

One thing is for sure: the study of other cultures, present and past, helps us to learn the importance of social conditions for

knowledge, which include all of the things we would find surprising as we hop from one culture to another.

27.

OBSERVATIONS ARE THEORY-LADEN

An important lesson we would learn through our imagined study-abroad experience is that it is hard to speak at length about anything without making it obvious that we have a particular theory about the world. Imagine walking down the road three thousand years ago with your northern African friend, Akil:

You: Boy, is it hot today! The sun is really beating down.

Akil: It is indeed! It's a good thing Re is so powerful.

You: Re? Oh, yeah, the sun. Why is it good that Re is powerful?

Akil: At night Re goes beneath the land to battle the forces of chaos. If he didn't fight so fiercely, we would have many more problems—food shortages, rebellions, fighting, you name it. Re's power helps to make sure life on the land continues as normal.

You: I agree the sun is really important. It sends

energy to our planet, warming our atmosphere and giving the plants energy to grow.

Akil: You talk so funny! You make it sound like the sun is just a big disk of fire.

You: It is—or at least a big fire sphere, many times bigger than the earth. And it doesn’t move “beneath the land.” It’s just that the earth turns and makes it look as if the sun is moving.

Akil: Obviously not! (He holds his hand up against the sun—or Re.) While Re is massive and powerful, I would guess he is about a half setat in size—plenty big enough to give the forces of chaos a good fight! And I don’t know why you would think the land is moving. Do you feel it moving? In your view, why would a big sphere of fire care about us and make the plants grow and keep our life free from chaos?

You: The sun just burns. It doesn’t care about anything. It just does what it does. Look, sometimes chaos happens, right? Even when the sun is shining?

Akil: Sure. The battles go back and forth, and sometimes the forces of chaos get an upper hand but never for long. How would you explain the fact that chaos is always defeated? How do you explain how the “sun” in your view helps plants to grow?

You: It’s complicated. Chaos gets defeated just because—well, there are many different cases, but wars have to end sometime, and peace has to happen. People

just get tired of fighting, I guess. Plants grow because of photosynthesis...

Akil: Foto Sin Theseus? Is he one of your gods?

We can imagine the discussion going much further and becoming ever more complicated as you and Akil try to fathom how you can both look at the one thing and see such different things. You see a massive star fueled by nuclear fusion, and Akil sees a divine person whose energy and concern for life infuses everything. In this sense, what each of you sees embodies a certain theory you believe. Philosophers call this “**the theory-ladenness of observation.**” It means that every observation carries some sort of theory along with it. The observation is connected to a background theory about what the observation is an observation *of*. More formally, we may say that “observations are theory-laden” means that the terms or concepts used in the observation have their meanings by virtue of some background theory.

Objection: But surely not *all* observations are theory-laden. There is a clear sense in which both we and Akil, in our imagined example, are seeing the same thing. We are both seeing a very bright disk in the sky, right? And then each of us has more to say about it. Akil says it is a sky-traveling divine person, while we say it is a huge, distant, uncaring star. But both of us will at least agree on the basic observation, right?

Perhaps this is so. But note that neither you nor Akil would count the claim “The sun is a very bright disk in the sky” as *knowledge*. Akil would say the claim is false, and perhaps even sacrilegious. You would insist that the claim is literally false, as the sun is not a disk, not very bright (relative to other stars), and not in the sky. But still, we might say, would you not both agree that the sun *looks like* a very bright disk in the sky? Perhaps, but you would both quickly explain that looks can be deceiving, and the truth is more complicated. So, if the claim is to count as an observation, it counts only as a *misleading* observation.

In other words, any claim that we confidently count as “knowledge” will be a claim that is tangled up with quite a lot of theory. The theory is in large part, if not entirely, a product of culture. Remember from the last chapter that, according to the scientific method, much of our scientific knowledge consists of hypotheses we have developed as we try to explain the natural world. These hypotheses do not develop in a vacuum but are drawn from our background learning, our community with other scientists, and our sense of what the scientific project is all about. In the imagined dialogue, for example, Akil’s question about why a big sphere of fire would “care” about us will seem to us like a wrong-headed

question—not the sort of question we are likely to pursue as a research project—since the very idea of astronomical objects like stars having “concerns” is well beyond the sorts of questions we are encouraged to ask. This sense we have about which questions are of the right sort and which ones are wrong-headed has very much to do with what we conceive the scientific project to be. That one question Akil raises says a great deal about the great distance between his worldview and our own.

28.

PERNICIOUS BACKGROUND THEORIES

England in the 19th century was proud of its scientific attitude and achievements. Steam engines, calculating machines, automated factories, and advances in medicine gave the Victorians much to brag about. Theirs was an age of dependable, fruitful, scientific knowledge. They knew, for example, that women are, by nature, weaker than men in mind and spirit and are prone to chemical and psychological imbalances. When excited by too much activity or difficult thought, women typically become hysterical, which is a psychological and physiological condition brought on by the sensitive nature of their reproductive organs. The only treatment is decreased activity, less exposure to new ideas, and doses of opium as prescribed by more rational men. Similarly, it was clear to the Victorians that evolution had endowed the English with greater skills and grit which had led to an empire on which the sun would never set as it had spread around the globe. Other races clearly had not evolved to an equal degree,

and it was the obligation of the white Europeans and Americans to help lead the lesser races to further degrees of civilization—for their own good, of course.



[From Wikipedia: “This cartoon depicts a representation of Rudyard Kipling’s famous poem ‘The White Man’s Burden.’ Originally published in February of 1899, the poem’s philosophy quickly developed as the United States’ response to annexation of the Philippines. The United States used the ‘white man’s burden’ as an argument for imperial control of the Philippines and Cuba on the basis of moral necessity. It was now the United States’ moral duty to develop and modernize the conquered lands in order to help carry the foreign barbarians to civilization.”]

The Victorians used these “scientific theories” regarding sex, race, biology, psychology, and evolution as justification for a

wide array of oppressive practices. But only in a few cases did knowledgeable people *knowingly* use these theories as some sort of “cover” for justifying racism and sexism. Rather, in most cases, knowledgeable people *really believed* they were seeing the world through clear lenses of science. They thought they were seeing women and people of color as they really were, or as nature had evolved them to be. It was perhaps regrettable (they might say) that nature was so unfair, giving so many benefits to some segments of humanity while leading other segments so backward and incapable, but that was why it was the duty of the superior humans to help along the inferior humans. As 19th-century scientists studied the physiologies, psychologies, and social structures of women and people of color, they “saw” what their background theories told them to expect to see: weaker, inferior creatures who could not help being who they were.

This is a clear case of observations being shaped and skewed by background theories that were pernicious and horrible. It is worth taking a moment to imagine what it would be like to be the target of these practices and theories. It would not only be the experience of sexism and racism, which is bad enough. The racism and sexism would be *built into* the culture: in the schools, in the medical books, in the training of all professionals, in popular lectures and newspaper articles, all promoted in just the way any scientific theory of atomic

elements or electrical power would be promoted. Any woman or a person of color who believed themselves capable of doing what white men could do would be denying scientific *fact*, being *irrational*, and refusing to believe what observations *plainly show*. Their rebellious attitude would be seen as a problem that needed fixing, either through drugs or through confinement in a prison-like asylum which would restore the deluded person back to “health.” And there would be no court of higher appeal to hear this person’s case as the entire society was equally “enlightened” by the science of the day. Everything in society would be telling a woman or person of color that they were crazy or stupid not to see themselves as inferior human beings.

We should all be grateful that we know better now (even while we also recognize that the legacies of these prejudiced views continue to shape practices and institutions). But if we stop at that point of gratitude, then we have not learned the full lesson. The full lesson is that the results of science *always* can be shaped and skewed by the prejudices, biases, superstitions, and inequalities of society. Or is it more accurate to say that the results of science *always will be* shaped and skewed by societal prejudice? For the only way in which science can be saved from these prejudices is if active steps are taken to confront and challenge those prejudices in the society at large. An individual scientist or group of scientists cannot merely resolve to *try*

hard not to be prejudiced. The Victorian scientists, after all, were trying very hard to be impartial and fair, and we can see where that led them. Rather, the societal prejudice as a whole must be challenged in order for the science that reflects that society not to be prejudiced.

Media Attributions

- [The White Man's Burden Judge 1899](#) © [Victor Gillam](#) is licensed under a [Public Domain](#) license

MORALITY OF KNOWLEDGE

This brings us to a range of important and difficult questions that connect our interest in knowledge with our interest in morality and social justice. We have seen that it is either impossible or extremely difficult for a society's knowledge *not* to reflect the society's own prejudices. Theories are shaped by a host of factors, and the attitudes and values of the surrounding society are counted among them. So, as a society begins to confront its own moral prejudices or skewed values, what effect should that have on the society's pursuit of knowledge? Should the pursuit of knowledge be constrained by a society's moral concerns? Or should knowledge be left to grow without restrictions or limits?

The German sociologist Max Weber (1864-1920) famously declared that science is itself "value-free," meaning that science never tells us what *should* happen but only what *happens*. A scientist can detail the process of nuclear fission and explain

what happens, but it is not the scientist's job to tell anyone whether they should build nuclear reactors or nuclear weapons. The scientists just tries to determine what is true; it is up to the rest of society, or its leaders, to decide what to do with that knowledge.

We might explore Weber's claim a bit further by considering the career of the German rocket scientist, Wernher von Braun. Von Braun engineered V-2 rockets, which killed thousands of British civilians, for the Nazis. He was a member of the Nazi party, wore an SS uniform, and was certainly aware that the rockets were being built by slaves in German concentration camps. Yet at the end of World War II, all was apparently forgiven as the US was keen to have him among their scientists. One might try to exonerate von Braun as a scientist just doing his job. As the songwriter, comedian, and social critic, Tom Lehrer, once sang sarcastically:

*Don't say that he's hypocritical!
Say, rather, that he's apolitical.
"Once the rockets are up, who cares where they come down?
That's not my department!" says Wernher von Braun.*

But this works only as sarcasm. Von Braun was not simply doing science under a regime that only happened to be the

Nazis. Being a Nazi was interwoven with the research he was doing. Indeed, his job as a scientist was to make deadly rockets to aid a monstrous political engine. This helps to demonstrate the more general point that science does not develop in some sealed environment that is insulated from society at large. Scientists are raised in societies and their attitudes are deeply shaped by those societies. Scientists do not “leave themselves behind” when they walk into the lab but carry with them their own attitudes, beliefs, conceptions, and prejudices.

We know from history that this can lead to very biased and inaccurate science. This means that ways in which a society’s morality can distort our beliefs should not merely concern us for *moral* reasons but also for *epistemological* ones. If we are interested in learning what is true, we should be concerned about the ways in which our society’s moral values distort our knowledge. In other words, if Wernher von Braun wanted an undistorted understanding of rocketry, he would have done well to pay some attention to how the rockets were being built and where they were meant to come down.

Objection: But that’s clearly **not true** in the case of Wernher von Braun. He was a true expert in rocket science, and while greater moral concern makes anyone a better human being, it is not at all obvious that greater

moral concern would have made him or anyone more knowledgeable about rocket science.

This is a good point. Still, one might ask whether the push among all technological nations for more advanced ballistic weapons was informed by moral concerns or by concerns to intimidate other nations and channel public funds into defense industries. Suppose, as a thought experiment, that the push had been to design missiles that could transport food and medical supplies to distant regions, with a possibility of re-using the rockets. If that had been the objective, would rocketry have developed even further than it did when the objective was only destructive?

This consideration is not mere fantasy. In fact, once Wernher von Braun became an American citizen, he became a very strong proponent of using rockets to travel safely into space. He continued to work on projects for the US military but with greater moral reservations and far less enthusiasm. One might argue that his later work, particularly with the Apollo program, generated far superior knowledge of rocketry *because* the objective was not restricted to better ways to blow up distant targets but to the more difficult task of sending humans into space without killing them. That's a lot harder to do and requires greater knowledge.

The point of reviewing von Braun's story is to suggest that the moral beliefs of a surrounding society determine how individuals see their world, what they value in it, which projects are possible or encouraged, and which are not. Though, clearly, it is not impossible to learn more about the world even in a thoroughly immoral society, a society that is more open to dialogue and to changing its moral attitudes will allow for a greater range of efforts to gain knowledge. A free society allows for free knowledge.

THE OPEN SOCIETY

Karl Popper (1902-1994) was a philosopher who wrote on both scientific knowledge and politics. His most famous work was a two-volume book entitled *The Open Society and Its Enemies* (1945). Popper's claim was that, all through history, societies that are run by elites who take themselves to have more valuable knowledge than other people always end up being repressive tyrannies. Plato's beloved republic was supposed to be run by philosopher kings, and in this society, Popper notes, there is state censorship, noble lies, eugenics, and very little human freedom. Hegel's society, Popper argued, requires individuals to subordinate their own interests and beliefs to the plans of an Absolute Spirit that worked through whomever happened to be king; and this unquestioning obedience, combined with the faith that God is behind whatever the state is doing, helped to make the rise of National Socialism possible in Germany. Marx believed that in the transition from capitalism to communism, there would need to be an interval when authorities controlled everything and re-educated the masses; this idea led to Stalinism. In general,

Popper argued, whenever people set themselves up as *knowing better than others*, an oppressive tyranny results.

Popper's alternative is the *open society*. An open society is one in which people are free to think as they like and say what they think. People are free to criticize one another's claims and ask for evidence and for justification. No one inherently has a greater claim to the truth. Each individual has the right to employ their own reason to determine their beliefs. A democracy is the best form of government for such a society as it allows for free and equal participation by all citizens. Obviously, in such a democracy not everyone will get their way, but everyone will have the chance to offer their own view, and a view will become dominant only by winning over the majority of citizens. So long as citizens are encouraged to exercise their own critical rationality, asking for evidence and justification and deciding on views that seem best supported by available knowledge, the society will generally follow the best available suggestions. Mistakes will be made, but no worse than what happens in any alternative to an open society.

This is obviously only the beginning of a rich discussion in political philosophy, but what is important for our purposes is the way Popper links **epistemic autonomy** with **social and**

political freedom. Epistemic autonomy is an individual's capacity to use their own reason and experience in determining what is true. The opposite would be being told what to believe without sufficient evidence or reason—belief at the point of a gun, in other words. Epistemic autonomy both *leads to* and *results from* social and political freedom. It *leads to* social and political freedom in two ways. First, individuals with epistemic autonomy want to have for themselves the freedom to use their reason and determine their own beliefs, and preserving this freedom for themselves will mean also preserving it for others (at least, so long as they do not regard themselves as having special privileges). Second, individuals may discover that the best way to use reason and evidence in determining their beliefs is through *free and open dialogue with others*, since others will have reasons and evidence that the individuals had not considered.

But it is also true that epistemic autonomy *results from* social and political freedom. If we allow citizens a maximal set of rights—in the words of John Rawls, a set of rights that is as broad as can be while extending the same rights to everyone else—then individuals will need to determine for themselves how they act, how they live, and what they believe. This is epistemic autonomy.

Objection: This is all well and good so far as it goes.

But even in self-professed free societies, there can be propaganda and persuasive advertising and all sorts of ways to manipulate people's beliefs. People may think they are epistemically autonomous when, in reality, they are being manipulated by powerful political or corporate interest groups.

Popper would hasten to agree. But how do we fight against such manipulation? Appointing some small group of people either to be in charge or to determine what public knowledge should be would only make the situation easier for those who seek to manipulate society to their own ends, for the manipulators now need not try to convince a majority of citizens but only the small group in charge. The best antidote to *bad information* is *more information*, Popper would say. Over time, the information that is more accurate will prevail over misinformation.

But let us consider a harder case in which citizen's freedoms might be seen as threatened by having more information. For many years in the U. S., standardized test scores have favored some groups of people over others. Males tend to outperform females, and Asian students outperform other racial or ethnic groups. This is normally regarded as indicating both biases in the tests and differences in the educational experiences of

people in these groups. But a few researchers have claimed that the differences in the test scores remain even when one compensates for the tests' biases and the differences in experience. In other words, some portion of the difference in test scores, they claim, really just has to do with the differences in sex or race or ethnicity. Some groups of people are smarter than others, at least according to the measurements of these tests.

Now suppose the view of this minority of researchers were to turn out to be true. Suppose there were some measurable difference in intelligence between these groups of people. Suppose we could set aside all of the well-founded concerns about standardized tests, disparities in education, and so on, and suppose there really turned out to be such a difference—a small difference, perhaps, but a genuine difference. It is a difficult conceptual possibility for us to confront because we know how such a result would be deployed in the service of sexism and racism. Elite college admissions and high-level employment opportunities would be skewed towards privileged groups—for after all, don't we want the best and brightest in these spots? And people would be shut out from opportunities over factors over which they had no control such as their sex or race. We might well worry that all of the hard work that has gone into the struggle for equality of

opportunity among historically disadvantaged groups would be wiped away with such a research discovery.

With that in mind, would it be wiser for researchers to suppress their discovery? If they saw that the discovery would be used to justify sexist or racist policies, would they not be morally obliged to block that discovery from becoming known?

Defenders of the free society would say that the discovery should not be suppressed. But, of course, that does not mean they would welcome sexist or racist policies. Rather, they might say, the research has given us a strange and wholly unexpected fact that will probably require further study to fully comprehend—but it has not given us any obvious reason for dismantling any of the moral progress we have made. It has been known for a long time that differences exist among groups of people for whatever reason; that is obvious. But we have learned to disregard those differences when it is a matter of social or political equality. Why should it be any different in this case?

The more general strategy of the defenders of the open society

is to let all information be out in the open so that it cannot work behind the scenes in the dark and undiscovered. If some group of people thinks a small difference in standardized tests would justify racist or sexist policies, let them bring that argument out into the open where it can be discussed, challenged, and refuted. The alternative is to allow the argument to fester unspoken in individuals' minds, governing their actions without ever being brought out in public display. The bright spotlight of public scrutiny will kill off the unreasoned beliefs and nourish the ones supported by reason and evidence.

Or that is the faith of the defenders of the open society, at any rate. But having seen what we have seen about the rationality of humans at various points in our history (think particularly of the Victorian scientists here), we may well worry whether “the bright spotlight of public scrutiny” always does the work it is supposed to do.

31.

QUESTIONS TO CONSIDER

1. If our scientific knowledge is shaped by social values and prejudices, then what about our moral knowledge? Is it better off in some regard? How confident can we be about the moral judgments we make about Victorians, for example?
2. Popper was very confident that the open society is tolerant of a wide range of views but was absolutely not tolerant of *intolerance*. In other words, anyone who wants to shut down others in expressing their views is not welcome in the open society. Is this a bug in his system or a feature?
3. Suppose I want to be a scientist but don't want to complicate my life with all of the moral concerns brought up in this chapter. I don't want to participate in evil; I just want to understand nature. Is there some kind of strategy I can follow to make sure I can do my work without worrying about its social implications?

FURTHER READING

The most influential author who has written about the ways in which knowledge is shaped by social conditions is undoubtedly Michel Foucault. But it is difficult to find accounts of Foucault's thinking that are easily approached by beginners. One might begin with the entry in the *Internet Encyclopedia of Philosophy* and follow up with Gary Gutting's *Michel Foucault's Archaeology of Scientific Reason* (Cambridge University Press, 1989).

Karl Popper, *The Open Society and its Enemies*, was published in two volumes in 1945. Discussions and summaries of it can be found at several sites on the internet. It includes criticisms of Popper's main ideas.

The views of the Victorians and the ways in which their own society shaped their science, and how the science shaped their

society, can be found in George W. Stocking's *Victorian Anthropology* (Free Press, 1987).

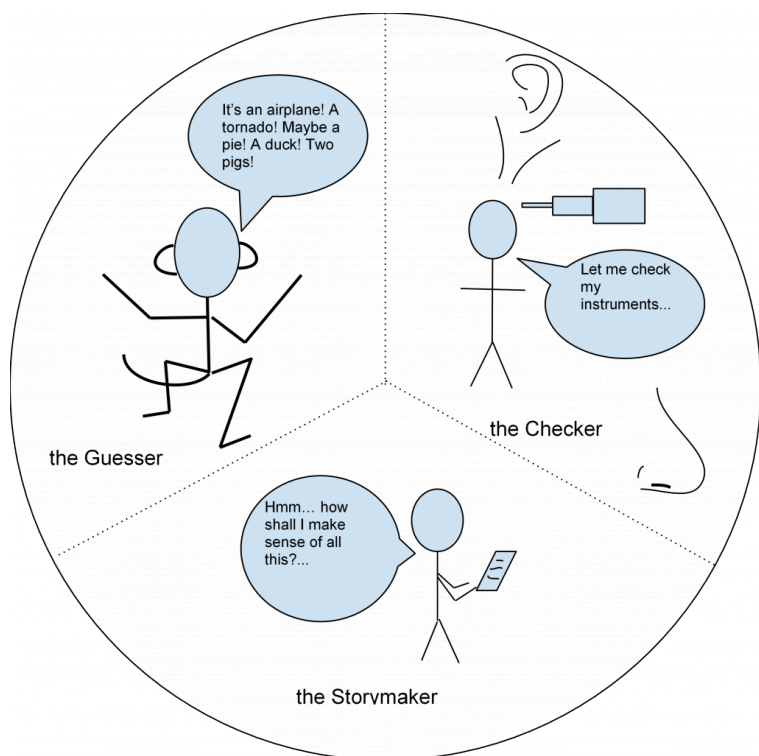
A fascinating account of the ways in which social prejudices affected the career of an amateur botanist and suffragette can be found in Tina Gianquitto's "Botanical Smuts and Hermaphrodites: Lydia Becker, Darwin's Botany, and Education Reform," *Isis* 104 (2): 250-277 (2013).

PART VIII

8. KNOWING OUR WEAKNESSES

Psychologists, philosophers, and cognitive scientists work together to create models of the human mind in order to try to understand how we process information. If we want to understand knowledge, it makes sense to have some understanding of the thinking system we are working with, even a schematic understanding, since then we can know where our system is strong and where it is weak.

To that end, I would like to offer an exceedingly simplistic model of the human mind so that we can begin to think about how our knowledge-gathering or belief-making process works and the various ways in which it can go wrong. The model I shall offer could be called the “Guesser-Checker-Storymaker” model.



The basic idea in this model is that there are three departments in our mind, and each of them has a different job to do. The job of the **Guesser** is to make all sorts of guesses about what is in our environment and even about what we ourselves are doing. It's a wild and creative department, always brainstorming new ideas that come seemingly out of nowhere. The job of the **Checker** is to use our senses to try to determine if any of the guesses coming from the Guesser have any connection to what we can see, hear, smell, taste, or feel. The main job of the Checker is to filter out the wilder guesses

coming from the Guesser and obtain a smaller set of guesses that seem possible given what we are experiencing. Finally, the job of the *Storymaker* is to take the plausible guesses and the information from our senses, and our memories as well, and then try to fold them all together into a coherent story about what is happening and what it all means. It provides the “finished copy” of what we think we know.

The three departments share information and affect one another. So, the Guesser might ask if we are seeing a duck on the pond. The Checker swivels the eyes toward the pond and examines more closely. On the basis of what the Checker sees, the Guesser makes more specific guesses such as whether we are seeing a Mallard or a Pintail—or maybe a goose or just a clump of sticks. Meanwhile, the Storymaker is rapidly putting together an account: “I see an object that might be a duck, and I have seen ducks here in the past, so it is not unlikely, but it’s not perfectly clear...” and this story, as it is being made, further affects the guesses that are being made and the ways in which the senses are being used. “Is this the same duck I saw here last Tuesday—the one with the funny feather sticking out of its head?...” Over time the three of them settle on a story—“Behold! I see a duck!”—before moving on to new guesses and new jobs to do.

Most of the time the system serves us quite well. Most of the time, we get things right. But there are also many ways in which the system can malfunction, leading to surprising results. Imagine what happens when the Checker does not have good access to the sense instruments—perhaps because the senses are “offline” (as they are when we are sleeping) or because their functioning has been affected by poor conditions, prismatic glasses, or hallucinogenic drugs. The Guesser keeps guessing away with all sorts of wild guesses about what’s going on, and the Checker does its best to confirm or disconfirm the guesses. But its functioning is impaired, so it is not very accurate. The Storymaker tries to keep up with the Guesser and the Checker, trying to weave together a coherent story from the information being provided. The result is a dream, or an LSD trip: objects keep turning into different objects (as the Checker keeps confirming wild guesses), and the story seems to make *some* sense at the time, but later on (when we are relaying the story to others), it will seem very strange and incoherent as the plot seems to keep changing. For example, at first I am looking at a duck on a pond, but then the duck is actually my brother, and we need to get to the airport because we are late for a flight. But the airport has no doors, and I am burdened with an enormous orange suitcase...

Or imagine a very intense and dangerous experience like

getting into a fight or getting mugged. All three systems are working quickly and furiously, fueled by adrenaline, looking for immediate threats and escape routes and frantically coming up with the best idea of what to do next. In the heat of the moment, the Checker may not take the time to notice what the other people are wearing or whether they have a beard or whether they are tall or short, the Storymaker may not be keeping record of the precise order of events as they unfold, and the Guesser may be screaming out all kinds of wild ideas in the hope of producing something that will help. Later on, when we tell our friends or the police about what happened, it may be hard to remember exactly how things went down, what the other person looked like, what they were wearing, or why we said the things we said. Eyewitness testimony is notoriously unreliable for this reason: the memories are made under extremely adverse circumstances.

Or imagine a Storyteller that does not pay attention to the information that is being provided by the other two departments. The Storyteller is totally occupied in putting together a story, perhaps one of events from the past or imagined events of the future, and it is putting together all sorts of details and consequences and flourishes. The Guesser is guessing away, as usual, and the Checker is checking away, as usual, but none of the information is being taken up by the Storyteller since the Storyteller is totally absorbed in its

own project. Then, a message suddenly comes through from the Checker: it appears we have been asked a question, and everyone around us is looking at us expectantly. What has been going on? We have no idea since the Storyteller has not been paying attention. We have been daydreaming. Things happening before our very eyes and ears have left no impression on us at all, though nothing was wrong with our Checker, and perhaps we were even nodding along with what other people were saying, though not keeping track of what was being said.

Generally, our cognitive system works very well. After all, it has served us well enough to allow us to survive this long in our evolutionary story. But its complicated nature means that it can malfunction from time to time. Moreover, it may be “engineered” to perform some tasks very well and others not very well at all. For example, perhaps our cognitive system works really well at processing information and maintaining social life in a small group of hunters and gatherers living on a savanna, but perhaps it does not do so well when asked to memorize passages and recite them backwards or to estimate probabilities.

Even apart from the various sorts of malfunctions described

above, it seems our thinking system does tend to make recurring sorts of mistakes in ordinary circumstances. We can call these mistakes **fallacies**, or ways of processing information that do not reliably yield true beliefs. We will consider seven such fallacies and also use our simplistic model of the mind to diagnose how these patterns of mistakes come to be made. Then, we will see if we can draw an interesting general conclusion from these fallacies and some guiding advice for steering clear of them.

Media Attributions

- Figure 8.1 © Charlie Huenemann is licensed under a [CC BY-SA \(Attribution ShareAlike\)](#) license

SEVEN FALLACIES OF HIGHLY-HUMAN THINKERS

Many sources have identified a host of fallacies we are prone to commit, but much of the discussion stems from a highly influential book, *Thinking, Fast and Slow*, by Daniel Kahneman. Kahneman provides his own list of fallacies we are prone to use. The following list overlaps a bit with his but also offers some other fallacies that seem to me very common and more relevant to this introduction to epistemology.


1. Anchoring. The thing we learn first often has an overly strong effect over the rest of what we learn, and we make decisions about what to accept or reject on the basis of that first thing. For example, perhaps we learn from Grandpa that the candy store charges too much, and thereafter, we insist that this is true. Every bit of contrary evidence that comes in (“But the chocolate bars are, in fact, cheaper than anywhere else!”) is brushed aside as just a fluke, or a ploy by the candy store to lure in more customers and then overcharge them for other

items. But what reason or evidence do we have for believing that the candy store is so devious? Really, it is just that we *first* learned one thing from grandpa, and then we stuck with it. Our Storymaker makes the first thing a crucial element in the story and uses it as a criterion for deciding what else to add to the stories we make. But of course, it can easily be that the first thing we hear or learn is not reliable—even if it comes from Grandpa!—and is not entitled to this kind of authority.

2. Confirmation bias. We seem to be wired to look greedily for evidence that supports whatever we already believe and to ignore any evidence that suggests otherwise. But this is, if anything, the *opposite* of what we would want to do if we wanted to be sure that what we believe is true. If we were to try to follow something like the scientific method in forming our own beliefs, we should look instead for good evidence that what we believe is *false*. And if we find none, then we may *tentatively* hold on to our belief until further evidence comes in. This bias toward trying to confirm our beliefs seems to result from the way our Checker interacts with our Guesser: the Checker looks to see if a guess is *correct* and does not look for evidence *against* the guess.

One illustration of our inclination toward confirmation bias is the “Wason selection task.” We are shown the cards in the diagram below. Each card has a color on one side and a number

on the other side. We are asked to determine whether the following rule is true: *every card with an even number on one side is red on the other side*. To determine whether this rule is true, we are allowed to turn over only two cards. Which two cards should we turn over?




Most people think immediately of turning over the “8” card and the red card, probably because the rule we are thinking of combines even numbers and red cards. So, we want to see whether that connection holds. But in fact, the only way to test the rule is to turn over the “8” card and the *brown* card. The red card will not tell us anything because the rule does not tell us whether cards with odd numbers might also have red on the other side. The rule also doesn’t tell us what happens with odd numbers. The rule only says that *if* we have an even card, *then* there’s red on the other side. So, to test the rule, we had better make sure the brown card *does not* have an even number on its other side.

It is hard for us to think this way because we are prone to think in terms of *confirming* a claim rather than *disproving* a claim or looking for contrary evidence. We look for the “even + red” combination to be true and do not think through what we would have to see in order for it to be proven false.

3. Dunning-Kruger effect. We tend to be more confident about our own expertise the *less* we know about something. As the philosopher Bertrand Russell said, “The whole problem with the world is that fools and fanatics are always so certain of themselves, but wiser people so full of doubts.” As you might imagine just from its name, the Dunning-Kruger effect has been studied in research settings by psychologists and has been put in the form of a graph that shows the relation between how much people know about something and how confident they are about their knowledge:

Dunning–Kruger Effect



So, people who know next to nothing about a subject are very confident in their beliefs about it. Then, with a little more knowledge, people realize that they know very little, and gradually, as they learn more and more, they become more “sustainably” confident in their knowledge.

The Dunning-Kruger effect is due to both the Guesser or the Storymaker in our model. If a topic comes up about which I know very little, my Guesser will go to work making guesses about the topic. My Checker has nothing to contribute since the topic is not about my immediate surroundings. My

Storymaker does not have much to offer since (again) this is a topic about which I know very little. So, the guesses I make will get a “free pass,” particularly if they happen to cohere nicely with somewhat-related beliefs I already have. No resistance is offered by any component of my cognitive system, and so I feel very confident of my guesses, like I am an expert. But in this case, my so-called knowledge really consists only in my *not knowing any better*.

4. In-group bias. We give greater weight to the experiences and reports of those who belong to our groups. The people in our groups are friends, family members, or co-workers whom we know and trust; it is hard for us not to trust and believe them. So, for example, I might read study after study that shows that vaccinations prevent disease, but the fact that my mother’s second cousin became extremely sick after receiving a vaccination when she was a little girl outweighs all of the evidence of the studies, and my entire family is set against any vaccinations as a result. Or, for another example, if I see on social media that all of my friends seem to share a political view, it will be difficult for me not to want to share that view with them. I trust them; they are like me; how can I disagree with them?

But of course *anyone* can be wrong, and some of these people may be our friends and family. There are excellent reasons for

trusting family and friends, but such strong trust becomes a liability in cases where what our group says is at odds with what stronger evidence suggests. In terms of our model, in-group bias seems closely related to **Anchoring**. The knowledge of what my group believes does not have to come to me first, but I give it a stronger voice or greater authority than other beliefs or considerations that come my way because it is coming from *my group*. My Storymaker regards it as a “vital element” to the story because my group, and belonging to my group, is a vital element of *my* story.

5. Out-group anti-bias. This comes along with in-group bias but is important enough to merit special attention. Just as we are likely to place *too much* trust in those who belong to our group, we are likely to place *not enough trust* in those outside our group. This is clearly demonstrated by the level of hostility on the internet toward people who are not in our groups. Anything that supports an outside group is seen as a threat to our group. The reasoning and evidence that supports the views of an outside group is rarely considered impartially and honestly, just as support for the views of our group is seldom subjected to critical assessment. Both in-group bias and out-group bias are products of an “us vs. them” mentality which skews our reasoning and ultimately puts us all in a weaker position with regard to knowledge. Forming epistemic groups makes us all ignorant.

6. Availability heuristic. It is difficult for us to continuously process all the information that comes our way. One shortcut for processing it is to make a “snap” judgment that what we are experiencing fits some sort of *model* or *template* (or *heuristic*) that we already have available. An obvious example of this is employing a *stereotype*, or making use of a ready-made list of characteristics in order to make judgments about an individual on the basis of their race, sex, or ethnicity. Unfortunately, our culture provides a very handy set of heuristics to use in judging people which allow us to draw false conclusions rapidly and easily.

Generally, in the “availability heuristic” fallacy, we adopt a model for understanding a problem or question in the hope of securing a fast and easy solution, but that model might not be the best one to use. We might insist that a certain stock value has to fall because “what goes up must come down.” In this case, we have assumed that stock values follow the same laws as projectiles. Or we might be afraid to swim in the ocean because we just watched a string of films about shark attacks. In this case, we have assumed that the films we watched provide a good model of what typically happens when people swim in the ocean. Or we might disparage a scientist who speaks with a certain accent because people with that accent are depicted in movies as being uneducated. In this case, we are assuming that movies provide accurate models of some

implausible connection between accents and intelligence.

In our model, this is the Storymaker's fault. Rather than take the time and effort to compose an accurate story from the available information, the Storymaker slaps on some handy story that is available and moves on to the next task.

7. Barnum effect. This effect is named after the great American huckster P. T. Barnum. Barnum realized that, in trying to deceive someone, you can count on the other person to meet you halfway. In some cases, we join in the effort to deceive ourselves—perhaps because we are being sold a flattering story or because the misinformation being presented to us fits so neatly with preconceived opinions we hold or allows us to draw conclusions we are already eager to draw. On some topics, we really don't mind being fooled.

This is seen most clearly in the business of telling fortunes or writing horoscopes. The fortune teller only needs to provide a vague outline, and most people will fill in the details for themselves. I can demonstrate this to you by showing off my own psychic powers: *I know **you**, the person reading this book right now. You feel a strong need for other people to like and admire you. You have a great deal of unused capacity which*

you have not yet turned to your advantage. You pride yourself as an independent thinker and do not accept others' statements without satisfactory proof. Yet, at times, you have serious doubts as to whether you have made the right decision or done the right thing.

Are my psychic powers not astounding?! But of course those last four sentences apply to anyone and everyone. If you were thinking about them as you read them, you probably thought of features in your life that fit the description. That is natural since, as we try to understand anything, we think about how the new information fits with what we already know. If the “new” information is about us, and if it is suitably general and vague, we will have no difficulty in thinking of ways in which the information fits with our knowledge of ourselves—particularly when the information sounds flattering—and we are deceived into believing that someone else has uncanny knowledge of our own private lives.

But the fallacy does not happen only in the presence of fortune tellers. Our Guesser and Storyteller love to fill in the blanks in any explanation. If someone provides us with a partial explanation, we will automatically start to fill in the gaps with whatever guesses seem to us to be plausible or to fit our other

beliefs or suspicions. But this means that significant parts of the story we end up with have been invented by us and may have no real connection to the truth.

Media Attributions

- [Figure 8.2 Huenemann](#) © [Life of Riley](#) is licensed under a [CC BY-SA \(Attribution ShareAlike\)](#) license
- [Figure 8.3](#) is licensed under a [Public Domain](#) license

REFLECTION ON THE FALLACIES

Obviously, we fall prey to more than just those seven fallacies. But seven turns out to be the most popular favorite number, and clearly, it is *not* any sort of fallacy (such as “wishful thinking”) to believe that our favorite number must be epistemologically significant, right?

These fallacies might be understood as consequences of the lifestyle we evolved to have. For most of our existence, *homo sapiens* have lived in small hunting and gathering groups. We forge strong bonds with other members of our group and show a certain amount of suspicion, or outright hostility, toward strangers. Life has been precarious with bad weather, food shortages, predators, and disease, so naturally we try to establish a lifestyle that is as predictable, as familiar, and as free of surprises as possible. It has been a matter of life or death for us. And so, the “evolved advice” has been to stick to what is familiar, what has worked in the past, what the rest of

our group believes, and whatever encourages us to stay within the group. And on the other hand, we reject what is strange, foreign, new, or comes from groups we don't know. Most or all of our fallacies can be tied to this mindset.

There is no arguing with success. This evolved advice has worked well for most of our existence. But of course we live in a radically different world now, one that we have not evolved to live in. Our groups are much bigger and spread across the globe; food comes to us from a complicated and intricate supply chain, and medicine is, for many of us, readily available; and our knowledge of the world and of ourselves is enormously more advanced, and many of the problems we face require difficult, scientific and technological thinking. Our survival in today's circumstances requires that we *out-think* the ways of thinking burned into us over evolutionary time, which is a hugely difficult challenge. We *know better*. Now we have to convince ourselves of that fact.

But how are we supposed to do this? How can we slow down or alter a style of thinking we have evolved to use? No simple remedy is available, but I will offer here four questions we might try to habitually ask ourselves to try to diminish the power that these fallacies exert over our thinking.

1. **What is my line of reasoning?** Sometimes making our thinking explicit is all it takes to make us realize it is fallacious. “Why am I going back to that same restaurant after just having had three bad experiences eating there?” Answer: “Because it would be really convenient for me to be able to go to that restaurant, and whatever would be convenient for me is likely to be true ... *oh, right, bad idea!*”
2. **Where did I get that idea?** This question is useful only if you can be very honest with yourself. Was it from Grandpa? Did you get this idea from something you read? Or someone you had a conversation with? On social media? Or some weird philosophy professor? Is it impossible for that source to be wrong, or for you to have misunderstood what the source said? If you do a little research, can you find better evidence for thinking the idea is wrong or, at least, not obviously right?
3. **How would things look if I were wrong?** This is an extremely useful question to ask, and we will discuss it further when we come to Bayesian reasoning. Ask yourself what evidence you would be seeing if what you are insisting on is *false*. Then, ask yourself if that is, in fact, what you are seeing. So, for example, suppose you

have ended up believing that the Moon landing never happened and all the reports and records of it are fake. Now try to imagine what you would be seeing if you were wrong about this, and people actually did land on the Moon. You would be seeing everything you in fact are seeing—video footage, recordings, books, movies, accounts in textbooks, and so on. There may also be some minor inconsistencies in the historical accounts or some fuzzy photographs since these commonly happen in all cases. Would you still also be seeing a small group of people steadfastly refusing to believe that people landed on the Moon? Yes, probably. You can pretty much always count on their being some small group of people denying what's true no matter the topic. This small exercise in imagination should be enough to lead you to think that there's a good chance you are wrong.

4. **How does this belief make me *feel*?** This might seem like an odd question to ask, but if you try to answer it honestly, it *might* reveal what motivates your belief. If the belief you are considering does not really make you feel one way or another, that is probably *good news*. For when our emotions are not exerting their powers over us, we are more likely to come to an unprejudiced assessment of the evidence. On the other hand, if the belief makes you feel important or sort of thrilled or impassioned or like you are an extremely special person,

then *watch out!* For it may be your desire to feel those emotions that is pulling you toward that belief more than any evidence or line of reasoning. Now, obviously, we can be wrong about emotionally-neutral information, and we can be right about ideas that excite us. But when strong feelings are attached to what we believe, we would be wise to slow down and consider more carefully what our line of thinking really is.

Asking ourselves these four questions will not counteract every line of fallacious thinking, and we may not always be able to ask ourselves these questions or answer them honestly. But they may help us to steer clear of the seven fallacies we discussed, in the following way:

	What is my line of reasoning?	Where did I get that idea?	How would things look if I were wrong?	How does this belief make me feel?
Anchoring	x	x	x	
Confirmation bias	x		x	x
Dunning-Kruger effect	x	x		
In-group bias		x	x	x
Out-group anti-bias		x	x	x
Availability heuristic	x	x	x	
Barnum effect	x	x	x	x

QUESTIONS TO CONSIDER

1. Do you have any examples from your own experience of any of the errors described in this chapter? If you are human, you probably do! Please offer that example, explaining clearly why it was an instance of some particular fallacy.
2. A very common fallacy that is not listed in this chapter is the fallacy of appealing to authority. It seems that the fact that Mr. Popular endorses Smile-eeze toothpaste should not rationally persuade me that Smile-eeze toothpaste is good for my teeth. But, as we all know, we have to rely on experts to give us information about the weather, medicine, load-bearing structures, and so on. How would you characterize the fallacy of “appeal to authority” to allow the good appeals to authority and disallow fallacious appeals to authority? After you offer that characterization, develop a solid objection or counterexample to it.

FURTHER READING

The Internet Encyclopedia of Philosophy's entry "Fallacies" offers a good overview of questions that can be raised about providing lists of named fallacies. It also presents a list of 230 named fallacies with brief descriptions and examples.

Daniel Kahneman's *Thinking Fast and Slow* (Farrar, Straus and Giroux, 2011) offers an illuminating discussion of two ways we think: in a faster mode with quick conclusions, and in a slower mode in which we carefully articulate each step along the way. We seem to be better at fast thinking, but fast thinking is unreliable in any complicated situation.

Hugo Mercier and Dan Sperber's *The Enigma of Reason* (Harvard UP, 2017) argues that the biases humans are prone to in reasoning are corrected when humans argue together in groups, and that reasoning is a social "superpower" rather than one belonging to individuals.

PART IX

9. HOW TO ARGUE WITH OTHER PEOPLE

“Man: An argument isn’t just contradiction.

Mr. Vibrating: It can be.

Man: No it can’t. An argument is a connected series of
statements intended to establish a proposition.

Mr. Vibrating: No it isn’t.

Man: Yes it is! It’s not just contradiction.

Mr. Vibrating: Look, if I argue with you, I must take up a
contrary position.

Man: Yes, but that’s not just saying ‘No it isn’t.’

Mr. Vibrating: Yes it is!

Man: No it isn’t!”

— Monty Python, “Argument Clinic” (1972)

Unfortunately, we do not lack examples of people arguing with one another. What we lack are examples of people arguing *constructively* with one another. But how can arguing be constructive? Isn’t any argument essentially a competition in which each side is trying hard to win and not to lose?

We should begin by introducing the philosophical sense of the term “argument.” According to the philosophers, an argument is not a competition between people who believe different things. An argument is instead a set of reasons (called *premises*) that are supposed to lead to a conclusion. If they really do lead to the conclusion, the argument is said to be *valid*, it *works*. Anyone who disagrees with the conclusion will then have to find fault with one of the premises leading to the conclusion.

Here is an example of an argument:

Anyone who is either a philosopher or a scientist believes in truths that go beyond what experience can show. For a philosopher believes that there are proper definitions of terms, like *truth* or *justice*, and a scientist believes in the laws of nature, which are generalizations that go beyond the particulars of sense experience.

Is it a valid argument? To determine this, we need to determine what claim is being argued for (the conclusion), and what reasons are being presented for it (the premises). In this case, the conclusion is the first sentence: “Anyone who is either a philosopher or a scientist believes in truths that go beyond what experience can show.” Then, the second sentence

gives us two reasons for thinking that the conclusion is true: “For a philosopher believes that there are proper definitions of terms, like *truth* or *justice*” (that’s reason or premise #1), and “a scientist believes in the laws of nature, which are generalizations that go beyond the particulars of sense experience” (that’s reason or premise #2).

This argument is valid, and the following, painstaking re-ordering of the argument will make it clear exactly why it is valid:

(Premise #1): A philosopher believes that there are proper definitions of terms, like *truth* or *justice* (and we should add here something that is pretty obviously being presumed, namely that what makes definitions “proper” is something that goes beyond what experience is able to show).

(Premise #2): A scientist believes in the laws of nature, which are generalizations that go beyond the particulars of sense experience (so they quite literally go beyond what experience is able to show).

(Conclusion): Therefore, both philosophers and scientists believe in truths that go beyond what experience is able to show.

If we drop out all of the information and pay attention only to the structure of what is being said, the argument looks like

this: All of the Ss are Bs, and all of the Ps are Bs, so all of the Ss and Ps are Bs. It is not a thrilling argument, but it is a valid argument. The two premises that are offered do in fact lead to, (or *imply*) the conclusion.

So, there is no arguing with the logic of the argument. If someone wants to resist the conclusion—let’s say they are a scientist who hates being put into the same group as philosophers—then they will have to reject one or both of the premises being offered. Perhaps they will deny that scientists, when they believe in the laws of nature, are believing in truths that go beyond what experience shows. If that is the case, then we will return to the discussion in chapter 6 of this book, raise the problem of induction, and our discussion continues. Or perhaps someone thinks philosophers are not concerned with proper definitions but are instead concerned with other things that do not go beyond what experience shows. In that case, we shall have to ask what this person thinks philosophy is, and our discussion continues.

Either way—and this is the point—*our discussion continues*. The understanding of the initial argument has helped us to focus more precisely on where we disagree. Is it about what scientists believe or about what philosophers do? So, the

discussion not only continues but continues in a *constructive* way since we are learning more about the different views and the reasons for holding the different views, or about the exact nature of the disagreement between people. This sort of result seldom comes from the shouting matches that are often recognized more popularly as “arguments.”

A philosophical argument is a cooperative effort to understand the reasons behind our disagreements. Sometimes the result is that a simple misunderstanding is cleared up. Sometimes one side ends up persuading the other because the reasons are made clearer, and on the basis of those reasons, someone is convinced that they should change their mind. Sometimes no one changes their mind, but everyone has a clearer picture of where other people are coming from or what their lines of reasoning are. Every result is an advance from where we were when we started.

There is a science of *logic*, or the exact nature of the ways in which premises lead to conclusions (or *validity*). It is extremely important not only in philosophy, but in mathematics, information science, and, really, any endeavor in which people are trying to extract more specific information from the information that is given. You should take a class in

logic if you haven't already. But we cannot cover all of that material in this textbook, so we will turn to something that is just important: how to argue with people philosophically so that our arguing is constructive and illuminating, not just frustrating and tiresome.

Taking the time to reorganize premises and conclusions in the painstaking way we just did, making everything transparently obvious, is hardly ever practical in daily life (though it is frequently done in philosophical essays). And, as was just said, logic is a science unto itself. But we can examine the attitudes we should be bringing to philosophical arguments and learn from them some lessons about the ethics of knowledge.

37.

ARGUMENTS AS OCCASIONS FOR LEARNING

The great Socrates (c. 400 BCE) argued that philosophical arguments never disappoint. His line of reasoning might be presented as follows: any two people either have knowledge about something, or they do not. If they both have knowledge, then they will agree, and they will not need to argue. If one of them has knowledge and the other does not, then the one with knowledge will teach the other. If neither of them has knowledge (but perhaps they falsely *think* they do), then a philosophical argument will soon demonstrate to both of them that, in fact, they do not have knowledge—which will be an improvement upon falsely thinking that they have knowledge. So, in all cases, philosophical arguments lead to somebody's improvement.

In Socrates's own case, he had been told by an oracle of the gods that he was the wisest of all humans. He could not see

how this could be true, so he went to the reputed experts of his day to try to find someone who was wiser than he was. What he found was that many people were regarded as wise, and regarded themselves as wise, but in reality did not have any wisdom. In the end, he thought that he was the wisest of all humans only in the sense that he knew he had no wisdom while so many other people falsely believed that they had wisdom when they didn't. In that regard, he was wiser than they were.

Socrates's arguments with the alleged experts were of the third sort mentioned above. Neither Socrates nor the alleged expert had knowledge, and the argument showed this to be true. Or it was *supposed* to show this. Very often the alleged experts refused to admit their own ignorance, and they began to regard Socrates as an annoying pest. He was eventually accused of impiety and corrupting the youth, and at his trial, he refused to stop doing what he was doing:

Gentlemen of the jury, I am grateful and I am your friend, but I will obey the god rather than you, and as long as I draw breath and am able, I shall not cease to practice philosophy, to exhort you and, in my usual way, to point out to any one of you whom I happen to meet: "Good Sir, you are an Athenian, a citizen of the greatest city with the greatest reputation for both wisdom and power; are

you not ashamed of your eagerness to possess as much wealth, reputation, and honors as possible, while you do not care for nor give thought to wisdom or truth or the best possible state of your soul?" Then, if one of you disputes this and says he does care, I shall not let him go at once or leave him, but I shall question him, examine him, and test him, and if I do not think he has attained the goodness that he says he has, I shall reproach him because he attaches little importance to the most important things and greater importance to inferior things. (Plato's *Apology*, 29d-30a)¹

It is a beautiful and moving speech. Of course, Socrates was in a life-or-death situation (which unfortunately ended up being death for him). We are not typically in that sort of situation. But there is something in Socrates's general attitude that is worthy of emulation. His attitude is that we should place a very high value on truth and that we should not be afraid of being wrong. We should enter into conversations and arguments about what is true and be willing to admit when we are wrong since then we will have learned something.

This Socratic attitude fits nicely with the "open society" envisioned by Karl Popper (as we saw in chapter 7). In an

1. Plato, "Apology," in *Plato: Five Dialogues*, (Indianapolis: Hackett, 1981).

open society, we are free to think as we please and say what we think. We are also free to criticize or object to what others say. Along with these freedoms, though, comes a responsibility to be reasonable and value the truth and admit to our own mistakes. An open society in which everyone just blabs whatever they wish to blab about will not be much of a society to be proud of. But one in which everyone is exchanging ideas and engaged in reasonable discussion of those ideas—in short, an open society in which everyone adopts the Socratic attitude—would be a truly great society.

38.

BEING FAIR, AND EVEN GENEROUS

Once we shift our attitudes about arguments and view them as occasions for learning and not as competitions, we will find that we do not have to mean or dismissive to our opponents. Indeed, we *should* not be mean to them, and we should not think of them as “opponents,” for they may turn out to be our teachers. We should at least regard them as *friends*, which is what Socrates typically does in Plato’s dialogues.

With that in mind, we will want our friends to present the best arguments they can, even when we disagree with their conclusions. We want this because we want to discover the truth. If we let our friends slide by with less forceful arguments, we may well miss an opportunity for learning something. Ideally, we want our opponents to be as rational, clear, and persuasive as possible so that we do not miss out on what they know. As J. S. Mill (1806-1873) once wrote, “Lord, enlighten thou our enemies. Sharpen their wits, give acuteness

to their perceptions, and consecutiveness and clearness to their reasoning powers: we are in danger from their folly, not from their wisdom; their weakness is what fills us with apprehension, not their strength.” We want our “enemies” to be as wise and clear and perceptive as possible, for they are not really our enemies, properly speaking. Ignorance or foolishness are the true enemies. If we want to learn the truth, we want everyone to be as sharp and perceptive as possible.

The contemporary philosopher Daniel Dennett has suggested three rules to follow when we criticize someone else’s argument:

1. You should attempt to re-express your target’s position so clearly, vividly, and fairly that your target says, “Thanks, I wish I’d thought of putting it that way.”
2. You should list any points of agreement (especially if they are not matters of general or widespread agreement).
3. You should mention anything you have learned from your target.

Only then are you permitted to say so much as a word of rebuttal or criticism.

Note that this does not necessarily mean saying that the other person's conclusion is *right*. Sometimes we are encouraged to interact with others in such a way that no one is ever said to be wrong about anything, and everyone is right, and we should all hold hands and give thanks for each other's company. It's a nice thought, perhaps, and of course we should always be kind and respectful to others. But we can respect others and follow Dennett's suggested rules while still disagreeing, raising forceful objections and criticisms, and insisting that the other person's view is false. It is tricky, to be sure, since no one ever likes being told that they are wrong. (Well, Socrates seems not to have minded.) *But the truth is at stake*. And if the truth matters, then we are not doing anyone any favors by saying they are right when they are not.

The trick is to show respect and even kindness while disagreeing and making clear exactly where our disagreement is. This is what Dennett's rules are about: being fair and generous to one's partner while also being as clear as possible about what is right, what is wrong, and *why*. This is the ideal of a philosophical argument.

It is one thing to understand the ideal and quite another to be good at achieving it. It is hard to be fair to those we disagree with and to refrain from making all of the unfair but devastating and clever remarks we might make at their expense. But whenever possible, we should view each argument as a trial with the very best lawyers appointed for each side. We want the “prosecution” to be aggressive in putting together the evidence and arguments for their side, and we want the “defense” also to assemble the evidence and arguments effectively for their side. Sometimes we have to switch back and forth between serving as prosecutor and as defense attorney as we try to make each side as compelling as possible.

FRIENDLY ADVERSARIES

Many court systems are *adversarial*. What this means is that the court features two “teams,” a prosecution and a defense. Each team tries its utmost to persuade a judge or jury that what they are saying is true. The prosecution tries to prove beyond any reasonable doubt that someone is guilty of having broken the law, and the defense tries to show that there are reasonable doubts about that person’s guilt and that the prosecutors have not proven their case. The article of faith that justifies such a context is that there is a *truth* about the person’s guilt or innocence, and the truth, or the facts, will enable one side or the other to provide a more compelling case to a judge or jury that is rational and impartial.

It is a noble ideal, and there are many ways in which our systems can fail to live up to it. Prosecutors might be extraordinarily talented and compelling orators who can persuade a jury to see things their way even when the truth is not on their side. The same goes for defenders who might be able to encourage great sympathy for the accused person or

suggest grounds for doubt where really there should be none. Evidence can be fabricated or lost on purpose. And the judge or jury may not be impartial but may have their own interest in either convicting someone or letting them go free.

But it is hard to think of a better system. Imagine systems of “justice” that are based on some other contest. We could let the accuser and the accused engage one another in armed combat or in an arm-wrestling match and “let God decide.” We could simply ask the whole community to vote on guilt or innocence without working through the evidence. We could flip a coin. Any of these systems would generate verdicts, but the problem is we have no reason to think that the verdicts that are generated will have any connection to whether the accused person is *really* guilty or innocent. A system in which evidence and reasonable arguments are presented, and we do what we can to make sure rules of evidence are followed, and everything is out in the open, and the judge and jury are as impartial as we can practically guarantee, is a system that should generate verdicts that line up with actual guilt or innocence.

Philosophers tend to see their own disputes as very much like the adversarial system used in courts of law. Arguments are presented, objections are raised, replies to the objections are

offered, the adequacy of those replies are assessed, and the discussion continues. The philosophers involved in the discussion are supposed to serve as judge and jury as well, which of course can mean that they are not completely unbiased as they present arguments and objections and are also supposed to rule on whether the arguments are compelling. But philosophers know they are supposed to be impartial in their rulings, and if they fail to be impartial, we can count on another philosopher to point this out quite forcefully. Daniel Dennett's rules given above can be seen as rules that are meant to keep us on the straight and narrow path of impartiality.

But *who is on trial*? If we follow the example of Socrates, we shall say that *our beliefs* are on trial. A “guilty” verdict means that the belief should go away, and an “innocent” verdict means that it need not go away; the belief is *defensible*. The aim of the philosophical trial is not to put anyone in jail, of course, but to help each other to have beliefs that are defensible. The system of arguments and objections is certainly adversarial, but the adversaries should be friends, as they are trying to help one another.

The same concerns raised above can be raised here: this is a noble ideal which we can fail to live up to in many ways.

Philosophers, like anyone, can become more concerned with scoring debate points than with getting at the truth, and some philosophers can make lousy arguments seem very compelling. But again, as with the court system, it is the best idea anyone has for arriving at defensible beliefs.

Note that both adversarial systems make the task of the prosecutor more difficult than the task of the defender. The prosecutor has to show that a person is guilty or that a certain belief is indefensible. The defender, on the other hand, does not have to show that the person is *innocent* or that the belief in question is *true*. That would be too much. In the case of courts of law, we make the job of the defender easier because, on the whole, we would rather have a guilty person go free (which is bad) than have an innocent person wrongly convicted (which is worse). In the case of a philosophical dispute, we would rather have someone believe something that is false but still defensible (which is bad) than face the challenge of proving something true before believing it (which is worse).

Why would it be worse to insist that we believe only what is proven to be true? Because, as we have seen through our study of skepticism, this will mean never being allowed to believe anything! We were able to escape the clutches of severe

skepticism only by allowing ourselves to have beliefs that are defensible, or at least beliefs that fit with the other things we believe. While it is perhaps a comforting fantasy to imagine having iron-clad arguments for everything we believe, that is only a fantasy, and the bulk of human life as we know it is lived in a wide range between things we are sure are false and things we are sure are true. (Would it really be comforting to have iron-clad arguments? My own suspicion is that such a life would be bereft of stimulating doubts, wonderings, and possibilities!)

40.

HELPING THE DISCUSSION TO CONTINUE

Having spent some decades leading philosophical discussions among undergraduates, I have come to appreciate the tremendous value in *helping the discussion to continue*. It's not simply because a teacher needs to keep students active for the length of the class. Rather, it seems to me that keeping a discussion alive and interesting is an effective way to teach two lessons that are otherwise impossible to teach: (1) that there are other views we had not thought of, and (2) that we can keep talking even though we disagree. These are important lessons for everyone to learn, for it is when we each think we have the only rational perspective and there's no use in talking about it that civil society breaks down.

It is difficult to help discussions to continue, and it takes a masterful teacher to do it well. I am not a masterful teacher in this regard, and I admire those teachers who can accomplish

the task. It requires allowing different views to develop while also keeping the discussion somewhat focused and progressive as it evolves from one question to the next. If it is done well, then a wide array of relevant perspectives are explored, but everyone feels as if the central questions have become clearer and more meaningful. A skillful teacher can work in real time to combine insights from various students and shape the discussion toward (what the teacher suspects will be) a fruitful outcome.

The same discussion-leading attitude can be carried out of the classroom and into all argumentative situations. People involved in arguments, in real life or in comment threads on the internet, are often concerned only to score points and win with bonus points awarded if the opponent is left crying and ashamed. Needless to say, in such a game the rational pursuit of truth is not likely to fare well. The results will only be hurt feelings, resentment, and greater stubbornness. But if the objective is not to score points and win but *to help the discussion to continue*, then a conversation begins to resemble a small-scale version of Popper's open society (as discussed in chapter 7). In this small society of discussion, we encourage the expression of a wide range of ideas and subject each idea to critical questions and problems but now with an aim to develop the conversation further and enlarge our understanding. Again, such discussions are not ones in which

everyone is right and no one ever says anything that is wrong or false, but the discussion allows for corrections to be made or arguments to develop in such a way as (once again!) *to help the discussion to continue*.

The skill of helping discussions to continue is hard to gain, and we often will fail at the task. But, like most skills, we get better at it the more we practice. And if we are committed to being good epistemic agents, not to mention good human beings, this is a skill very much worth developing through continued practice.

41.

QUESTIONS TO CONSIDER

1. Suppose you are arguing with someone who has a belief that is morally repugnant (for example, that people with low IQs are slaves by nature). Should you really follow Dennett's advice and try to "re-express your target's position so clearly, vividly, and fairly that your target says, 'Thanks, I wish I'd thought of putting it that way'"? If so, why? If not, where do you draw the line between following Dennett's advice and not following it?
2. An experiment: find a suitable comment thread on the internet—not one that gains nobody's attention, and not one that is huge, but one in which individuals are likely to reply to one another's comments. Try to engage with the comments constructively, following the rules Dennett offers. Report your results.

FURTHER READING

Plato. “Apology.” In *Plato: Five Dialogues*. Translated by G. M. A. Grube. Indianapolis: Hackett, 1981. Plato, *Apology*. If you have already read it, you should read it again. And again.

Daniel Dennett’s advice about arguing is found in his book *Intuition Pumps and other Tools for Thinking* (W. W. Norton, 2013). The book also contains other interesting and creative ideas.

Also, for a discussion of what makes “trolling” wrong, one should consult Rachel Barnett’s “translation” of a work seemingly by Aristotle, entitled “On Trolling,” published in the *Journal of the American Philosophical Association* 2(2), 193-195. doi:10.1017/apa.2016.9

PART X

10. BAYESIANISM AND WHAT IS LIKELY

Our beliefs are not static. Ideally, our beliefs will change as new information becomes available. Otherwise, we are unreasonably stubborn. But our beliefs should not *always* change since the new information might not matter, or it might not be trustworthy. “A wise man proportions his belief to the evidence,” David Hume once intoned, but it does take some special sort of wisdom to know exactly *how* to proportion our beliefs to the evidence.

In a way, this is the central, practical problem in knowing: how should I revise my beliefs when new information comes in? When should I find new information compelling, and when should I be skeptical? When should I change my mind, and when should I hold fast? Anyone who has a good method for figuring this out would be an expert knower.

In fact, there is a method available for figuring this out—at least, in a wide range of cases—and that is what this chapter is about. The method is **Bayesianism**. But it is a little tricky to understand, so we will have to take some care in building up to it. We will begin by considering David Hume’s discussion of miracles and when it is rational for someone to believe the report of a miracle.

43.

DAVID HUME AND MIRACLES


We have encountered David Hume before when we were discussing the problem of induction. Hume is known as a great skeptic because of this problem and because he did not think there is a rational solution to it. But Hume was far from an irrational philosopher. He believed that once we admit that human beings are conditioned to expect that nature will continue to follow the patterns it has followed in the past, we can begin to think more carefully about what we should or should not believe given the patterns of our past experience. He knew that sometimes we can be surprised when nature doesn't do what it is expected to do. But such surprises are rare (thank goodness!), and we ought to be careful when we hear from someone else that some surprising event has happened.

As Hume recognized, we all make our judgments on the basis of our past experience. He tells the story of “an Indian prince” who was told by some visiting Europeans that where they came

from, water freezes in the winter. The prince had never experienced or even heard of such a thing. In his experience, water was always liquid, and the idea that it could turn into a solid that people could walk on seemed ludicrous. Yet here were strangers in his court claiming that water could become as hard as stone. The prince faced a choice: he could believe these strangers and accept the seemingly outrageous claim that water can turn into a solid, or he could continue to believe what all of his previous experience showed—that water was, is, and always shall be liquid—and suspect these strange newcomers as trying to pull a fast one on him.

What is the rational thing to believe? We know the truth: water can freeze. But when we ask what is rational *for the prince* to believe, we are asking what he should believe given his previous knowledge and experiences. The idea that water turns solid was totally “unprecedented” in his experience (he had no evidence of it having happened before). But the idea that strange people from strange lands might be less than truthful, or might even make up astonishing things in an attempt to impress him, was not at all unprecedented. Maybe for a prince at that time it was even a common occurrence. So, Hume observed, the prince “reasoned justly” when he concluded that water does not freeze and that these visitors were not telling the truth.

We can see in this case a kind of rational *weighing*. Imagine a balance scale. On one side we put the likelihood that water freezes; on the other side we put the likelihood that strangers tell false and fantastic stories. The second likelihood was greater according to the prince's experience and weighed more, so the balance was tipped in favor of denying the strangers' claims.



Hume went on to consider a pair of more complicated cases. Suppose we find historical reports that say that on the first of January in the year 1600 there was darkness over the whole Earth for eight days. Now that is quite a claim, and if it were

just a single report, we would be wise to suspect that the report was false. But suppose it is not just a single report. Suppose that as we gather reports from around the world, from all sorts of societies, we find the same thing being reported in different languages and in different calendar systems. The reports all agree on the details: eight days of darkness beginning on the first of January 1600 (as Europeans reckon it). In this case, Hume thinks, we would be rational to accept the truth of the reports, and we should then start trying to figure out what strange sort of eclipse or weather phenomenon might explain the eight-day darkness.

Why should we accept the reports? Consider again the weighing analogy. On the one side we put the claim that the whole world was covered in darkness for eight days. That is quite extraordinary. On the other side we put the claim that independent societies around the world all came to report the eight-day darkness, *even though, in fact, the world was not covered in darkness for eight days*. That is also quite extraordinary. But which is *more* likely or *less* extraordinary? In Hume's estimation, it is more unlikely that observers from around the world would all agree on something that did not happen than that some strange astronomical event or weather event happened. It's more likely that something weird happened in nature than that all of these reports would be false.

But here is the second case Hume considers. Suppose we are reading records from English history and we find reports that on the first of January 1600, Queen Elizabeth died and was confirmed dead by court physicians. A successor was put on the throne, as would usually happen in such cases. Then, the reports say that, one month later, the queen arose from the dead, resumed the throne, and governed England for three more years. The records, let us presume, are the sorts of records historians typically rely on as they try to trace all the details of English political events.

What is rational for us to believe in this second case? We might think that it is like the eight days of darkness and that we ought to accept that nature does some pretty wild things sometimes, including bringing British monarchs back from the dead. But Hume writes that he would not have the “least inclination” to believe the reports. Why not? Hume answers that “I should not doubt of her *pretended* death and of those other public circumstances that followed it: I should only assert it to have been *pretended*, and that it neither was, nor could possibly be real.” In other words, Hume thinks it more likely that something tricky was going on. The queen’s death was faked, and her resurrection was faked, possibly to gain some sort of political advantage, for who would not be faithfully obedient to a queen who arose from the dead?!

So, in this case we are weighing the possibility of real death and real resurrection one month later from being dead against the possibility of some rather large-scale hoax being perpetrated on the English public. It is difficult to perpetrate such a large-scale hoax. But that is nothing compared to the difficulty of resurrecting someone who has been truly dead for a month. So, Hume concludes, it is far more likely that there was a hoax due to “the knavery and folly of men” than that the queen truly arose from the dead.

Why doesn't Hume also suspect a hoax in the eight days of darkness case? We might consider just how massive and difficult such a hoax would have to be. How would you trick people from different societies all around the world to believe that the world was dark for eight days? Or perhaps we need not go that far. Perhaps we only need a dedicated team of hoaxers to infiltrate every society around the world and doctor all the records to make it seem as if there had been eight days of darkness. But either way, that is a truly massive hoax. At some point it becomes more likely to judge, in this case, that the strange event really happened than that such a hoax happened. But the English hoax is not nearly so massive. We perhaps need only a dozen people to agree to a cover-up operation and keep

quiet about it. That is difficult, to be sure, but (again) not nearly so difficult as resurrecting someone from the dead.

Objection: But we should remember the Indian prince and the fact that he came to the wrong conclusion about water freezing into ice. Perhaps in the right circumstances people can be resurrected after being dead for a month.

Hume would grant the objection. Nature certainly can be surprising. But, like the prince, we can only make judgments about what is rational for us to believe given our previous experience. The prince was rational to deny the reports of the visitors given his previous experience. We are similarly rational, given our experience, to deny the queen's resurrection. Of course, we might be wrong. But, *given our experience*, there is going to have to be a lot more evidence to make the death and resurrection of the queen more likely than the possibility of a hoax.

Hume discusses these three cases while arguing for what was, in his day, an extremely radical and dangerous claim: that a rational person should not believe in the miracles reported in

the Bible. The Bible presents reports of extraordinary events like a flood covering the entire globe, the Red Sea parting, angelic visitations, resurrections, and so on. The Bible itself is a collection of ancient documents written by various people at different times, and none of the extraordinary events are corroborated by any other texts. We typically do not believe other ancient texts, like Homer's *Iliad* and *Odyssey*, when they report the interventions of the gods or visits to the underworld or creatures like Circe or the Cyclops. Why, then, should the Bible be treated as any more truthful in such matters than any other ancient text that describes similarly extraordinary events?

Hume puts his point quite forcefully (and exhibiting an anti-semitism which was all too common, even among so-called Enlightenment thinkers):

Here we are first to consider a book, presented to us by a barbarous and ignorant people, written in an age when they were still more barbarous, and in all probability long after the facts which it relates, corroborated by no concurring testimony, and resembling those fabulous [fable-like] accounts which every nation gives of its origin. [...] I desire any one to lay his hand upon his heart, and after a serious consideration, declare *whether he thinks the falsehood of such a book, supported by such a testimony, would*

*be more extraordinary and miraculous than all the miracles
it relates ...*¹

The last point which I have put into italics demonstrates what Hume is weighing as he considers whether to believe the Bible's reports of miracles. On one side, we have the possibility that the reports are false. On the other side, we have the possibility that all those events really happened and have been accurately described. Which is more probable? We have plenty of examples of ancient texts (and not-so-ancient texts) presenting false accounts. There is nothing extraordinary about this. But the events being reported—the flood, the parting of the sea, and so on—are about as extraordinary as any report could possibly be. Indeed, they are deemed as *miracles*, meaning events that are never observed to happen in the ordinary course of nature. In Hume's view, we should not hesitate to deny the Biblical reports.

Hume offers a general rule to follow in all these cases we have been examining: "That 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

1. David Hume, *An Enquiry Concerning Human Understanding* (Indianapolis: Hackett, 1993).

endeavours to establish.” In other words, don’t believe the reports of an extraordinary event unless the falsehood of that report would be even more extraordinary than the event it reports.

Media Attributions

- [Likelihood Scale](#) adapted by Charlie Huenemann is licensed under a [CC BY-SA \(Attribution ShareAlike\)](#) license

BAYESIANISM

We can now use what we have seen in Hume's discussion to understand the basic idea of Bayesianism. In all the cases Hume discusses, he asks us to use our previous experience of people's behavior and the workings of the world when we try to decide whether to believe a report. The Indian prince bases his decision on his own previous experience, the eight days of darkness case asks us to consider just how hard it would be to hoax the whole world, the Queen Elizabeth case asks us how hard it would be to hoax a smaller public, and the Bible case asks us how strange it would be for an ancient text to report things that did not happen. All of these cases require us to first consider what we first take to be generally true from our experience and then to assess the new information in relation to that general experience. Paraphrasing Hume, we ask whether the falsity of the new information would be *more extraordinary*, given our experience, than believing what the information says, given our experience.

This is the core of the Bayesian method. The method is named

for Thomas Bayes (1701-1761), an English mathematician, philosopher, and minister. Instead of viewing events as probable or improbable on the basis of how frequently they happen, Bayes asked how confident we should be about some new bit of information given our other beliefs, or what degree of probability we should place on the new information. It is therefore a *subjective* view of probability since it focuses not on the “real chance” of an event happening but on how likely we should think such an event is.

Bayes expressed his method through a precise mathematical theorem. We will not need to delve into that theorem to understand what the core idea is (in case you are interested, the theorem is given and explained briefly at the end of this chapter). But Bayes’ theorem is most often explained in terms of a situation in which someone is being tested for having a certain disease, and some math is always used in these cases. In the interest of providing a general introduction to Bayesianism, we will take the time to work through one of these typical medical cases to see how it works. Then, we will make the underlying idea more general.

1. A medical example

Suppose Patrick has been exposed to some rare but frightful disease. Only one person in a thousand has it. He has not

shown any symptoms, and has no other reason to think he actually has it, but he is worried, and so he goes to a doctor to be tested. Alas, the test comes back *positive*. He wants to know how accurate the test is. The doctor carefully explains that the test is accurate 90 times out of 100, and only generates false positive results in 5 cases out of 100. It is the worst day of Patrick's life, for he now believes he has a 90% chance of having this terrible disease.

But should Patrick be this concerned? This is where Bayes' theorem does its work, and the results are surprising. Ordinarily, before any tests or anything else, how worried should Patrick be about having this disease? Only one person in a thousand has it, so Patrick should not be worried very much. But now his positive test result comes in, and the test "gets it right" 90 times out of 100. Following the core idea of Bayesianism, Patrick needs to assess the test results in the light of his previous knowledge, namely that the disease only affects one in a thousand.

So what should Patrick consider? He needs to think how this test result changes his earlier guess about how likely it is he has the disease. He needs to compare how likely it is he is in the group of people who have the disease and would test

positive and how likely it is he is in the group of people who do not have the disease but still end up testing positive. If we have a million people, there will be 900 people who have the disease and test positive (that's 90% of the 1,000 people who will have the disease in a population of a million people). On the other side, the test offers a false positive result 5% of the time, which means that out of the one million people we are considering, 49,950 people do not have the disease but would still test positive for it (that is 5% of the 999,000 people who do not have the disease). So really, Patrick should consider which is more likely: that he is among the 900 people who have the disease and test positive? Or that he is among the 49,950 people who do not have the disease but still test positive?

We are now looking at a smaller group of 50,850 people out of a million who would test positive for the disease whether they really have it or not. Patrick should consider whether he is in the diseased group of 900 out of this 50,850, or in the undiseased group of 49,950 out of 50,850. All other things being equal, the odds of being in the first group is about 1.8%. The odds of being in the second group is over 98.2%. **So, the positive test result tells him that he should update his chances of having the disease from about a one in a thousand to about a two in a hundred.** Admittedly, this is a significant increase. Patrick should be more worried than he was before the test. But he should not be *a lot* more worried.

He certainly should not believe there is a 90% chance that he has the disease!

Here is our reasoning, laid out in a table:

In a population of 1,000,000 people, how many have the rare disease?	Of the 999,000 <i>without</i> the disease, how many will test positive?	Of the 1,000 <i>with</i> the disease, how many will test positive?	Patrick tests positive; what are the odds he is in either of these groups?
999,000 do not have the rare disease	49,950		49,950 out of 50,850, or ~ 98.2%
1,000 do have the rare disease		900	900 out of 50,850, or ~ 1.8%
(total number of people who test positive = 50,850)			

This is a surprising result (and no small relief to Patrick!). It comes about because Patrick is reminded that the disease itself is very rare, and when the rarity of the disease combines with the imperfection of the test, the numbers end up being not as threatening as one might otherwise think. When we jump from the claim that a test is accurate 90% of the time to the conclusion that Patrick, testing positive, has a 90% chance of having the disease, we are forgetting about just how rare it is for anyone to have the disease in the first place. Once we remind ourselves of that background knowledge, and we do our math, we have a much more accurate sense of how worried Patrick should be. There are over *fifty times* as many people who *do not* have the disease and test positive than there are people who *do* have the disease and test positive—just because the disease is so rare.

Suppose Patrick figures all this out, and he is somewhat relieved, but he still knows that he has a greater chance of having the disease than before. He would like to know with greater certainty whether he actually does have the disease. What should he do? **He should be tested again.** Remember, the number of people (out of a million) who had the disease and tested positive was 900, and the number who did not have the disease and tested positive was 49,950. Let's imagine

this whole group taking the test again. Again, we are assuming no one is showing any symptoms, and we have only the test results to worry about. Out of the 900 people with the disease, 810 will test positive again (for the test correctly catches the disease 90% of the time). Out of the 49,950 who tested positive without the disease, only about 2,500 will test positive again (for the test gives false positives only 5% of the time). So, now we will have a smaller group of 3,310 who have tested positive a second time. If Patrick gets a second positive result, then he is either among the diseased 810 out of 3,310, or he is among the undiseased 2,500 out of 3,310. There is a 24% chance of being in the first group (diseased), and 76% chance of being in the second (undiseased). This is more worrisome, as he now has a 1 in 4 chance of having the disease. If he tests positive a third time—and you can do the calculation on your own!—his chance of having the disease is 95%.

By the way, if the disease were far less rare—affecting 1 in 100 people, for example—and all of our other numbers stayed the same, getting a positive result on the first test should tell Patrick that he has a 15% of actually having the disease; a second positive result should tell him that his chance of having the disease is 89%.

The moral of this story is that *you should always get a second opinion*. It also shows that the new information we receive has to be positioned correctly relative to our broader knowledge of the world. When a new piece of information comes your way, you need to remember what your previous experience of the world is and let that previous experience guide how much importance you attach to the new piece of information.

2. Another example

Let's turn to another sort of case to see Bayesianism in action. Suppose you read that there has been a UFO sighting. Someone driving on a deserted highway at night reports they saw a glowing disk descend from the sky, hover over the ground for a minute, and then shoot back up into the sky. What should you believe? In particular, should you believe that this sighting is compelling evidence for the claim that Earth has been visited by intelligent extraterrestrials?

We should begin by considering what our previous experience says about the likelihood of alien visitors. From what we know, the universe is a very big place, and spaceships can only travel so fast (less than the speed of light), which means that it takes a very, very, very long time to travel from planet to planet or from solar system to solar system. There probably are

intelligent beings elsewhere in the universe just given how big it is, but the likelihood that they are anywhere near us is very, very low. So, our previous experience suggests that the likelihood of alien visitors is *extremely low*, perhaps on the order of one in a billion or one in a trillion or even less.

Now, let us consider the new information, the UFO report. We need to weigh the extremely low probability of alien visitors against the likelihood that this report is true. It is an extraordinary report; people do not often report similar experiences. What sort of probability should we attach to it? Let's consider it. On the one hand, people do very often offer true reports of what they experience. That's normal. But on the other hand, sometimes people lie about extraordinary experiences. Sometimes they *seem* to have extraordinary experiences, but the experience is due to psychological stress or mishap. Sometimes such experiences might be caused by strange weather phenomena, rare distortions of light, or anything other than alien visitors. There are many alternative explanations.

We have to combine all of these possibilities into an overall estimate of how likely we think it is that the person actually saw an alien spaceship. If we guess that all of these possible

explanations for their experience—really seeing a ship, lying, psychological stress, weather—are all *equally likely* (which is a generous guess!), we might think the chance of the experience really being a sighting of an alien spaceship is about 1 in 4. Put another way, by our estimate, it is three times more likely that someone *wrongly believes* they have seen an alien spaceship than that they have truly seen one.

Now we put that 1 in 4 chance in the broader context of the overall extremely low probability of alien visitors. (Note that this case is similar to the case of Patrick and the very rare disease.) The chances that aliens have visited earth, and that this person actually saw them, is extremely low (one in a billion, say), mainly because the likelihood of alien visitors, given what we know about the vast distances of space, is extremely low, and the likelihood that the person falsely believes themselves to see an alien spaceship is so much greater. So, in all, **the report of a UFO should not cause you to significantly change your belief about alien visitors.** It is far more likely that the report is coming from some other cause.

But what if *lots* of people report such experiences? This might change our belief, depending on the details. If by “*lots*” we

mean thousands of reports coming from individuals driving at night on lonely highways, our beliefs should not change all that much. It is still far more likely that their experiences are coming from something other than actual alien visitors. But suppose by “*lots*” we mean *millions* of reports from people driving on highways and in crowded sporting events and public assemblies, plus observations of alien spaceships on radar and from orbiting satellites and from professional astronomers and even a video recording from the International Space Station of aliens cruising by and waving from their spaceship window. (It is going to take *a lot* if we are to overcome the initial extremely low probability of alien visitors.) If these are the reports we are considering, we are in a case like Hume’s eight days of darkness. The likelihood of so many reports coming from some cause other than alien spaceships becomes extremely low, and we should regard the possibility of alien visitors as far more likely.

Again, the general rule David Hume offered can provide a quick assessment. Which would be the greater “miracle”? Is it more of a miracle for the reports to be explained by some other cause or for the reported events to have actually happened as described?

45.

QUESTIONS TO CONSIDER

1. Hume finds it far more likely that there was some sort of conspiracy in his Queen Elizabeth case than that she was resurrected from the dead. But doesn't this just show that he is not being open-minded? Shouldn't an epistemic agent be impartial?
2. Suppose that in our UFO example, someone regards the initial probability of intelligent life forms visiting Earth as being pretty high—because there are so many credible reports of UFOs! What might a good Bayesian say to them?
3. Bayesianism suggests that one should not believe that extremely improbable events happen. But extremely improbable events *do* happen from time to time. Isn't this enough to show that Bayesianism gives bad advice?

46.

FURTHER READING

As mentioned, Bayes' Theorem is usually presented as a formula. The basic question is, "What is the probability of {some event} given this evidence?" The idea is to make the probability *proportional to* both the probability of the event on its own and the probability that the evidence would exist if the event happened and to make the probability *inversely proportional* to the probability of the evidence existing, even if the event didn't happen. So, we have this:

$$\text{Probability of event, given this evidence} = \frac{(\text{probability of the evidence existing if the event happens}) \times (\text{probability of event on its own})}{(\text{probability of the evidence existing anyway, even if the event does not happen})}$$

Or, letting A = the event, B = the evidence, and " $P(A|B)$ " meaning "the probability of A given B ", we have:

$$P(A | B) = \frac{P(B | A) \cdot P(A)}{P(B)}$$

There are *many* excellent videos about Bayes' Theorem on YouTube. One that is exceptionally clear is from 3blue1brown, "The medical test paradox: Can redesigning Bayes rule help?", <https://youtu.be/lG4VkPoG3ko>.

A fascinating book about the powerful effects of Bayes' Theorem is Sharon Bertsch McGrayne's *The Theory That Would Not Die: How Bayes' Rule Cracked the Enigma Code, Hunted Down Russian Submarines & Emerged Triumphant from Two Centuries of Controversy* (Yale UP, 2011)

Hume, David. *An Enquiry Concerning Human Understanding*. Indianapolis: Hackett, 1993. Hume's argument against believing reports of miracles is chapter 10 of his *Enquiry Concerning Human Understanding*. It is a classic!

Media Attributions

- Bayes Formula © Charlie Huenemann is licensed under a [CC BY-SA \(Attribution ShareAlike\)](#) license

- Figure 10.2 © Charlie Huenemann is licensed under a [CC BY-SA \(Attribution ShareAlike\)](#) license

PART XI

11. EPISTEMOLOGY AND THE INTERNET

Nobody can be told what the matrix is. You have to see it for
yourself.

— Morpheus, in *The Matrix* (1999)

Much of epistemology, like much of philosophy, focuses on what stays the same in human experience regardless of technological progress: we know the truths about ordinary things, like how many people are in our family, or whether it is raining, or whether we are hungry. The truths about such ordinary things hold as well for 21st-century astronauts as for neolithic people. But this does not mean that technology cannot offer new kinds of philosophical questions and provide new possible philosophical answers. The wise philosopher is the one who can reliably discern what stays the same and when a difference really makes a difference.

The advent of the internet in the late 20th century is a difference that makes a difference. It has changed economic and political landscapes, and it has deeply transformed our sense of art, culture, and communication. Most of the world's population now takes for granted that they are able to locate themselves quite precisely on the globe and come into near-instantaneous communication with nearly anyone else on it, and they also can access virtually the entirety of human knowledge and human history through a device they carry in a pocket. Perhaps it is as true as ever that truth is truth and knowledge requires justification, but what is believed, what is known, how it comes to be known have changed about as dramatically as can be imagined.

But along with this profound transformation in knowledge, there is a rising tide of *false knowledge*, or claims that seem to be knowledge but, in fact, are not. Paradoxically, this has made knowledge both easier and harder to get, since we have such easy access to a great domain of knowledge that is thoroughly entangled with lies, deceptions, distortions, and misinformation. To be responsible knowers, we need to develop new skills, new questions, and new sensibilities, for we are living now in a world quite different from the world in which we evolved.

We will begin by trying to put the epistemological changes brought on by the internet into a historical perspective. Then, we will turn to the crucial topic of algorithms and the role they play in our interactions with the internet. We will then conclude with some observations of the challenges forced by internet epistemology.

47.

INFORMATION IN HISTORICAL CONTEXT

Humans have existed in their modern form for around 200,000 years, but only for the last 5,500 years have humans been writing anything down. So, for 97% of our history, our knowledge has been limited by our *memory*. People can remember a lot—particularly when memory is routinely developed and exercised and when lives depend on it—but even so, the capacity of human memory is sharply limited. Even less information can be reliably passed along from generation to generation when memory is the only storage resource since there are only so many stories that can be invented and passed along. This means that for most of our history there have been sharp limits on preserving old information and making use of it in new contexts. Not that there was *no* preservation of information, of course; traditions, epic stories, folk wisdom, religious lore, cultural practices, and ethical mores all are ways in which knowledge from the past can be carried into the future without writing anything down. But when these are all there is, the survival of any item of knowledge is very precarious. Evidently, human memory,

together with these cultural practices, have provided enough for us to “get by,” but they have provided scarcely more than that.

So, for the great majority of our species’ existence, **nature itself has served as a limit on the flow of information.** Brains can only remember so much; only a portion of that memory is encoded in songs and stories and passed along to new generations; only a portion of those songs and stories continue to be retold over subsequent generations. There may have been individuals who attempted to suppress stories or distort them, but their attempts at controlling information are negligible when compared to the inherent limitations of nature.

The development of writing (c. 3500 BCE) punched a hole through these natural limits, as information could then be stored in something more durable than a human brain. Still, very few humans learned to write and read, which meant that there still remained a tight control over what information was preserved and exactly how it was preserved—in other words, which stories were written down and how they were written. Praises for kings, the details of economic transactions, and religious myths seem to have been common subjects for the

earliest encoders of information (or *scribes*). Nevertheless, writing itself allowed for more information to be passed and stored from generation to generation, which assisted in the growth of more complex societies. This increase in stored information coincided with more extensive trade networks and more cultural commerce among civilizations, generating even more information, of which more and more came to be written down through the works of humanists and philosophers. A notable result was that the amount of information accessible by humans came to be greater than the amount of information actively known by humans. Literate societies “knew” more than their populations did because information was stored in libraries of unread books.

Nearly 50 centuries later, the invention of the printing press (c. 1450) tremendously amplified the production and dissemination of information, and the flow of information became a flood. More and more people had access to more information as literacy rates increased. While political states and religious authorities still exerted some control over what was published, these controls were overcome by an ever increasing popular demand for information which was seemingly without limit. More people read, owned, stored, and even wrote books on ever greater varieties of topics—some fiction, some nonfiction, some ludicrous or scandalous, and a great many falling into all of these categories at once. It quickly

became impossible for anyone to control information or to control *disinformation* or falsehoods posing as information.

The development of information technologies throughout the 20th century, culminating in the creation of the internet, represents a jump in the history of human information processing capacities that is greater in scope to the development of writing and the invention of the printing press. *No other innovation has given more people easier and faster access to more information.* There is very little that is known that cannot be shared instantly with anyone with a connection to the internet. The greatest repositories of information in the history of the world are literally at our fingertips.

But the creation of the internet has led to more problems, or, at any rate, monstrous enhancements of older problems. The chief problems have to do with *searching*, *filtering*, and *controlling*. For centuries there have been the problems of tracking down a particular text in one library or another; of sorting through irrelevant information to find what is relevant; and of limiting the spread of information that is considered false, misleading, or dangerous. In the past, these problems were confronted through information management systems

(such as the “call number” system used in libraries) or through the specialized training of scholars in schools and universities or through political or religious attempts at censorship. But these traditional, human-based efforts at managing information are no match for the modern engines of information production. (Measurement on these matters is tricky, but by one estimate, the world produces 2.5 quintillion bytes of data *each day*. Whatever exactly that means, it is a lot!) There is simply no way that humans can manage all the information that is available. Hence the need for **algorithms**.

ALGORITHMS

An **algorithm** is any sort of routine procedure. A cake recipe is an algorithm for making a cake: if you mix together the ingredients, put them in a pan, put the pan in an oven, and wait for some time, at the end, you will have a cake. Some paper assignments given to students can be done algorithmically: state a thesis, present some arguments, write up a conclusion, and turn it in. The virtue of an algorithm is that it breaks a larger project into a set of smaller steps which, when done in the right order, complete the project. It does not matter who performs the smaller steps so long as they are done correctly. In some cases, the steps are small enough and simple enough for a mindless machine to do them. In that case, when a simple machine can perform an algorithm, we call the algorithm a **program**. One might think of the laborious “recipe” we follow when we perform long division. That algorithm, or one like it, can be performed by a mindless calculator, thank goodness.

With the advent of computers in the 20th century, more and

more of our information has been coded into data that computers can process and manage, which means that more and more of the information at our disposal is processed by algorithms. As I type these words, an algorithm is taking the electrical signals from the keys I press and storing them as numbers; other algorithms are turning those numbers into commands to light up a few pixels on my screen so that I can see what I write. And of course, matters get increasingly complicated from this point forward as I send a file to you through the air and over some wires and your machine receives the signals and you pull up the document and read it. The algorithms in ordinary laptop computers are like complicated factories of routines, all performing their narrow operations so as to produce overall effects that we take for granted—until, that is, something goes wrong and a file won't load and we curse the machines for being so stupid!

When we turn to the internet to search for the things we are interested in, armies of algorithms take in the information we give them and search for other collections of stored information that “match” what we are looking for (at least according to the programming of the algorithms). The algorithms, of course, do not know what they are doing; they are mindlessly following recipes which (if all goes right) end up with results that satisfy us. Since we cannot count on the algorithms to have any common sense or to know what they

are doing, the programmers of the algorithms have to rely on certain tricks that will get the algorithms to do what we want. Search algorithms will look for the websites that most people have ended up going to when they typed words similar to what we typed; they will rank websites according to how popular they are by some measure or other; they may even take a peek at your own history to try to gauge which sites are more likely to satisfy your interests. It is far from foolproof, of course. If you are interested in the historical and cultural background of cockfighting, be careful what search terms you employ or you may be presented with images not strictly relevant to your inquiry.

Exactly how search engines do what they do is a closely guarded secret because search engines are *very big business indeed*. They are big business because the companies that provide search engines use them as opportunities to provide you with information you did not exactly ask for. This is *advertising*. If you are interested in baking cakes, you might also be interested in buying special baking pans, a stylish apron, or a new mixer, and advertisements for such products might appear somewhere on your screen. Companies hoping to sell these products pay search engine companies to put those ads on the screens of people who are likely to be interested in

the products. The strategy is far more focused and far more effective than placing an ad in a *newspaper*.¹

But the business model of search engine companies does not stop there. As you search for items, search engines also gather data about *your* interests. These data are compiled together with data from all other users so that high-level algorithms can discern larger patterns of human interests and behavior. It may turn out, for example, that people interested in recipes for carrot cake are also more likely to be interested in folksy aprons, and also more likely to be interested in magazines celebrating rural lifestyles. Perhaps they also tend to vote Republican and have pro-life views. Perhaps they also are more likely to buy domestic automobiles and air fryers. I am making up these correlations for the point of illustration, but search engines are actively gathering data to make far more secure assessments as to what sorts of people like what sorts of things. These *metadata*—or data about the data reflecting people’s online activities—are the real source of wealth for search engine companies. The aim is to know people better than they know themselves. Search engine companies do not exactly sell

1. *Newspaper*: a 20th-century artifact made of paper on which was printed news, advertisements, and comics, and delivered to people’s doorsteps; you can find images of such ancient relics on the internet.

this information outright, but it plays a central role in a complex, multi-level process in which advertising space is auctioned off to companies in an automated process known as *real-time bidding*. Basically, real-time bidding is an algorithm that sells access to user's information to other algorithms so that client companies can mount more effective ad campaigns.

It probably seems to you that this sort of information may be valuable, but it cannot be the most valuable thing in the world. But you are wrong. In 2017, *The Economist* announced that information had surpassed oil as the world's most valuable resource. There is more money to be made in gathering information about people's buying habits than there is in selling them any particular thing.

Algorithms are like vast ant colonies. An individual is fairly simple and robotic in its behavior. But when assembled into great colonies, those simple robots can manage to accomplish extraordinarily complicated tasks. Indeed, they can accomplish *any* task that can be broken down into simple steps. It practically does not matter how many simple steps need to be performed since we have limitless supplies of “ants” to put to work. The production, storing, searching, filtering, and control of information in the modern world—and even the

buying and selling of it—is done primarily by algorithms; the role of humans now is to be sources of information and to consume in an economy based upon it.

There are two broad lessons to be gathered from this discussion of algorithms. The first lesson is that algorithms are mechanical and mindless in the sense that they do not know what they are doing and proceed according to rules in robotic fashion. This means they can make “mistakes” (from the point of view of our own expectations) without realizing it or without anything going wrong in their programming. As we will see, this also means that algorithms can be “gamed” or manipulated in clever ways so as to produce disinformation to users. The second lesson is that there is a lot of power and wealth connected to algorithms and tremendous incentives to find ways to gain control over them. Were Francis Bacon alive today, he would say that “Algorithms themselves are power.”

...

KNOWING THROUGH THE INTERNET

In many ways, for many purposes, the internet provides the most accurate resources of information ever available to human beings. A case in point is Wikipedia. Let us first admit that many valid criticisms can be made of Wikipedia. On unpopular or relatively obscure topics, Wikipedia merely reproduces seriously dated publications that are in the public domain. On some topics, the presentation is shaped by amateurs with uninformed and peculiar points of view. Topics can be hijacked by political operatives. Information from more authoritative sources is not always accorded greater value. There is no systematic practice of fact-checking. Many of the entries read as if they were composed by disorganized committees of volunteers (which they are). And more criticisms can also be made; a long list of them, in fact, can be found in the Wikipedia entry entitled “Criticisms of Wikipedia.”

But despite these criticisms, and despite persistent injunctions of college professors against using Wikipedia as a source, *it is without a doubt the greatest single source of knowledge ever assembled*. The more responsible epistemic agent will always balance whatever is said on Wikipedia with a broader survey of other more authoritative sources, but for a quick and *mostly* accurate overview of the *widest* array of possible topics, no other encyclopedia even comes close. Even if Wikipedia is a second-rate (and sometimes third-rate) resource on each particular topic, there is no other base of knowledge that can come close to its range of coverage and general level of accuracy.

And, of course, Wikipedia is not the only available source of knowledge on the internet. The internet gives us access to first-rate scholarly journals, news media, blogs by true experts of obscure matters, maps, lectures, and so on without limit. The internet makes it easier than ever before to take a broad sampling of different accounts of nearly any topic, and to form judicious opinions based upon that diversity of resources.

But this is only, at best, *half* of the story. As recognized at the end of the previous section, internet searches can be “gamed,” and there are powerful incentives for distorting what users find

as a result of their searches. This is shown most dramatically in instances of exploiting *data voids*.

As expansive and comprehensive as the internet is, there are topics that have a very minimal presence on the web. It may be a set of words that is not commonly used, or it may be a person or event or little town about which people have very little to say. We can call such a neglected entity a “data void,” which simply means that there is not much information on that particular topic on the internet. If, for whatever reason, someone wishes to tell some particular story about these data void entities, they can tell that story in multiple places throughout the web, and mindless algorithms will direct users to that story if they happen to search for that entity. A data void is thus an opportunity to establish and control a narrative.

For example, the term “crisis actor” was for some time a data void. No one searched for that pair of words. But, according to recent media and internet researchers, at some point malicious individuals seized the term and populated the internet with many false stories about people who were hired to pretend to be victims of mass shootings. Multiple websites offered seemingly genuine accounts by “crisis actors” who admitted to having portrayed victims of various faked or staged mass

shootings. Efforts were then made by the malicious individuals to get the phrase “crisis actor” mentioned on some national media outlet. The efforts succeeded, and when viewers of the media went to search the term “crisis actor” they found multiple accounts from different sources of people admitting to playing roles in staging fake shootings. This gave support to baseless conspiracy theories about the government staging mass shootings. Similar data void hijackings have been executed with the terms “collusion hoax,” “black on white crime,” and “pizzagate.”

This is an example of malicious agents exploiting the mindless operations of algorithms to bring baseless conspiracy theories to a broad audience’s attention. But there are also less outrageous attempts at doing the same thing. Political organizations can promote specific terms and slogans and make sure that they direct the public narrative by establishing websites that become the go-to sites for searches employing those terms: this is called *strategic keyword signalling*. These are efforts to “game” the mindless functioning of algorithms so as to exert influence over what broad communities take to be truth.

Of course, as a student in epistemology that has thought

through skepticism and the Grand Deception Doubt, you will naturally wonder whether “strategic keyword signalling” is itself an exploitation of a data void. Perhaps the experts in media studies are doing exactly what they are accusing the people behind “crisis actors” as doing! And a terribly destructive seed of skepticism is thereby planted. We may begin to suspect that we can no longer trust *anything* we find on the web. “Anyone can ‘prove’ anything by posting false information and manipulating the internet’s algorithms,” someone might think. “We cannot ever know anything, and so we might as well choose to believe *whatever story we like best*.” These extreme doubts will only be nourished by our observation of how much power and wealth there is in the control of algorithms. There are strong incentives to control the knowledge of individuals, and the private companies hosting search engines have unparalleled power and incentives to exert that control. With so much at stake, how can information not be thoroughly biased and skewed toward the interests of those in power?

But as students of epistemology, we also know how to begin to think our way through these extreme doubts. If we adopt a basic Humean or Bayesian outlook on the information we are coming across, we shall start to assess the occasions where we have solid reasons for doubt and those where the reasons are less solid. The claim that public shootings are entirely

fabricated as the result of massive government conspiracies is an *extraordinary* claim and should require *extraordinary* evidence—far more extraordinary than a small collection of obscure websites. The claim that *everything* on the web is fabricated is even more extraordinary, and finding evidence for such a claim through conspiracy-theory websites leans decidedly in the direction of being a self-refuting justification. The more plausible claim is that there is genuine information and disinformation on the web, and patient inquiry and reasoning is required to sort the more likely from the less likely or the reliable from the unreliable—as has always been the case in human knowledge.

50.

SOME ADVICE

Perhaps the most important advice one can give regarding the problem of what one should trust when trying to gain knowledge through the internet comes from a school of philosophy known as *hermeneutics*. Hermeneutics is the study of interpretation. Particularly, it is the study of how to interpret texts whether those texts are books or websites. Authors have always been tricky and strategic in providing information. Sometimes they are trying to advance one special cause, and sometimes they are trying to appear as advancing one cause while really—“between the lines”—advancing some other cause. Readers have to be very critical and self-reflective when they read if they are to grasp what is really being said in a text.

The French philosopher of hermeneutics, Paul Ricoeur (1913-2005), advocated a **hermeneutics of suspicion** especially when reading the tricky texts of Marx, Nietzsche, and Freud. These authors were very conscious of the effects words can have and often, Ricoeur argued, seemed to say one

thing while really meaning something else. They deliberately set out puzzles and paradoxes in order to prompt readers to think for themselves and reach some further conclusion that was not stated explicitly in the texts themselves. But setting aside the challenges of reading such difficult philosophical authors, for our purposes we might consider adopting a hermeneutics of suspicion toward our readings of online texts.

In adopting a hermeneutics of suspicion, we must first recognize that in any exchange of information, the provider of the information is attempting to cause the receiver to adopt some belief. If we are suspicious, we will begin to ask *who* is getting us to believe *what* and *for what purpose*:

Who? Who is providing this information? Are they in a position to have this information? Are they relying on other sources? How reliable are those sources?

What? What is this provider trying to get me to believe? Is that claim, given my prior experience, likely to be true? What sort of initial probability would I put on that claim, before reading the account of this provider, and to what extent should this new information change that probability?

For what purpose? Why would this provider want me to gain this belief? Is it this provider's "job" to

simply communicate truths — or are they working for some other cause?

The hermeneutics of suspicion asks us to adopt the sort of attitude toward information that one might adopt in buying a used car from a stranger. We will want to gauge the character of this stranger, what they are selling us, and for what purpose.


Very often, of course, a stranger is selling us a used car that has no hidden flaws, merely because they no longer need the car and they would like a fair amount of money for it. That's a good and very common scenario. Similarly, in many instances, the providers of information on the internet are simply trying to provide accurate information based on reliable sources so that readers gain accurate information. Reputable news media on the internet—the Associated Press (AP), for example—are dedicated to providing accurate news information. That is their brand, and they command a large population of readers precisely because they are seen over time to be reliable providers of information. Their “business model” is based on providing accurate and reliable information. It would be very hard to find one set of claims that the AP is trying to get its readers to believe, since stories are provided on a wealth of topics from a variety of angles, and there is not a single theme that runs through them all.

But other cases are far less straightforward. Some providers of information on the internet clearly provide only information that is meant to encourage a narrow range of beliefs in readers, and it is relatively easy to identify what those beliefs are. The sources that are relied upon are only sources that are interested in promoting those beliefs, and there is little or no discussion of other sources or support for other beliefs. These sites are like sellers of used cars who only point out the positive features of the car in question and refuse to answer any questions about recent repairs, mileage, oil changes, etc.

As in every case of trying to determine what sources to trust, we are always working from some initial information we have about the world and assessing new information on the basis of our existing information. Someone who is antecedently convinced that big news media like the AP are only offering stories to justify those in power, and, for example, are refusing to report widespread alien landings and the efforts of those aliens to take control over national governments, will provide a very different assessment of the reliability of those media. Or, somewhat more plausibly, if someone holds that big news media are thoroughly embedded in capitalistic economic structures, and they refuse to highlight the evils and injustices of those structures, they also will give less credence to the

reports of those media. We must always begin with the worldview we have, make our assessments from that starting place, and be willing to change our minds as new evidence is presented and as that evidence warrants changing our beliefs. (This is the lesson of Bayesianism.) A hermeneutics of suspicion should always be running in the background as we take on the endeavor of trying to gain knowledge through the internet.

For what it is worth, we may consider the judgments of media experts who take on the task of sorting through various news outlets and making assessments as to their accuracy and reliability. Ad Fontes Media, for example, is a crowd-funded organization that uses a team of about 20 analysts with varying political perspectives to rate the accuracy and objectivity of hundreds of stories from dozens of major news outlets. In the end, they provide a chart displaying their findings. Here, for example, is the chart as of January 2021:



Of course, this chart will only be authoritative to someone who antecedently agrees that Ad Fontes Media and the analysts they use are in positions to make assessments of objectivity and accuracy. But we might see what results if we confront Ad Fontes Media with our suspicious questions. (In what follows, I will provide answers as made available on Wikipedia since by this point it should be clear to everyone that I have drunk the Wikipedia Kool-Aid!)

Who? “Ad Fontes Media, Inc. is a Colorado-based media watchdog organization primarily known for its Media Bias Chart which rates media sources in terms of political bias and reliability. The organization was

founded in 2014 by patent attorney Vanessa Otero with the goal of combating political polarization. Ad Fontes Media uses a panel of analysts across the political spectrum to evaluate articles for the Chart.”

What? The Ad Fontes media chart is meant to show that some news sources are more reliable than others and some more biased than others. For the 2020 chart, for example, “nearly 1800 individual articles and TV news shows were rated by at least three analysts with different political views (left, right and center). There were 20 analysts, [and] each reviewed about 370 articles and about 17 TV shows.”

For what purpose? “Otero [the founder] sees the Media Bias Chart as an ‘anchor’ that counteracts political polarization in news media and aspires for Ad Fontes to become a ‘*Consumer Reports* for media ratings’. She compared low-quality news sources to junk food and described sources with extreme bias as ‘very toxic and damaging to the country’.”

Both Wikipedia and Ad Fontes Media are about as crowd-sourced as information sources can be. This suggests that, in the judgment of a *very broad consensus*, the chart represents an honest effort to communicate information about media bias for the sake of having smarter consumers of information. Anyone with serious doubts as to their impartiality should be expected to justify those doubts using all that we have learned

about Humean and Bayesian probability with a robust but sensible hermeneutics of suspicion.

Media Attributions

- [Figure 11.1](#) © Ad Fontes Media is licensed under a [All Rights Reserved](#) license

51.

QUESTIONS TO CONSIDER

1. Where do you get your news, and how do you sort out reliable from unreliable news sources? Can you formulate good rules to follow?
2. Some have argued that social media platforms (such as Twitter or Facebook) should be held responsible for the material they allow to be published and shared on their sites. What do you see as the advantages and disadvantages of such a policy?
3. Some have argued that the flood of information we experience has made our understanding shallower or less complete. “We have more facts, but less understanding.” Can you explain why this might be so? Would we know *more* if we had *less* information?

FURTHER READING

Michael Golebiewski and Danah Boyd, “Data Voids: Where Missing Data Can be Easily Exploited,” Data & Society, report, October 29 2019. <https://datasociety.net/library/data-voids/>

O’Connor, C & Weatherall, J.O. (2019). *The Misinformation Age*. New Haven, CT: Yale University Press.

Nguyen, C.T. (2018). “Epistemic Bubbles and Echo Chambers.” *Episteme*. DOI:10.1017/epi.2018.32

Tarleton Gillespie, “The Relevance of Algorithms,” in *Media Technologies: Essays on Communication, Materiality, and Society*, edited by Tarleton Gillespie, et al., MIT Press, 2014. (pp. 167-193). Excellent, clear account of the questions we should be raising about algorithms.

PART XII

12. CONSPIRACY THEORIES

“Just because you’re paranoid
doesn’t mean they aren’t after you.”

— Joseph Heller, *Catch-22*

At some level, people who believe in conspiracy theories are being paradigmatically rational. They are seeking out evidence for a claim, they are providing explanations grounded in reports and observations, and they construct a theory that is consistent and makes sense of the relevant data. And sometimes, of course, there really are conspiracies in which the conspirators successfully (for a time) manage to hide what they are doing or what they have done (Watergate is one such example). But the term “conspiracy theory” is normally used to refer to theories that are extremely implausible and even irrational. So, what has gone wrong? How does it happen that people with such praiseworthy epistemological virtues end up promoting implausible, irrational theories?

As we will see, a conspiracy theory emerges when someone accepts a claim as incontrovertibly true and then employs their rational abilities to reinterpret what they read or see so as to support that claim. So, there is a great degree of rationality employed in any conspiracy theory; it's just that it is in the service of an implausible end.

THE "SEED BELIEF" MODEL

As we will see, there can be many ways that a conspiracy theory comes to be, and people may have many different sorts of motivations in coming to believe them. There may not be a single general theory of conspiracy theories that captures the essential nature of them all. Indeed, there is good reason to think there is no set of features that all “bad” conspiracy theories have in common since some conspiracy theories turn out to be true. But we will consider one “rational reconstruction” of conspiracy theories that seems to be true for a great many of them.

We might see a conspiracy theory as beginning with a “seed belief,” or one that gets the theory planted. The seed belief may be that it is simply too difficult to send humans to the moon; that it is worrisome to put a chemical like fluoride into drinking water; that Elvis Presley is too important to have died from drug abuse; that a small group of terrorists could not have organized the terrible events of 9/11; that the Earth simply could not be spherical; that Barack Obama must be from Africa; that the vapor trails left by jet airplanes must

be dangerous; and so on. Each seed belief may seem to have a degree of plausibility, given some basic beliefs and without looking any further into the matter. A person may have some separate motivation for wanting the seed belief to be true (Elvis may have meant *a lot* to them, for example, or they may wish to have some reason for rejecting the legitimacy of an African-American president). For now, we will set aside the question of exactly *why* the seed belief gets planted; let's assume that, for some reason, it does.

The seed belief quickly sets down some strong roots by making use of the human tendency toward **confirmation bias** (as discussed in Chapter 8). Humans very naturally seek out evidence for beliefs they have, and it takes additional effort for any human to seek out evidence that goes *against* their beliefs. And it is relatively easy to find confirming evidence for many beliefs. Most information is either irrelevant to the seed belief (in which case it is at least consistent with the belief), or with a bit of further interpretation, the evidence can be interpreted in such a way as to be consistent with the belief. So, for example, there are newspaper reports of Elvis's death. But stories in newspapers can be faked, or reporters can be given false reports from the police or the coroner. There was a funeral for Elvis, but funerals can be staged. Elvis isn't seen in public anymore, so he could be in hiding. And when it eventually happens that some people report seeing someone who looked a bit like Elvis

(though he had lost weight and had grown a beard). **Aha!** The King is alive and in disguise! Confirming evidence!

Once the seed belief has secured itself in the mind of the believer, the believer will have to face a very awkward question. The question is *why there are not more people who share the seed belief*. The seed belief is supported by a lot of confirming evidence, after all, as the believer has discovered. The only explanation (so it seems) is that other people are being actively *mised*. They are constantly being fed some story that simply isn't true in an effort to keep the seed belief from spreading to more people. And so there must be a **conspiracy** that is manipulating people into false belief. A group of conspirators would do this only if they had something to hide, of course, which indicates that the conspiracy must be malicious. And because the conspiracy is so successful at keeping the seed belief from spreading, it must be an extraordinarily widespread and intelligent conspiracy. There is hardly any evidence for its existence, which may just show how crafty the leaders of the conspiracy are!

But the dedicated individual will be able to find evidence for the conspiracy by looking for clues in the right places and by discovering the absence of information where the theorist

believes there should be some. And through this process of finding evidence first for the seed belief, and then for the conspiracy keeping others from sharing the seed belief, the seed belief itself becomes *justified*. The theorist has interpreted what they have seen, read, and heard in such a way as to support the seed belief, which is no longer just a belief that happened to take root in someone's mind but is now documented with scores of newspaper articles, YouTube interviews, and independently-published books by the small circle of other people who share the seed belief. And now comes the clincher: why would there be all of this evidence if the seed belief were simply false? Why would people take so much trouble to find evidence for something that didn't happen? Why would this dedicated individual devote so much energy and effort to documenting something that isn't real? There is only one way to explain all of the work people are putting into justifying their belief: the seed belief must be *true*.

Two further patterns of reasoning also help a seed belief to grow strong roots: the conviction that a single explanation is better than a set of independent explanations, and the conviction that when someone benefits from an event, they must be causally responsible for it.

MONOCASUAL EXPLANATION, AND *CUI BONO* INFERENCES

Humans tend to understand events through stories, and stories are easiest to follow when there is a single line of causes. It is easy to follow a story of the form “A caused B, which caused C, which caused D.” It is much harder to follow a story of the form “A caused B, and meanwhile C caused D, and D kept E from happening, and when B happened without E also happening, F was the result.” For this reason, we are likely to favor stories with a single line of causes—a “**monocausal explanation**”—over a story that requires keeping track of separate lines of causation. (We might see monocausal explanations as one version of the **available heuristic** bias.)

Sometimes, of course, a monocausal explanation is perfectly in order, and there is no need to overcomplicate things. But very often we push for a monocausal explanation when the truth is more complicated. We can see this in many popular history

books which push for a monocausal explanation of why “the West” conquered the rest of the world, rather than the other way around. World history over the centuries of course includes many thousands of separate events, forces, pressures, and shifts, but recounting them all makes for an extremely complicated story. It is much nicer to be able to tell a single comprehensive story.

We can see the human inclination toward monocausal explanations at work in conspiracy theories. So, for example, if someone wanted to understand what is happening in our economy, they would have to study the behavior of markets, the supply of labor, effects of consumer demand, the roles of tariffs and regulations, and so on. It is safe to say there is not a single person who can master all of these separate causal economic influences in detail, and anyone who tried to offer even a vague general picture would be telling a very complicated story with many separate lines of causality, a story that would be very hard to follow. It is far more satisfying to simply believe in a secret group of people who control everything for their own economic interest, “a ruling global elite.” With such a simple theory we can explain every economic development as being caused by a single entity. No need to do all of that research! And when we see groups of powerful people assemble for meetings of organizations like the World Trade Organization, our suspicions will be

confirmed. They *must* be the ruling global elite. (Of course, this is not to say that the WTO is not enormously influential, but not even the WTO can control everything.)

In this example, a more plausible view is that our economy is an exceedingly complicated system with many agents and many causes. It is true that rich people have more advantages (that is what “rich” means), and certainly, they play pivotal roles in making big investments and enacting certain policies, and organizations like the WTO promote their interests. Poor people do not have nearly as much power, and they usually suffer as a result of those investments and policies. Many rightful criticisms can be made of many features of the overall system, and economists and public policy experts are busily raising these criticisms in mounds of articles and books. But none of these truths imply that there is a single, organized group of people “behind it all.” In fact, the very complexity of the overall system suggests that there could not possibly be a single group “behind it all.” How on earth could anyone manage “it all”? Nevertheless, all of that being said, it surely is tempting to believe there must be a single group “behind it all.” For then we have an easy story to tell, an easy story to understand, and someone to blame.

Conspiracy theories—and especially conspiracy theories of economics—also often trade upon “*cui bono*” inferences. “*Cui bono?*” is Latin for the question, “Who benefits?” Many events, of course, are to the benefit of certain people or institutions or at least are seen as beneficial to them. If we are likely to believe in monocausal explanations, then when we see some group benefit as the result of some action or change, we might well *infer* that the action or change was *caused by* that group. So, for example, spending on the military increased as a result of the terrorist attacks on 9/11, and the spending brought considerable benefit to military contractors. *So, were military contractors behind 9/11?* The state of Israel was established as a result of the horrors perpetrated on Jewish people by the Nazis in World War II and establishing a homeland for the Jews was a great benefit for them. *So, did the Jews fake the Holocaust, in order to get a homeland?* These are *cui bono* inferences, and they are more easily made when monocausal explanations are also believed—in other words, when there is a single group of people “behind it all” who will benefit from the events.

While raising a *cui bono* question is often important and can help to sort out questions of motivations of certain actors in a situation, the mere fact that someone benefits from a change never itself implies that they intentionally brought about the change. It might be true that, very often, humans act to benefit

themselves, but it does not follow from this that *whenever* humans benefit it is the result of their own strategic actions. Sometimes events benefit a group, but that group is not responsible for those events. And in the cases of 9/11 and the Holocaust, *none* of the people who benefitted by subsequent results would say that the benefits were worth the cost.

WHERE SEED BELIEFS COME FROM

So far, we have an account of conspiracy theories as growing from seed beliefs and being nourished by confirmation bias, a bias toward monocausal explanations, and the temptation to make *cui bono* inferences. The biases and temptations are perhaps easily enough accounted for. Our evolutionary past has not given us perfect reasoning capacities. But where do the seed beliefs come from? How does a person come to believe and then insist upon a belief that should be discarded after even a little bit of research?

Sometimes, a seed belief may result from some of our other weaknesses in reasoning such as **anchoring** or **in-group bias**. For example, we first hear that vaccines cause diseases or that Barack Obama is African or that UFOs are commonly seen from friends and family members we trust. The credence we give to those beliefs only has to be strong enough to get us to look for further evidence in support of the belief, or to

defend them against others' objections, and then the other features of our cognitive machinery will kick in to add more support for the belief making it stronger and stronger the more we defend it. What began as a simple belief becomes a well-defended theory and then an unshakeable conviction. In these cases, a conspiracy theory is seen as something like a parasite or infection that lands upon a hapless believer, and the believer's own cognitive machinery is harnessed to give the invader more strength and vitality. "Curing" the infection will be difficult as it will require somehow re-orienting the believer's entire cognitive system to cause them to recognize the invader as an invader and to begin to recover from the disease by re-examining one's entire structure of beliefs.

An example of this sort of re-orientation can be found in the case of Derek Black, the son of a grand wizard of the Ku Klux Klan, who was a young and rising star of white nationalism. Black was raised in an environment encouraging the belief that the races should be separate from one another and that the United States should be "kept pure" as a white nation. He believed that a vast government misinformation campaign was behind popular attitudes in favor of desegregation, multiculturalism, and antiracism.

Black went to college and encountered many people who wanted nothing to do with him or his beliefs. But a small Jewish community at the college invited him to weekly dinners and patiently offered evidence against his views. In an interview Black recalls the discussions:

I would say, “This is what I believe about I.Q. differences. I have 12 different studies that have been published over the years, here’s the journal that’s put this stuff together, I believe that this is true, that race predicts I.Q., and that there were I.Q. differences in races.” And they would come back with 150 more recent, more well researched studies and explain to me how statistics works, and we would go back and forth until I would come to the end of that argument, and I’d say, “Yes that makes sense, that does not hold together, and I’ll remove that from my ideological toolbox, but everything else is still there.” And we did that over a year or two on one thing after another until I got to a point where I didn’t believe it anymore (“Derek Black”, Wikipedia entry).

Over a year or two, Black became convinced that the entire framework of beliefs in which he had been raised was false. He publicly renounced his beliefs even though this came at the cost of alienating him from his family.

Black's case is one in which his seed belief in white nationalism was caused by his home environment, and his belief in a conspiracy theory grew from that seed belief. In other cases, however, a seed belief is chosen precisely because it is widely rejected by a culture that an individual rejects for one reason or another. In these cases, the more extreme the seed belief, the better because the individual wants to distinguish themselves from the wider culture judging the belief as "crazy." In these cases, adopting the seed belief is a conscious act of rebellion against prevailing norms, and the conspiracy theory that grows from the seed belief is meant as an indictment of the prevailing culture that so confidently denounces the seed belief.

The recent growth in Flat-Earthers may be an example of this. A Flat-Earther believes the Earth is a flat disk and that there is a conspiracy of scientists and others who brainwash people into thinking that the Earth is a globe. There were Flat-Earthers many centuries ago, but the view seems to have resurfaced in small communities in the 19th century, and a Flat Earth Society was formed in the 1950s. This community has been very small historically. But with the growth of the internet, the community of Flat-Earthers has grown considerably. In recent years, many celebrities have at least said they doubt the official account of the Earth's shape. Most famously, the rapper B. o. B. announced his skepticism and in 2017, started a campaign to send multiple satellites into space to document the true

shape of the Earth. (Previous satellite missions apparently were not trustworthy.)

There is something in the recent Flat-Earther phenomenon that goes beyond a seed belief “accidentally” taking root in someone’s set of beliefs. Many Flat-Earthers are eagerly embracing a belief *precisely because* it is at odds with the belief of a dominant culture; the act of rejecting a belief as obvious as the belief that the Earth is a sphere is a way to confront and deny the authority of the surrounding culture. The surrounding culture (in this case) affirms the value of science and the value of a history of progress in knowledge, but it is also a culture from which many people feel alienated, perhaps for ideological, political, racial, or economic reasons. In this case, accepting the seed belief is an act of rebellion against that surrounding culture. It is a declaration that the dominant culture has no authority over the beliefs of the individual.

The deliberate acceptance of implausible seed beliefs also seems to be the primary cause of the QAnon conspiracy theory. The fundamental QAnon seed belief is that the world is being run by a conspiracy of people who rape and eat children, worship Satan, and that Donald Trump is the only person who can save the world from them. (It is sometimes

further claimed that the Satan-worshipping pedophiles are members of a lizard race that lives below the surface of the Earth.) The conspiracy theory associated with this seed belief maintains that the U.S. government and governments around the world have managed to keep these facts undiscovered—except for one brave and anonymous individual, known as “Q,” who was somehow able to thwart the global conspiracy and post messages about it on 4chan, an online forum for various hate groups. None of the early adopters of this seed belief were raised to believe it, and never has any positive evidence been given for it; rather, people embraced the ludicrous belief precisely because it was utterly ludicrous, and in so doing, declared their epistemic emancipation from the entirety of world media, educational institutions, scientists, and governmental bodies.

Recall the Baconian claim that knowledge itself is power. Some conspiracy theorists develop their theories in order to develop a base of power to challenge the power of a dominant culture’s knowledge. In such cases, proving the falsehood of the beliefs will do nothing to diminish the believers’ confidence in their theories. The deeper issues of power inequities will have to be addressed. Better epistemic practices will not be directly helpful.

56.

REAL CONSPIRACIES

Of course, understanding that many conspiracy theories are false and believed for not fully rational reasons does not show that there never are any actual conspiracies. Some of the more infamous and true conspiracy theories in the United States include:

During Prohibition (1920-1933), the U.S. Treasury Department poisoned industrial alcohol in an attempt to discourage bootleggers from using it to make alcoholic beverages. But apparently not all bootleggers were concerned with public health, and they produced and sold the beverages anyway, resulting in thousands of deaths. The government secretly continued the practice until the end of Prohibition, despite knowing its effects.

In 1932, the U.S. Public Health Service conducted an experimental trial of a treatment for syphilis on several hundred African-American men in Tuskegee, Alabama, without securing their informed consent. Men with the

disease were never given adequate treatment for it and were never fully informed of their role in the experiment.

In 1972, President Richard Nixon authorized a break-in at the headquarters of the Democratic National Committee in the Watergate Office Building and then unsuccessfully tried to cover up his administration's involvement.

In each of these cases, portions of the U.S. government conspired to commit harmful actions while keeping their role secret. This meets the letter definition of a “conspiracy” which is when any group of people have a secret plan to do something illegal or harmful. Given that definition, we can plausibly suspect that many actions by many governments result from conspiracies or secret plans to cause harm, particularly in the areas of espionage and counter-intelligence. Coming up with these conspiracy theories— and determining how well they are supported by available evidence—is the job of watchdog organizations, investigative reporters, and (later) historians.

So, there most definitely are conspiracies, and some conspiracy theories are true. But not all of them are. So, we are brought once again to the difficult epistemic challenge of trying to sort out the true from the false or the reasonable from the

unreasonable. Are there any rules or indicators to help us distinguish between plausible and implausible conspiracy theories?

There are no rules that will reliably sort the true conspiracy theories from the false ones, but the following rules may serve as a set of helpful indicators:

1. The bigger and more powerful the conspiracy is supposed to be, the less likely it is real. Anyone who has managed a sizable group project knows how hard it is to get people to coordinate their efforts. The task becomes even harder, or impossible, if the shared effort is to cover up some harmful or immoral secret. So, the more people who must be included in the conspiracy, the bigger the lie that must be told, and the more harmful the thing being kept secret is, the more unlikely it is that a conspiracy will succeed.

2. Hanlon's Razor: "Never attribute to malice that which is adequately explained by stupidity." Many times, big events happen for stupid reasons or for not really any single reason at all. We tend to think that significant events must have significant causes, but in fact, the universe does not pay

attention to what we regard as significant. For example, the explosion of the Hindenburg was a horrific disaster. A huge, gas-filled airship burst into flames, killing dozens of people. Some think it *must* have been sabotage because the event was so horrific. But it is more likely that the explosion was caused by static electricity, lightning, or engine failure. Sometimes, significant events happen for relatively unimpressive reasons.

3. Beware of claims of conspiracy that cannot be falsified.

The frustrating aspect of thorough-going conspiracy theories is their seeming unfalsifiability. No matter what happens or whatever is uncovered, it will end up being used as proof either of the theory or of just how crafty and manipulative the conspirators are. There is practically nothing that could prove to a Flat-Earther that the Earth is not flat, for example, as every contrary bit of evidence is rejected as mere propaganda or as improperly-interpreted data.

4. Positing a conspiracy should be an explanation of last resort. Given how difficult it is to maintain any sizable conspiracy and how common it is that significant events happen for insignificant reasons, positing a conspiracy should be an explanation of last resort. If there is no more natural or plausible explanation of some event, and if we are sure that

all the evidence is genuine, then perhaps we must posit that some conspiracy is at work. But one must work with great honesty and objectivity to determine whether the evidence to be explained is genuine and whether there really is no more plausible explanation. One of the most popular quotes employed by conspiracy theorists themselves is from Arthur Conan Doyle's character Sherlock Holmes: "When you have eliminated the impossible, whatever remains, however improbable, must be the truth." Conspiracy theorists proudly proclaim this dictum of a fictional detective as their justifying principle. The problem is that most conspiracy theorists are actually eager to reach for the improbable, and the "impossible" explanations they have eliminated are perfectly possible, at least, once one has sorted out the real evidence from false or implausible reports.

These four considerations may help alert us to false conspiracy theories in some cases. But, again, sometimes there are real conspiracies, and the theories about them are true. In the end, the best advice that can be given is to take each theory on its own merits, assessing the plausibility of its claims given our prior beliefs, our general experience of the world, and the basic attitude in Bayesian reasoning. We need to compare the likelihood that we would be seeing the so-called evidence for the theory if the theory were true against the likelihood that we would be seeing it anyway even if the theory were false.

57.

QUESTIONS TO CONSIDER

1. Arguing against conspiracy theorists is often frustrating because the theorists seem to be able to handle any evidence that goes against the theory. But then again, someone with a *true* theory should be able to handle any evidence that goes against the theory. So what's the difference?

2. Being willing to change your mind when there is good evidence against what you believe is usually considered an epistemic virtue. But are there cases in which someone should “stick to their guns”—that is, continue to defend their beliefs despite good evidence to the contrary?

3. (Paper assignment) Explore your favorite conspiracy theory. (A list of them can be found on Wikipedia's page “List of conspiracy theories.”) Write a short paper in which you briefly explain the conspiracy theory, and then, examine the belief

using what you have learned in this class. You might reflect on what motivates someone to believe the theory, how the theory's believers handle evidence against their theory, ways in which that defense of the theory is rational, and ways in which belief in the theory is irrational.

58.

FURTHER READING

Coady, David. 2007. "Are Conspiracy Theorists Irrational?" *Episteme: A Journal of Social Epistemology* 4 (2(Special Issue: Conspiracy Theories)): 193–204.

Kurtis Hagen (2018) "Conspiracy Theories and the Paranoid Style: Do Conspiracy Theories Posit Implausibly Vast and Evil Conspiracies?", *Social Epistemology*, 32:1, pp. 24-40. Hagen argues that the answer to the question is "no" and that each conspiracy theory must be judged on its own merits.

Saslow, Eli (2018). *Rising out of Hatred: The Awakening of a Former White Nationalist*. New York City: Doubleday. (This is the story of Derek Black.)

This is where you can add appendices or other back matter.