

Understanding Knowledge

Michael Huemer

Text copyright ©2022 Michael Huemer, all rights reserved.

Cover photo: Washington Square Park, New York, ©2019 Michael Huemer,
CC BY 4.0. Back cover photo: Barn Owl by Sheribeari, CC BY-ND 2.0.

ISBN: 979-8359316699

Brief Contents

Preface	xv
PART I: GENERAL ISSUES ABOUT KNOWLEDGE & JUSTIFICATION	1
1. Introduction	2
2. What Is Knowledge?	11
3. More Logical and Semantic Debates	38
4. The Structure of Knowledge	56
5. Grounds of Foundational Justification	81
6. Meta-Knowledge	101
7. Taxonomy and Paradigms of Knowledge	116
PART II: SOURCES OF KNOWLEDGE	130
8. Perception & Skepticism	131
9. Theories of Perception	150
10. Pure Reason	168
11. Memory	193
12. Induction	201
13. Testimony	239
PART III: AREAS OF KNOWLEDGE	250
14. Scientific Knowledge	251
15. Moral Knowledge	274
16. Religious Knowledge	303
PART IV: APPLIED EPISTEMOLOGY	323
17. Irrationality	324
18. Critical Thinking and Trust in Experts	341
19. Peer Disagreement	357
Afterword	382
Glossary	384

Preface

Why Read This Book?

This is an introduction to the field of epistemology, which studies philosophical questions about knowledge, rational belief, and stuff like that. There are a fair number of other introductory epistemology books, but this one is better than all of them, so you should buy this one. Here's what is great about this book:

- i. The writing.* It is written in a clear, simple style. It should be easier to read and won't put you to sleep as fast as other epistemology books. (Admittedly, I am less skilled than most epistemologists at curing insomnia.)
- ii. The subject matter.* Epistemology is super-important, because it reflects on what makes beliefs justified and how we can know things in general, which is of central import for all other inquiry, whether in philosophy or in other fields. It's also fun to think about how you know you're not a brain in a vat.
- iii. The price.* Most epistemology books, from traditional publishers, are \$30–\$60 (even for e-books!). The ones that are intended as textbooks are especially expensive. By self-publishing this one, I can keep the price reasonable. I can also write in a more friendly and less turgid style, which traditional publishers wouldn't like.
- iv. The author.* I'm smart, I know a lot, and I'm not confused—which means you can probably learn some interesting things from this book without running into too many confusing nonsense passages.

About the Author

I've heard that it's rude to brag about oneself. On the other hand, you might legitimately want to know who I am before deciding whether to buy a textbook by me. So here is some information to show why I'm a good person to write this book:

I got my BA in philosophy from UC Berkeley. I got my PhD in philosophy from Rutgers University, which was then the #3 philosophy department in the U.S. (they later moved up to #2).¹ I am a tenured full professor at the University of Colorado at Boulder, where I have taught philosophy for over 20 years. As of this writing, I have published more than 70 academic articles in various fields, including more than 30 in epistemology. I've published articles in all 5 of the top 5 philosophy journals. (In philosophy, by the way, the good journals reject 90–95% of submissions.)

I have written eight and a half books (one is co-authored) before this one and edited a ninth, including two epistemology books. See the Afterword for abstracts.

My Approach in Writing This

That's enough about me. Now here are some comments about my approach in writing this:

1. I have included chapters on the most important and interesting topics in epistemology. Some topics are included because students like them (like brain-in-a-vat skepticism); others are included because I think that it's particularly important for your intellectual health to avoid confusion about them.
2. I give a basic presentation of each issue, including what I consider the most important and interesting arguments that can be explained reasonably briefly. (In each case, there are many more complicated and nuanced views and arguments to be found in the academic literature.) When you read these arguments, don't just memorize them and move on, as students sometimes do. Spend some time thinking about why you do or don't agree with them.
3. All of these are issues that people disagree about. In each case, my presentation aims to be (and I think is in fact) *objective*, but not *neutral*. That is:
 - a. I give each view a fair hearing, presenting its case as strongly as I can (given space constraints), in terms that I think are faithful to its proponents' intellectual motivations. I do not select evidence, distort people's words, or use any other tricks to try to skew the assessment of any of the philosophical theories.
 - b. I do *not*, however, promise a *neutral* presentation—one that just reports other people's ideas without evaluation (which I would find incredibly

¹ See the Philosophical Gourmet Report, <http://www.philosophicalgourmet.com/>. This is the most widely used set of rankings in philosophy.

boring). I am going to tell you what I think, and I am going to defend it with logical arguments that try to show you why that view is right.

If you don't like that, this isn't the book for you. Go get another book, like maybe my anthology.²

Vocabulary Words

Periodically in the text, you'll see an expression in bold, such as **epistemology**. This indicates an important philosophical vocabulary word, which will appear in the glossary at the back of the book.

Recycled Philosophy?

Some of the text in this book has appeared before, in my general introduction to philosophy, *Knowledge, Reality, and Value*. I'm totally reusing the epistemology chapters from that book, with some amendments and modifications. There is, of course, much more epistemology in this book.

Acknowledgements

I would like to thank Ross Levatter, Ari Armstrong, and Jun Lin Zeng for their helpful comments on the manuscript, which helped to correct numerous mistakes and shortcomings. I'd also like to thank Iskra Fileva for general awesomeness, God for not being a deceiver, and the evil genius for not existing. Naturally, none of these beings are to blame for any errors that remain. Any such errors are most likely *your* fault. Yes, you, the reader. Because you're dreaming, and you dreamed up this book with errors in it! Why did you do that? Next time, dream a perfect book.

² *Epistemology: Contemporary Readings* (2000).

1. Introduction

In this chapter, I'm going to tell you what epistemology is and why you should study it. I'll also give you some conceptual and terminological background.

1.1. What Is Epistemology?

Epistemology, a.k.a. “the theory of knowledge”, studies the nature of knowledge, whether and how we know what we think we know, whether and how our beliefs are justified, and stuff like that. (The word comes from the Greek root *episteme*, meaning “knowledge”.)

Examples: What is the definition of “know”? How do we know that we can trust the five senses? How do we know what is right and wrong? Are all beliefs justified by observation, or are some things justified independent of observation? We'll address those questions (and so much more!) in this book.

We'll talk about the meaning of “know” in the next chapter. But right now, I want to clarify that when epistemologists talk about knowledge, we are generally talking about *propositional knowledge*, which is the state of knowing something to be the case, or knowing a fact—for example, knowing that it is raining, or knowing that 257 is prime.

There are other uses of “know”—e.g., you could be said to know a person or a place (as in, “I know Spongebob” or “I don't know Westeros very well”), or to know how to do something (as in, “I know how to ride a manatee”). Epistemologists generally are *not* talking about those things when we ask about the nature of “knowledge”.

1.2. Why Is Epistemology the King of All Fields?

1.2.1. Moore's Paradox

I'm going to start with a little puzzle that G.E. Moore famously discussed.

Interlude: G.E. Moore

George Edward Moore (known as “G.E. Moore” due to Britishness) was a twentieth-century British philosopher. He is known for his disarmingly simple

appeals to common sense (like the time he tried to refute external-world skepticism by showing people his hands), his staunch defense of ethical intuitionism, his tediously long discussions of word meanings, his repetitive belaboring of simple points, and his general opposition to philosophical skepticism. Also, he's known for Moore's Paradox, to be explained presently. Moore's Paradox was so named by Wittgenstein, who thought that the discovery of this "paradox" was Moore's greatest contribution to philosophy.

Imagine that you ask me what the weather is like outside, and I reply:

It is raining, but I don't think that it is raining.

There would be something badly wrong with my answer, right? It seems like a nonsensical thing to say, which is why you've probably never actually heard anyone say something like that. But what exactly is wrong with the statement?

In particular, would my statement be *contradictory*? Many people, on first hearing the sentence, answer "Yes". But they are wrong. A contradictory sentence is one that would not be true in any possible circumstance. Yet my answer above *could* be true. There is a possible circumstance in which it is raining outside and at the same time I fail to believe that it is.

You can perhaps see the point more clearly by noticing that there would be nothing wrong with *someone else* saying, "It is raining, but Mike Huemer doesn't think that it is." That person could even be completely correct in saying that. But the person who says that is asserting the same proposition that I assert when *I* say, "It is raining but I don't think that it is." The same proposition cannot be both contradictory and non-contradictory. So it must simply be non-contradictory, regardless of who asserts it.

Sentences like "*p* but I don't believe that *p*" and "*p* but I believe that not-*p*" are called "Moore-paradoxical sentences".³ So here's the puzzle: What exactly is wrong with Moore-paradoxical sentences, and why do they *sound* contradictory even though they are not?⁴

Here is G.E. Moore's answer to the puzzle: Whenever you make an assertion, you are *implying* that you know the proposition that you assert. If you don't know it, you shouldn't be saying it. (*Note:* You're not *saying* that you know the thing in question; you're merely *implying* it. Sort of like how if I say, "Have you stopped harassing manatees?", I haven't actually said that you have harassed manatees in the past, but I definitely implied it.) So if you say that it is raining, you are implying that you know that it is raining. However, for you

³ We also include as "Moore-paradoxical" such statements as "*p* but I don't know whether *p*", "*p* but I have no reason to think that", etc.

⁴ I think it is this puzzle (rather than the Moore-paradoxical sentence itself) that Wittgenstein intended to name "Moore's Paradox".

to *know* that it is raining, you have to at least *think* that it is raining. Therefore, when you say, “It is raining but I do not think that it is raining”, the second half of the sentence contradicts what the first half implies. The “it is raining” part implies that you know it’s raining, while the “I do not think that it is raining” part entails that you do *not* know it’s raining.

Questions you might raise:

1. Is it really true that asserting p always implies that one *knows* that p ? What if you think p is probably true, but you’re not sure; are you then barred from saying anything?

Answer: No; what you do in that case is make a weaker statement, such as, “I *think* that p , but I’m not sure” or “ p is *probably* true” or “ p *seems to me* to be the case”. And notice that all of these are actually asserting things that (plausibly) you *do* know: You may not know that p is true, but you can at least know that you *think* that p , or that p is *probably* true, or that p *seems to you* to be true. So all this is consistent with Moore’s claim.

2. Why say that asserting p implies that you *know* that p , rather than just that asserting p implies that you *believe* that p ? The latter would suffice to solve the original puzzle.

Answer: Because sentences like “It’s raining but I don’t know whether it’s raining” sound nonsensical in the same way as “It’s raining but I don’t believe that it’s raining.”

Let’s say G.E. Moore is right.⁵ Great, we explained a minor puzzle about a weird sentence that no one ever uses. So what? Well, the puzzle is interesting because the solution includes a striking general thesis about statements: *All statements are implied knowledge claims.*

I put that in italics because it’s interesting and I’m making a big deal about it. Any time anyone says anything—whether in philosophy, or science, or ordinary life, or religion, or any other intellectual context—they are implicitly claiming to have knowledge of that whereof they speak. Epistemology is interesting because it studies the nature of this thing, knowledge, that everyone is constantly (implicitly) laying claim to. In that sense, the field of epistemology studies the conditions for *any* claim in any field to be apt. That makes epistemology of sweeping, fundamental importance for all intellectual inquiry.

1.2.2. Epistemological Problems

Here’s another thing about why you should study epistemology. You may not be in a position to know this yet, but I’m just going to tell you this: There are

⁵ For more on the solution to the puzzle, see my paper, “Moore’s Paradox and the Norm of Belief” in *Themes from G.E. Moore* (2007).

a lot of very hard problems in epistemology. These problems call into question a huge portion of our putative knowledge. It is easy to construct arguments for the conclusion that we know none or almost none of the things we normally take ourselves to know. And these aren't just silly fallacies, but arguments whose premises are almost certainly going to seem obvious to you, and whose conclusions are almost certainly going to seem to you to follow from those premises. Trained philosophers have trouble rebutting these arguments, and there is in fact no consensus on what's wrong with them. And what this means is that most of us have some beliefs or intuitions about the nature of knowledge that are *inconsistent* with a huge portion of our normal judgments about what we know.

That is a very interesting and intellectually troubling situation. It suggests that something has gone very wrong in most people's understanding of knowledge—which, given the general importance of knowledge (as just discussed above), has the potential to mess up our thinking about all sorts of other matters.

My claim is not that in general, one must first have sound epistemological views in order to think clearly and cogently in other areas. My claim is that having *unsound* epistemological views can easily screw up your thinking in other areas—and that a great many people are in fact attracted to unsound epistemological views, which they have not carefully examined. That's why you should study epistemology.

1.3. Conceptual Background

Here are some notes about some important concepts that you should already know about if you've already taken some philosophy courses. But if you haven't, then you probably don't know them.

1.3.1. Propositions

Propositions are things that can be true or false—but wait, I need to distinguish three sorts of things that can be true or false.

- i. Sentences.* Sentences are sequences of words like what you're looking at right now. Not *all* sentences can be true or false; e.g., questions or commands cannot be. Only assertive sentences, or proposition-expressing sentences, can be true or false. For instance, "It is raining" is true or false; "Is it raining?" and "Make it rain!" are not.
- ii. Beliefs.* Beliefs are a kind of mental state, a state of thinking something to be the case. They are typically expressed using assertive sentences. They need not actually be expressed, though; you could just think silently to

yourself that it is raining. The thought must be either true or false. This contrasts with, e.g., emotions, desires, or sensations, which are neither true nor false.

- iii. *Propositions.* **Propositions** are the sort of things that beliefs and statements are *about*. When you have a belief, there is something that you believe to be the case; when you make an assertion, there is something you are asserting to be the case. That thing is a “proposition”. Propositions are sometimes thought of as ways the world could be (possible states of affairs), or ranges of possibilities.

A proposition should not be confused with a belief, since the proposition is the thing *that* one believes, the thing one’s mental state is about, not the mental state itself. When you believe a proposition, that proposition is called the **content** of your belief. By the way, people can have different attitudes to the same proposition: One person may *believe* that we will colonize Mars, while another merely *hopes* that we will, a third *doubts* that we will, a fourth *is glad* that we will, and so on.

A proposition also should not be confused with a *sentence* or phrase in a particular language. The proposition is not the *phrase* “that we will colonize Mars”; it is the thing that that phrase *refers to*. (Compare: The Eiffel Tower is not to be confused with the *expression* “the Eiffel Tower”; the Tower is *the referent* of that expression.) The sentences “We will colonize Mars” and “Nous allons coloniser Mars” have something in common. (The second one is the French translation of the first.) They are obviously not *the same sentence*, but they do say the same thing—that is, they express *the same proposition*.

1.3.2. *The Forms of Propositions*

Propositions have structures—that is, they have different kinds of components, which can be connected to each other in different ways. The structure is often referred to as the “form” of the proposition.

The simplest kind of proposition has a simple subject-predicate form. That is, there is a thing the proposition is about (the “**subject**”), and there is a way that thing is said to be, or a property that is ascribed to the thing (the “**predicate**”). Example: [Donald is angry]. *Note:* I often use square brackets like that, to refer to propositions.⁶ This proposition is about Donald, and the way he is said to be (the property ascribed to him) is *angry*.⁷ So Donald is the

⁶ More precisely, this is my convention: If you take a sentence that normally asserts a proposition and enclose it in square brackets, the whole expression becomes a singular term denoting the proposition that the original sentence normally asserts.

⁷ All my hypothetical examples are purely fictional, and any resemblance to any actual persons, living or dead, is entirely coincidental.

subject, and the property of being angry is the predicate. Notice that neither of these things by itself is a proposition.

Some propositions are compound, meaning that they have other propositions as components. Example: [If Donald is angry, then he is dangerous]. In this case, there are two simple propositions, [Donald is angry] and [Donald is dangerous], which are combined using an “if-then”. Sentences or propositions like this (using “if ... then”) are known as “**conditionals**”. The “if” part (in this case, [Donald is angry]) is known as the “**antecedent**” of the conditional. The “then” part (in this case, [Donald is dangerous]) is known as the “**consequent**” of the conditional.

Another type of compound proposition is a **conjunction** (an “and” statement), for example, [Donald is angry and dangerous]. The two parts are called the “**conjuncts**”. In this case, the first conjunct is [Donald is angry], and the second conjunct is [Donald is dangerous].

Another type is a **disjunction** (an “or” statement). The two parts are called **disjuncts**. So in the disjunction [Jesus is a liar, a lunatic, or the Lord], there are three disjuncts. The first disjunct is [Jesus is a liar], the second disjunct is [Jesus is a lunatic], and the third disjunct is [Jesus is the Lord].

We also sometimes talk about **negations**, which are propositions that deny another proposition. For instance, [Jesus is not a liar] is a negation; specifically, it is *the negation of* [Jesus is a liar].

1.3.3. Arguments

An **argument** is a series of statements, some of which are supposed to provide reasons for others, where the whole series is meant to justify a particular conclusion. (We also sometimes speak of an argument as the series of *propositions* expressed by such statements.) These are the different parts of an argument:

- i. *Premises*: The **premises** of an argument are the statements that are used to support the other statements. We reason *from* the premises. Premises are usually chosen to be things that are widely accepted or would seem obvious to most people. (If not, then you may need further arguments to support the premises.)
- ii. *Conclusion*: The **conclusion** of an argument is the statement that the argument is meant to justify. We reason *to* the conclusion. The conclusion is usually something that is initially controversial or non-obvious. (Otherwise, we would not need the argument.)
- iii. *Intermediary steps*: Sometimes an argument has intermediary steps. These are steps in between the premises and the conclusion that help you to see how the premises support the conclusion.

Now here are some characteristics an argument can have:

- i. *Valid or invalid*: An argument is said to be **valid** (or “deductively valid” or “logically valid”) when the premises entail the conclusion; that is, *it would be impossible (in the sense of contradictory) for all the premises to be true and the conclusion to be false.*

Note: This is not the ordinary English usage of “valid”; this is a special, technical usage among philosophers. Virtually all philosophers use the word this way, so you have to learn it. In this sense, “validity” does not require the premises of the argument to be correct, or reasonable, or even consistent. The only thing that is required is that it not be possible that the premises all be true *and* the conclusion be false.

Example: “Socrates is a fish. All fish live on Mars. Therefore, Socrates lives on Mars.” That’s valid, because it could not be that Socrates is a fish, and all fish live on Mars, *and* that Socrates doesn’t live on Mars.

- ii. *Sound or unsound*: An argument is said to be **sound** when it is valid (in the sense given above) *and* all of its premises are true. (In this case, of course, the conclusion must also be true—you can see that if you understood the definition of “valid”.) An argument is unsound whenever it is invalid *or* has a false premise. This is also a technical usage, not the ordinary English usage, and again, philosophers take the stated definition perfectly strictly and literally. Example: “The sky is blue. If the sky is blue, then it isn’t green. Therefore, the sky isn’t green.” That’s sound.

- iii. *Circular or non-circular*: We say that an argument is **circular** or **begs the question** when the premises contain the conclusion, or they contain a statement that is so *similar* to the conclusion that you couldn’t believe the premises without already believing the conclusion, or they contain a statement whose justification *depends upon* the justification of the conclusion. Example: “God exists; therefore, God exists.”

Here is a more realistic example: “Everything the Bible says is true, since the Bible is the word of God. And we know the Bible is the word of God because the Bible *says* that the Bible is the word of God.”

By the way, the phrase “beg the question” *does not mean* “to raise the question”! It means “to give a circular argument”. It fills me with a burning rage when I see people abuse the phrase “beg the question”, so don’t do that.⁸

The above categories are used for assessing arguments, especially for discussing what might be wrong with a given argument. If you have a deductive argument, then the argument needs to be valid, sound, and non-circular. If it is invalid, or

⁸ Really? No, not really. I’m exaggerating my annoyance for comedic effect.

has a false premise, or begs the question, then it's a bad argument. (It's also bad if we merely lack justification for believing one of the premises. But we don't have a separate term for that.)

1.4. Some Symbols

The following symbols are commonly used in epistemology (and philosophy generally).

<u>Symbol</u>	<u>Meaning</u>
S	Commonly used to stand for any person or conscious being (“ S ” for “subject”; a “subject” is just a conscious being).
p, q, r, \dots	Stands for some proposition. You can use any letter. Capital letters are also sometimes used. $p, q,$ and r are epistemologists' favorites. Example: We might say, “ S knows that p only if S believes that p .” (A subject knows a proposition only if they believe that proposition.)
$\sim p$	It's not the case that p . Often read “not- p ”.
$(p \vee q)$	Read “ p or q ”. Note that the “or” is normally read inclusively, so this means “ p or q or both.”
$(p \& q)$	p and q .
$(p \rightarrow q)$	If p then q .
\therefore	Read “therefore”. This is often placed before the conclusion of an argument.

Example: Consider the argument, “Either Skeezix or Ted ate the goldfish. Ted didn't do it. Therefore, Skeezix did it.” This can be symbolized like so:

$$\begin{aligned} &(s \vee t) \\ &\sim t \\ &\therefore s \end{aligned}$$

That's a valid inference.

1.5. Conclusion

Epistemology studies philosophical questions about knowledge and justified belief. Knowledge is interesting and important because whenever you assert anything, you are implying that you *know* that thing (this is shown by the absurdity of asserting “ p but I don't know whether p ”). Also, there are lots of very hard puzzles that arise in epistemology, which show that a lot of us must be pretty confused about knowledge. That's a big thing to be confused about.

To understand what follows, you should know these terms: *proposition*, *content* (of a belief), *subject* (of a mental state), *subject* (of a proposition), *predicate*, *conditional*, *antecedent*, *consequent*, *conjunction*, *conjunct*, *disjunction*, *disjunct*, *negation*, *argument*, *premise*, *conclusion*, *valid*, *sound*, *begging the question*.

7. Taxonomy and Paradigms of Knowledge

How can we best classify the different kinds of knowledge? Epistemologists rarely spend time discussing this, but there is a fairly standard taxonomy. I'm going to describe it, criticize it, then propose another possible classification.

7.1. A Traditional Taxonomy

7.1.1. *Inferential vs. Non-inferential Knowledge*

Traditionally, we divide knowledge into inferential and non-inferential knowledge.

Inferential knowledge is knowledge that is justified on the basis of one or more other beliefs. The paradigm would be knowing a mathematical theorem by proving it from a set of axioms. In that case, you have knowledge based on *deductive* inference. There is also knowledge based on *inductive* inference. For instance, say you meet a bunch of rabbits in different circumstances, notice that they are all furry, and this is how you come to know that all rabbits are furry. Your knowledge that all rabbits are furry is inferential because you inferred it from the premise that this particular rabbit is furry, and that other one is furry, and so on. As a result, if your premises were unjustified, then your conclusion would be unjustified. Also, if (enough of) your premises were *false* (whether or not they're justified), then your conclusion (even if true) would not constitute knowledge (this is the lesson of the Gettier examples of §2.4).

Non-inferential knowledge (also called **foundational knowledge**) is knowledge that is *not* justified on the basis of any other beliefs. (Of course, this assumes foundationalism, which is okay since all sensible people are foundationalists.) Good examples of non-inferential knowledge would be the knowledge that you are in pain (when you are) and the knowledge that $1+1=2$. Some cases are controversial, e.g., it's controversial whether knowledge about

things you perceive with the five senses is foundational or inferential (see ch. 9).

7.1.2. *Empirical vs. A Priori Knowledge*

We also traditionally divide knowledge into empirical (or “a posteriori”) knowledge and a priori knowledge. (*Note:* The expression “a priori” is *one word*. There is no word “priori” by itself. There is only “a priori”, which is an *adjective* applying to beliefs, items of knowledge, and ways of justifying a belief. I have to mention this because some students seem to think that there is a noun “priori”.)

Empirical knowledge (sometimes called “a posteriori knowledge”) is knowledge that is justified (directly or indirectly) by observation. That includes observation by the five senses; it is also usually taken to include introspective “observation”, i.e., the direct awareness of your own conscious mental states. Good examples of empirical knowledge would be your knowledge that you’re in pain (when you are), your knowledge that there is a cat on the refrigerator (when you see this), and our knowledge that humans evolved by natural selection.

A priori knowledge is knowledge that is justified *not* on the basis of observation. This could include innate knowledge (Plato believed in this, but few philosophers today do) or knowledge acquired by some non-observational faculty. Good examples would be the knowledge that $2+3=5$, the knowledge that all grandsons are male, and the knowledge that nothing can be entirely red and also entirely blue.

Now here’s another thing I have to clarify because I know some of you readers are already making this mistake: The distinction between “empirical” and “a priori” knowledge is *not* about how you *acquire concepts*. So don’t say that some item of knowledge is empirical because you *acquired the concepts* through observation. For example, don’t say that “all grandsons are male” is empirical because we acquire the concepts “grandson” and “male” through experience. That’s not what we’re talking about. (John Locke and David Hume were obsessed with how people acquired concepts, but contemporary epistemologists hardly ever talk about that.) What we’re talking about is what is the *justification* or *evidence* for a given item of knowledge. “All grandsons are male” is considered to be known a priori because you don’t have to *justify* it by citing observations that you’ve made of grandsons (or of anything else). (You just have to understand the meaning of “grandson”.)

The above distinctions yield four possible kinds of knowledge: (*i*) foundational empirical knowledge, (*ii*) inferential empirical knowledge, (*iii*)

foundational a priori knowledge, and (iv) inferential a priori knowledge. All four of these exist; see examples below.⁶¹

7.1.3. Four Cognitive Faculties

Our “cognitive faculties” are general, knowledge-gathering capacities that we have. There are at least four that epistemologists traditionally recognize: (i) sensory perception (including sight, hearing, taste, touch, and smell), (ii) introspection (the capacity for awareness of your own mental states), (iii) memory, and (iv) reason (a.k.a. “the understanding” or “the intellect”). The last one (reason) is the faculty that enables us to make inferences. Some (sensible and wise) epistemologists also say that this same faculty enables us to recognize certain self-evident truths, such as “the shortest path between any two points is a straight line”, “if A is better than B and B is better than C , then A is better than C ”, and “nothing can be red all over and also blue all over” (see ch. 10). The capacity for recognizing these self-evident truths is sometimes called “**intuition**”, and the mental state you’re in when you see such a self-evident truth is often called “an intuition”. So I could say that I have the intuition that $1+1=2$, or that my intuition tells me $1+1=2$. Please note that most contemporary philosophers consider intuition to be a *part* or *aspect* of the faculty of reason, not an *alternative* to reason.

Perception and introspection both count as types of observation, so the faculties of perception and introspection both produce foundational, empirical knowledge. Reason, when applied to the material provided by perception and introspection, produces inferential, empirical knowledge. Reason also (according to sensible and wise epistemologists) produces some foundational, a priori knowledge through intuition. When we reason from those intuitions, we obtain inferential, a priori knowledge.

7.2. Traditional Paradigms of Knowledge

A “paradigm” of a category is something that is an excellent example of that category. A sparrow is a paradigm bird, the Empire State Building is a paradigm skyscraper, etc. So the paradigms of knowledge would be excellent examples of knowledge. It’s useful to have some paradigms in front of us, partly just to make sure you understand all the above kinds of knowledge, and partly because it’s easier to address philosophical questions about knowledge by thinking about specific examples. On the Wittgensteinian theory of concepts introduced earlier (§2.6), we form the concept of knowledge in the first place by grouping

⁶¹ Again, I’m assuming foundationalism because I’m sensible.

together cases that strike us as similar, and we can identify the contours of the concept by reflecting on particular examples.

So here's a good example: Say you have a ripe tomato in front of you. It's nearby, directly in your line of sight, well-lit, unobstructed, and there is nothing wrong with your eyes or your brain. In this case, you would know by direct observation that there is a red, round thing in front of you. That's a traditional paradigm of *foundational, empirical knowledge*.

For a different kind of example, think about a scientist's knowledge that water is H_2O . The scientist does an experiment in which he sticks a negative and a positive electrode, both connected to a battery, into a sample of water. Bubbles of gas start forming at both electrodes, and the quantity of water starts decreasing. On further testing, the gas coming from the electrodes turns out to be hydrogen and oxygen. Measurements show that the total mass of gas produced equals the quantity of water that was lost. The scientist infers that water is composed of hydrogen and oxygen. (This is a real experiment that can be done. There's also an experiment in which you can burn hydrogen and oxygen and create water.) That's a paradigm of *inferential, empirical knowledge*.

Here's another sort of example that has influenced many thinkers throughout history: You're reading Euclid's *Elements*, when you come upon the proof that the interior angles of any triangle add up to 180° . When you read Euclid's axioms, you see that they are self-evidently true (this is *foundational, a priori knowledge*). When you go through the proof, you acquire *inferential, a priori knowledge* that the angles of a triangle sum to 180° .

Interlude: Euclid's *Elements*

Euclid's *Elements* is among the most influential books in history. In it, the ancient Greek mathematician Euclid presented proofs of many theorems in geometry and number theory that had been discovered as of the time he lived (around 300 B.C.). Many generations of students since then have been taught geometry through Euclid. If you learned geometry in high school, you can thank/course Euclid for most of the content of that class.

Euclid starts out with some definitions (e.g., "an *obtuse* angle is an angle greater than a right angle"), postulates (e.g., "one can draw a straight line between any two points"), and axioms (e.g., "things that are equal to the same thing are equal to each other"). The postulates and axioms are generally self-evident and in no need of justification. (*Note:* The things he calls "postulates" are specific to geometry, whereas the "axioms" are more general truths of mathematics and logic. But this distinction does not matter, and most people just use "axiom" for both.) From the starting definitions, postulates, and axioms, Euclid proceeds to systematically enumerate and prove 468 different theorems.

Many thinkers since Euclid have been deeply intellectually impressed by the beautifully systematic and rigorous edifice of knowledge set forth in the *Elements*. Many view the book as an ideal example of human knowledge, and some have set out to produce similar systems for other subjects. For instance, Spinoza (the famous 17th-century Dutch philosopher) tried to produce a similar system of knowledge of metaphysics, with definitions, axioms, and theorems in the style of Euclid.

The above are the sorts of examples that philosophers often give when thinking about the nature of knowledge and the different kinds of knowledge that we have.

7.3. Some Hard-to-Classify Cases

7.3.1. Recognition

Let me give you a different sort of example from the ones that are usually discussed by epistemologists. Say I'm getting a phone call, and I don't know whom it's from. I pick up, say "Hello", and the voice on the other end says, "Hey." As soon as I hear the voice, I recognize it: It is my friend Sue calling. She doesn't have to *say* who it is or give any other identifying information; I just recognize the voice.

Now suppose you ask me: "*How do you know* that it's Sue?" I would say: I know because it *sounds* like Sue. "And how exactly does Sue sound? What properties does her voice have that distinguishes it from all other voices?" Well, I don't have anything useful to tell you about that. I can't say exactly what properties of the voice pattern I'm responding to, and I certainly can't describe it in a way that would enable someone who hadn't heard it to recognize it. But I'm very familiar with Sue's voice, and so I know it when I hear it. That's really all I can tell you.

Now let's think about how this example fits into the taxonomy of knowledge. My recognitional knowledge that the voice I'm hearing is Sue's is empirical, not a priori, since it depends on my sense of hearing. But is it foundational or inferential?

It doesn't seem to be foundational, since the knowledge that *that's Sue* presumably depends upon my hearing certain properties of the voice, which are more directly perceived. (This is different from, say, seeing that something is red, since you can directly see that something is red; you don't have to perceive the redness *through* perceiving any other properties.) On the traditional view, the immediately observable properties of a sound would be things like pitch, loudness, and timbre, not "being made by Sue"—just as the immediately

visible properties of an object are its color and shape, not things like “being a tomato”. So it seems that this must be an example of inferential, empirical knowledge.

But that doesn’t seem quite right either. I didn’t really *infer* that the voice was Sue’s. It’s not as though I listed the properties of the voiceprint, cross-referenced them with previous samples of Sue’s voice, then said to myself, “This voice has properties *A*, *B*, and *C*; only Sue’s voice in the past has generally had *A*, *B*, and *C*; therefore, this is probably Sue.”

You might want to say that perhaps I *unconsciously* or *implicitly* made this inference, rather than making a conscious, explicit inference. But this is still a bit odd. Normally, if someone is going to make an inference, the person must *believe* the premises of that inference and *see how* those premises support the conclusion. If we want to allow for unconscious or implicit inference, we might say that though one need not explicitly, consciously affirm the premises, the premises should at least be things that one *would* endorse *if* they were made explicit, and that one could then see how they supported the conclusion. In the present case, neither of these things is true. Suppose someone wrote down a precise and accurate description of the voiceprint, and I read and explicitly understood the description. If they asked me, “Hey, are these the properties of Sue’s voice? And are these the properties of the voice you just heard?”, I would have to say, “I don’t know.”

Now, you might say that, since people do in fact recognize voices, obviously there must be *some* mechanism in the brain that identifies the owner of a voice based on some (no doubt very complicated) set of properties of the voice. And I agree; there is some complicated neural processing that goes into voice recognition. But to say that *my brain* does some complicated *neural processing* is not to say that *I* make *an inference*. In order for something to count as an inference that *I* make, I think it must involve premises that I endorse, or at least would endorse if I thought about them, and for which I would see how they supported the conclusion.

My point is that this example doesn’t fit the taxonomy comfortably. It is an ambiguous case, or on the borderline between inferential and non-inferential knowledge.

7.3.2. Categorization

Recognition of particular people or things, as in the above example, is one interesting kind of case. A related type of case is the recognition of a *type* of thing. When we look around us, we not only perceive particular objects and events; we automatically *classify* the things we observe as falling into various categories. We see this object *as a person*, that one *as a table*, and so on.

As in the case of recognizing an individual, this type of recognition occurs automatically, and one usually cannot fully describe its grounds. Check out this famous illustration, the “duck-rabbit” (figure 3).




Figure 3: The duck-rabbit

You can see this picture as either a duck looking to the left or a rabbit looking to the right. You may be able to shift between them at will. Notice that these are two different *experiences* you can have, even though the colors and locations of the lines that you see on the page don't shift and even though your beliefs don't change either. You do, however, have to have some background knowl-

edge and experience; a person who had no concept of a duck could not perceive the duck aspect of the image.

This example illustrates that the human mind automatically classifies the things we observe according to the concepts we possess at the time, and that this classification is reflected in our experience, not simply our *beliefs*.

In this case, the object you're seeing is literally neither a duck nor a rabbit; it is just ink on the page. Hence, you don't know that it's a duck or a rabbit. But now take a case where you see a real duck. You immediately recognize it as such and hence *know* that it is a duck. What kind of knowledge is this?

There are similar issues here as in the previous case. Your classification of the duck obviously depends upon its shape and color properties, so it would be odd to call it foundational knowledge. At the same time, you didn't exactly *infer* that it was a duck, and you probably could not describe the characteristic visible qualities of ducks in sufficient detail to enable someone who hadn't seen one to recognize them and distinguish them from all other kinds of birds. So it would also be odd to call “that's a duck” inferential knowledge.

7.3.3. Judgment

Here is another interesting type of case. Say you're on the jury in a murder trial. Your task is to determine whether the evidence presented by the prosecution proves beyond a reasonable doubt that the defendant committed the crime. A lot of evidence is presented by both sides. Each side also tries to cast doubt on the other side's evidence, for example, suggesting that some witnesses might have a motivation to lie, or might not have seen events perfectly clearly, or might have imperfect memory. In the end, you have to weigh up all the

evidence on both sides in your mind, form some sense of the overall weight of evidence for guilt, and decide whether it passes the “reasonable doubt” threshold. Needless to say, some people might be better at this than others, some might be more impartial than others, and different people might come to different conclusions.

Let’s say that when you do this, you correctly and justifiedly come to the conclusion that the evidence *does* constitute proof beyond a reasonable doubt. Suppose, in fact (as might be the case), that you *know* that it does. What kind of knowledge is this?

Well, you have empirical knowledge of *what the evidence in the trial was*—you got that by observing the trial and remembering it. But *given* that the evidence was what it was, the knowledge that *that constitutes proof beyond a reasonable doubt* looks to be a priori; it is not itself based on some other evidence. (In other words, you know empirically that the evidence was a certain way, but you know a priori that *if* the evidence is that way, then there is proof beyond a reasonable doubt.)

Now, is your knowledge foundational or inferential? Well, the knowledge of [this set of evidence constitutes proof beyond a reasonable doubt] doesn’t look much like the paradigms of foundational, a priori knowledge—it isn’t a self-evident axiom. But it also doesn’t look much like the paradigms of inferential knowledge. You look at the whole set of evidence and try to *weigh it* in your mind. This weighing isn’t exactly going through a series of inferential steps. We *could* say that it’s a one-step inference where the premises include all the details about the evidence and the conclusion is “the evidence constitutes proof beyond a reasonable doubt”, but then we’d have to acknowledge that this does not follow any of the familiar inference rules that philosophers talk about, and that it is a very different type of inference from those usually discussed. Notice that it is neither a formal deductive inference, nor inductive, nor inference to the best explanation.⁶²

7.3.4. *What Makes a Good Taxonomy?*

You might wonder: So what? Nearly all categories have borderline cases. There are things on the borderline between red and orange, between tall and non-tall, between an adult and an adolescent, etc. Usually, we just go ahead and use these categories anyway, without worrying about how to classify every case. If there is some particular purpose for which we need all cases to be definitively classified, then we can always stipulate a more precise definition. For instance,

⁶² It is not inductive, since the conclusion is not a generalization from multiple particular cases. It also is not inference to the best explanation, since there is no premise stating that the conclusion is the best explanation for something.

for voting purposes, we need to definitively classify everyone as either an adult or not an adult; hence, we just stipulate a cutoff of 18 years of age. No problem.

But I want to point out how my examples differ from most borderline cases. Usually, borderline cases are *rare*, and that is important for the usefulness of our conceptual schemes. If you are classifying objects of some kind, you want a classification scheme that most objects fit comfortably, i.e., such that in most cases it is clear which of the categories the object falls into. E.g., most people fall comfortably into one of the categories, {adult, adolescent, child}. That scheme for classifying people would be less useful if most people were on the borderline between adolescent and adult, or between adolescent and child.

But *my* examples of borderline cases above *are not unusual*. An enormous range of human knowledge, perhaps most of it, is like those examples. As we navigate the world, we are *constantly* recognizing particular people and objects, classifying things into categories, and making judgments that weigh up complex sets of evidence. So the traditional taxonomy of knowledge has the shortcoming that an enormous range of actual knowledge doesn't comfortably fit into it. An enormous range of human knowledge fails to resemble the usual paradigms.

This doesn't mean the taxonomy is wrong. A classification scheme can't really be wrong; it can only be more or less useful. The traditional taxonomy is less useful than one might have hoped, because it was designed to fit a certain set of paradigms that do not in fact resemble a large class of actual cases of knowledge.

Operating with this taxonomy and these paradigms might also mislead people. When we try to formulate principles about the nature of knowledge, we tend to be guided by certain paradigms and by the examples that fit our conceptual scheme comfortably. Thus, there may be a tendency to think that all knowledge must have certain characteristics that those examples have. This might explain some of the felt force of arguments for skepticism. For instance, a budding skeptic might start out by thinking about the sort of examples given in section 7.2 above, decide in light of those paradigms that all knowledge must consist of direct observations, certain and self-evident axioms, or things deductively or inductively inferred from them. This skeptic might then notice that, say, the judgment [it is *prima facie* wrong to break your promises] doesn't fit any of those categories, whereupon the skeptic assumes that no one can know [it is *prima facie* wrong to break your promises]. In brief, the skeptic may think, the wrongness of promise-breaking doesn't resemble the theorems of Euclid, so it must not be real knowledge.

I suspect that something like that happens to many philosophy students. So it may be helpful to use a different set of paradigms, one that includes cases of recognitional knowledge, categorization, and judgment.

7.4. Noticing, Judging, and Calculating

Here is another way of classifying ways of knowing: We know things by noticing, judging, or calculating.

Noticing that p is something that typically occurs automatically, as when you see a tomato in front of you (in good viewing conditions) and you simply notice that there is a tomato. You can also notice things in an a priori way, e.g., when you think about it, you might notice that 3 is less than 7. In “noticing”, something immediately seems to you to be the case, without the need to weigh any alternatives. This category includes exercises of our basic perceptual capacities as well as acquired perceptual abilities (like the ability to recognize a specific person’s voice or the ability to perceive something as a duck) and intuition.

Judgment differs from noticing in that judgment involves weighing alternatives. It does not occur automatically but requires effort. One has the experience of *deliberating about* or *deciding* what to believe, rather than simply being presented with a fact. Judgment is usually based on a complex set of information, but one has no algorithm for going from that information to the conclusion. This kind of cognition may be specific to humans (and intelligent aliens, of course). Examples include weighing up the evidence in a trial to decide what verdict is warranted; evaluating the aesthetic merit of a piece of art based on its observable qualities; morally evaluating an action in light of the reasons for and against it and the set of surrounding circumstances; and evaluating the general plausibility of a philosophical theory.

“Calculation” is my term for a certain sort of inference—roughly, it’s the sort of inference that a computer does. (This is a slightly technical use; I don’t just mean doing math, and it’s not limited to deductive inference but can include probabilistic reasoning.) Calculation is reasoning that doesn’t demand exercises of judgment; at each stage, a certain step has to be accepted, with no need to weigh evidence or reasons for and against. Unlike noticing, calculation is volitional and effortful (one must decide to do it), but unlike judgment, there is no sense of *deciding* what to believe, apart from the decision to do the calculation.

By my read, much of the philosophical tradition, especially including the skeptics, has overemphasized calculation, underappreciated judgment, and underappreciated some forms of noticing. Importantly, we often know things—in perfectly ordinary, everyday cases that are not at all controversial—

without being in a position to *describe how* we know them. Sometimes, philosophy students are disturbed and puzzled by the inability to describe how particular things are known, but once you realize how completely normal that situation is, what you should *really* find weird is the cases in which we *can* describe how we know.

Apropos of this, some of the popular “paradigms” of knowledge are actually bizarre examples, in a certain sense. Mathematical proofs and scientific experiments are the *weird* cases, in the sense of being atypical instances of knowing, despite the enormous amount of attention that philosophers have given them. A normal human being can live an entire life without ever doing a single mathematical proof or scientific experiment—as many people in fact did throughout human history. Yet no normal person goes through life without doing such things as recognizing a friend’s voice, or recognizing a familiar type of object, or judging a body of evidence. That is worth remembering when we want to theorize about the nature of knowledge.

7.5. Top-Down Versus Bottom-Up Cognition

When you see the sort of proofs in Euclid’s *Elements*, it’s very natural to have the reaction, “That is *real* knowledge.” It is then natural to suppose that, to truly know any other subject matter, we must construct Euclid-style proofs—or at the very least, that we should strive to come as close to doing this as possible.

But the example of Euclid is highly misleading about how human knowledge works. Euclid’s geometrical system is a *top down* system of knowledge: It starts from a small number of very general, abstract principles which are obvious and certain on their own; the rest of the system is then inferred from them by a series of calculations (in the sense of “calculation” explained in §7.4). And that is of course a genuine example of knowledge. But it is also, as I say, an extremely atypical kind of knowledge. Hardly any human knowledge works like a mathematical system.

Most human knowledge is *bottom up*. That is, one starts from a large number of cognitions about specific cases or particular individuals. When one has enough cases, one can start to see patterns and general rules. One then starts to formulate abstract principles based upon the cases. Once one has these abstract principles, one may then go on to use them to resolve certain difficult cases. Importantly, one can’t skip the first step, that of gathering judgments about cases. If one tries to *start* from the abstract principles, there is almost a 100% chance of going wrong, often disastrously so. Human beings just are not that smart; we can’t formulate the correct general principles of

almost any subject without looking at examples. We're almost always going to commit oversights and confusions.

By the way, I have not adopted this point of view as an initial, a priori axiom. On the contrary, my natural inclinations are more in line with the top-down approach to knowledge-seeking. It is only through reflection on many examples of knowledge over the last few decades that I have eventually come to see this pattern—that most successful knowledge-seeking activities take a bottom-up approach.

That is part of what we learn from the history of natural science. Before modern science, there were many centuries during which the received scientific theories in different areas were things that we now know were absolutely and totally wrong. One of the reasons for this was that people started out by thinking abstractly about what seemed plausible in each area, rather than *looking and seeing* what was happening in particular cases. E.g., people thought that when you had a fever, this was because you had too much blood. This seemingly made sense, since blood is warm and since people look redder (from more blood in the skin) when they are warm than when they are cold. So the remedy for a fever was to drain blood from your body. To modern ears, that sounds stupid, but that's only because you already know something about the correct account, because someone else taught you. If you didn't already know anything, then the medieval account would probably strike you as plausible, as it struck many smart people before modern medicine was developed. Of course, the theory is 100% wrong—a fever has nothing whatever to do with having too much blood, and bleeding patients is *harmful*, not beneficial. What you have to do is gather a lot of data about particular cases before you start formulating theories about what causes or cures diseases.

Here's another illustration. Let's say that you have graduated college with a physics major. You know all the abstract physics theories that are relevant to explaining, say, how airplanes work—Newton's laws of motion, the conservation of energy, even the Navier-Stokes equations of fluid dynamics. Importantly, though, assume that you know *only general, abstract theory*. No one has taught you about any specific airplane designs, nor have you made any detailed observations of particular airplanes, nor have you talked to any airplane designers. Now let's say that, based solely on your knowledge of physics, you decide to design a new airplane. You draw up the design, pay some people to have it built, and then try finding some people to fly your new airplane for the first time. How do you think it would go?

Well, I am not getting in that airplane. It's probably going to crash, if it gets off the ground at all. I say that even though all kinds of airplanes regularly fly safely all over the world today, and even though I gave you by stipulation all the relevant modern *theoretical* knowledge that explains why those airplanes

fly. The design of modern airplanes, cars, bridges, dishwashers, computers, and innumerable other things, is the product of *decades of experience*—of watching particular designs fail in different ways, figuring out what happened, then modifying the design and trying again. The first versions of anything suck. If you don't believe me, talk to any engineer.

Trying to produce a system of knowledge about some subject matter by reasoning from first principles, without relying on judgments about cases, is like trying to design a new airplane based on theoretical physics, without knowing details about any actual airplanes. Your system is almost certainly going to fail in some major way. If it's a philosophical system, though, there's a good chance that you won't recognize that it failed, because there won't be any decisive empirical test as there is in the case of an airplane, so when someone points out the problems with your theory, you can come up with rationalizations to keep holding on to it. That has happened to many, perhaps the vast majority of, philosophers throughout history.

So why did the top-down approach work so well for Euclid?

It didn't. The way that Euclid presents his system does not in fact reflect the way in which all that mathematical knowledge came about. The Greeks did not learn all those theorems by laying down a handful of simple axioms at the start and then systematically deducing their consequences. Human investigation of geometry began in Egypt around 3000 B.C., and for centuries the field consisted of a miscellaneous collection of principles and formulas, often devised for specific practical purposes in surveying, construction, and astronomy. The formulas were largely empirical, based on trial and error. As a result, they are often approximately but not exactly correct. One ancient document, for instance, gives the area of a circle as the square of eight ninths of the diameter (this is off by about 0.6%).⁶³ So that sort of thing went on for *2,700 years* before Euclid showed up in 300 B.C. and systematized what had been discovered. For perspective, the time between the start of geometry and the writing of Euclid's *Elements* is about 400 years longer than the time between Euclid and us. My point here being that even our knowledge of geometry doesn't work the way you would think from reading Euclid's *Elements*.

7.6. Conclusion

The traditional taxonomy divides knowledge into foundational empirical knowledge (which comes from sense perception and introspection), foundational a priori knowledge (which comes from intuition), inferential empirical knowledge (which results from applying reason to observations), and

⁶³ This is the Ahmes papyrus from Egypt circa 1500 B.C.

inferential a priori knowledge (which results from applying reason to intuitions). Traditional paradigms of knowing include things like seeing a red round thing right in front of you, doing a mathematical proof starting from self-evident axioms, and making an inference to the best explanation from a scientific experiment.

The tradition tends to overlook a wide range of ordinary cases, such as recognizing a friend's voice, automatically categorizing a familiar type of object, and making a judgment based on a complex body of information. These cases are hard to classify as they do not resemble the paradigms of either foundational or inferential knowledge. They also show that it is perfectly normal for individuals to know things without being able to describe how they know. The traditional taxonomy and paradigms have the potential to mislead people into thinking that cases that don't resemble the paradigms are weird and suspect.

An alternative set of paradigms would include cases of recognition, categorization, and judgment. An alternative taxonomy would divide ways of knowing into noticing, judging, and calculating. The tradition underemphasizes judgment and certain kinds of noticing.

Most human knowledge is bottom-up rather than top-down. The top-down approach is extremely unreliable in most areas. That is, rather than starting from general abstract principles and deducing implications, we need to start from many cognitions about particular cases before drawing generalizations.

Glossary

Here are definitions of all the important vocabulary words that appear in boldface throughout the text. In parentheses, I include the section of the book where they were introduced.

- A priori knowledge*: Knowledge whose justification does not depend on observations. (§7.1.2) Contrasted with empirical knowledge.
- Access internalism*: The view that knowledge requires justification, and justification is entirely determined by factors that one can be introspectively aware of. (§3.3.1)
- Adaptation*: A feature of an organism that has a genetic basis and that exists because it promoted inclusive fitness in the organism's evolutionary ancestors. (§15.8.1)
- Adjustable parameters*: In a theory: Quantities whose hypothesized values can be adjusted (without giving up the basic theory) to accommodate the available data. (§14.3.3)
- Aesthetic evaluation*: The evaluation of things in terms of their beauty, artistic merit, or other aesthetic qualities. (§15.1.2) Contrasted with epistemic and practical (including moral) evaluation.
- Agrippa's Trilemma*: The trilemma among three possible structures for a series of reasons: It must stretch back infinitely, or go in a circle, or end in something that one has no reason for. Used by the ancient philosopher Agrippa to argue for global skepticism. (§4.1)
- Analytic philosophy*: A style of philosophy that arose in English-speaking countries in the 20th century, emphasizing clarity of logical argumentation. (§2.1) Contrasted with continental philosophy.
- Analytic*: Of a sentence: True by definition; the negation of an analytic sentence is a contradiction. (§10.1.2) Contrasted with synthetic statements.
- Antecedent*: The first clause in a conditional; in [If *A* then *B*], the antecedent is *A*. (§1.3.2)
- Appearance*: The mental state one is in when it seems to one that something is the case. Also called a "seeming". (§5.2.1)

- Argument from cynicism*: A fallacious tendency to accept certain kinds of theory due to the pleasure one derives from being cynical. (§16.4.2)
- Argument*: A series of statements in which one is supposed to be supported by the others. (§1.3.3)
- Attraction/aversion heuristic*: A method of making judgments that relies on one's feelings of attraction or aversion to something to judge how good, bad, right or wrong the thing is. (§15.6.5)
- Attributor factors*: In the theory of contextualism: Features of a speaker's situation that might affect whether the speaker is correct to say that S "knows" something, e.g., how important it is to the speaker that S be right, what alternative possibilities have been mentioned in the conversation, etc. (§3.1.2) Contrasted with subject factors.
- Awareness*: A mental state that non-accidentally roughly-correctly represents something. (§9.2.1)
- Basing condition*: The principle that a belief counts as knowledge only if the belief is based on the factors that provide justification for it. (§9.4.4)
- Bayes' Theorem*: The theorem that, for any b and e , $P(b|e) = P(b) \times P(e|b) / P(e)$. (§12.3.2)
- Bayesian account of confirmation*: The view that e confirms b if and only if $P(b|e) > P(b)$. (§14.1.4)
- Bayesianism*: An approach to epistemology that seeks to explain all cogent non-deductive reasoning in terms of the principles of probability theory, especially Bayes' Theorem. (§12.5.1)
- Beg the question*: To give a circular argument. (§1.3.3)
- BIV*: Abbreviation for "brain in a vat". (§2.5.4)
- Bleen*: The property of being blue if observed before 2100 A.D. and green otherwise. (§12.8.1)
- Brain in a vat*: A brain that is being kept alive in a vat of nutrients and artificially fed electrical stimulation to simulate life in the real world. The brain-in-a-vat hypothesis is the hypothesis that you are a brain in a vat. (§2.5.4)
- Burden of proof principle*: The thesis that those who make positive claims have the burden of providing evidence, and that there is a presumption in favor of negative claims. (§14.3.1)
- Calibrated*: Said of a set of credences wherein, when one has a credence of about $x\%$ in a proposition, that proposition tends to be true about $x\%$ of the time. (§18.4)

- Categorical imperative*: A proposition stating that some action is right or wrong independent of one's goals or desires; an obligation that one must follow regardless of what one wants. (§15.1.2)
- Causal theory of reference*: The view that one can only have an intentional state referring to x if one has had some causal interaction with x , or with things in terms of which x could be described. (§8.3.3)
- Certainty skepticism*: A form of skepticism that says we lack knowledge because our beliefs are not absolutely certain. (§8.2.4) Contrasted with justification skepticism.
- Circular*: Of an argument: Having premises that contain the conclusion or depend for their justification on the conclusion. (§1.3.3)
- Classical interpretation of probability*: An interpretation according to which the probability of an event is the ratio of the number of possible situations in which the event occurs to the total number of possible situations. (§12.3.3)
- Closure principle*: (a) Closure for knowledge: The principle that if one knows p , and p entails q , then one is in a position to know q (or something like that). (b) Closure for justification: The principle that if one has justification for p , and p entails q , then one has justification for q (or something like that). (§2.4, n10; §3.2.1)
- Cognitive faculties*: Faculties by which we acquire (what we usually take to be) knowledge, e.g., sensory observation, memory, reason, introspection, and intuition. (§6.2.1)
- Coherence theory of justification*: The view that justified beliefs are justified because they are supported by a system of beliefs that fits together well, where this is a matter of the beliefs in the system supporting each other, not contradicting each other, explaining each other, etc. (§4.3.1)
- Coherentism*: See coherence theory of justification.
- Compositionality of meaning*: The principle that the meaning of a sentence is determined by the meanings of the component parts and the structure of the sentence. (§10.3.3)
- Conceptual scheme*: A system for mentally grouping things into categories and distinguishing them from other things. (§2.1)
- Conciliatory view*: Any view according to which, when epistemic peers disagree, each should significantly adjust his credence in the direction of the other person's credence. (§19.1)
- Conclusion*: The proposition that an argument is supposed to support. (§1.3.3)
- Conditional*: A sentence/proposition of the form "If A then B ". (§1.3.2)
- Conditionalization*: A way of updating beliefs upon acquiring new evidence, e , in which, for each proposition h , you set your new credence in h to what was

previously your conditional credence in b given e , i.e., you set $P_{\text{new}}(b) = P_{\text{old}}(b|e)$. (§12.5.3)

Confirm: To provide some support for; to render (a proposition) more likely. (§14.1.1)

Confirmation bias: The common tendency to look only for evidence supporting a theory rather than looking for evidence against it. (§17.4.4)

Confirmation holism: The view (associated with W.V.O. Quine) that a statement cannot be confirmed or disconfirmed in isolation; one can only test a belief system as a whole (or: one can only test a statement *given* a background belief system). (§10.4.1)

Conjunct: One of the clauses in a conjunction; in $[A \& B]$, A and B are conjuncts. (§1.3.2)

Conjunction: A sentence/proposition of the form “ A and B ”. (§1.3.2)

Consequent: The last clause in a conditional; in $[If A then B]$, the consequent is B . (§1.3.2)

Content (of a belief): The proposition that one believes. (§1.3.1)

Contextualism: A theory according to which the standards for someone to count as “knowing” a proposition shift depending on the conversational context; i.e., the correctness of knowledge attributions depends not only on subject factors but also on attributor factors. (§3.1.1)

Continental philosophy: A style of philosophy that became popular on the European continent, especially France and Germany, in the 19th-20th centuries. Often less clear and logical than analytic philosophy, with which it is contrasted. (§2.1)

Contingent: Said of a proposition that could have been true and could have been false; that is, it is neither necessary nor impossible. (§8.1)

Control group: In an experiment: The group of things, people, etc., that did not receive the intervention whose effect one is trying to test. (§18.3.1)
Contrasted with the experimental group.

Credence: A person’s degree of confidence in a proposition; subjective probability. (§12.3.3)

Critical thinking philosophy: The view that one should attempt to think through controversial issues for oneself, rather than trusting experts. (§18.1)

Deduction: A form of reasoning in which the premises are supposed to entail the conclusion, such that the premises could not all be true *and* the conclusion fail to be true. (§12.1.1)

Deductivism: The thesis that only deductive reasoning is cogent; entails inductive skepticism. (§14.4.3)

Defeasibility theory: The theory that knowledge is justified, true belief with no (genuine) defeaters. (§2.5.7)

Defeasible: (a) Of a belief: Having only defeasible justification. (b) Of a (source of) justification: In principle subject to defeaters, i.e., capable of being outweighed or neutralized by further information. (§5.2.1)

Defeater (for a proposition, p): (a) In the defeasibility theory of knowledge: A true proposition that, if added to one's beliefs, would result in one's no longer being justified in believing p . (§2.5.7) (b) A proposition that one believes or has justification for believing that gives one reason to doubt p , where p would otherwise be justified. (§4.5.1) *See also* rebutting defeaters; undercutting defeaters.

Descriptive proposition: A proposition that is not evaluative. (§15.1.1)

Direct realism: The view that perception gives us direct awareness and non-inferential knowledge of the external world. (§8.3.5) Contrasted with indirect realism.

Direct awareness: A state of awareness of something that does not depend upon awareness of anything else. (§9.2.2) Contrasted with indirect awareness.

Disjunct: One of the clauses in a disjunction; in $[A \text{ or } B]$, A and B are disjuncts. (§1.3.2)

Disjunction: A sentence/proposition of the form " $A \text{ or } B$ ". (§1.3.2)

Disjunctive conception of experience: The view that there is no mental state in common between normal perceptual experiences and hallucinations, hence "sensory experience" is a disjunctive kind. (§9.2.4)

Disjunctive syllogism: A form of argument in which one takes a disjunction as a premise, rejects all but one disjunct, then concludes that the remaining disjunct must be true. (§4.4.1)

Disjunctivism: *See* disjunctive conception of experience.

Dogmatism: The common tendency to under-adjust one's credences in response to evidence that undermines one's initial beliefs. (§17.4.2)

Doxastic justification: The justification of an actual belief in virtue of its actually being based on something that provides propositional justification. (§2.3.3)

Doxastic voluntarism: The view that we can voluntarily control our beliefs. (§16.1.2)

Empirical knowledge: Knowledge whose justification depends on observations. (§7.1.2) Also called a posteriori knowledge. Contrasted with a priori knowledge.

Empiricism: The view that there is no synthetic, a priori knowledge; the view that all substantive knowledge of the world depends on observation. (§10.2.1) Contrasted with rationalism and Kantianism.

- Entrenched*: Of a predicate: Frequently used in actual people's inductive reasoning. (§12.8.3)
- Epicycle*: In Ptolemaic astronomy: A smaller circle that planets were held to be moving in, where the small circle itself moves in the main orbit around the Earth. Often used as a metaphor for ad hoc complications added to a theory. (§14.3.2)
- Epistemic circularity*: The alleged mistake of using a belief-forming method in coming to the conclusion that that method is good. (§6.3.2)
- Epistemic evaluation*: The evaluation of beliefs (or similar things) in terms of their epistemic justification. (§15.1.2) Contrasted with practical (including moral) and aesthetic evaluation.
- Epistemic evidentialism*: The view that the epistemically justified attitude about any proposition at any given time is solely determined by one's evidence. (§16.1.1)
- Epistemic peers*: People who are about equally well-positioned to evaluate a given issue, in terms, e.g., of their intelligence, relevant knowledge, and time spent thinking about the issue. (§19.1)
- Epistemic probability*: A type of probability that indicates the degree of justification that one has for a given proposition. (§12.3.3)
- Epistemic reason*: A reason for believing or not believing something of the sort that shows the would-be belief to be probable or improbable. (§2.3.2)
- Epistemically justified*: Rational in the sense of being supported by sufficient epistemic reasons, or being sufficiently likely to be correct. (§2.3.2)
- Epistemology*: The study of philosophical questions about knowledge, justified belief, and things like that. (Preface)
- Equal Weight View*: The view that, when epistemic peers disagree, each peer should give about equal weight to the other person's opinion as to his own. (§19.1, §19.2.1)
- Equivalence Condition*: The thesis that if p and q are logically equivalent, then whatever confirms p confirms q . (§14.1.3)
- Ethical intuitionism*: The view that there are objective, evaluative facts which are irreducible and knowable on the basis of ethical intuitions. (§15.7.1)
- Evaluative proposition*: A proposition that positively or negatively evaluates something; a proposition that entails that something is good or bad in some respect. (§15.1.1)
- Evidentialism*: See epistemic evidentialism; moral evidentialism.
- Experimental group*: In an experiment: The group of things, people, etc., that received the intervention whose effects one is trying to test. (§18.3.1) Contrasted with the control group.

Expressivism: The view that moral statements fail to express propositions and instead serve to give commands, express emotions, or express some other type of non-cognitive attitude. Also called “non-cognitivism”. (§15.2.1)

Externalism: The negation of internalism. *See* access internalism; internal state internalism; semantic internalism.

External world: The world outside one’s own mind. (§4.4.1)

External world skepticism: The view that we don’t know, or aren’t justified in believing, any contingent propositions about the external world. (§4.4.1, 8.2)

Fallibilism: The view that knowledge can be uncertain; i.e., one’s knowing p is compatible with $\sim p$ ’s having a non-zero probability. (§15.7.7)

Falsifiable: Of a theory: Testable; capable of being proven false, or at least shown to be likely to be false, *if* it were false. (§8.3.4, §14.2.1)

Fideism: The view that we ought to believe religion based on blind faith, as opposed to evidence. (§16.1.3)

First-order evidence: Discussed especially in cases of disagreement: Evidence directly about the matter in controversy. (§19.1) Contrasted with second-order evidence.

Foundational belief: A belief whose justification does not depend upon reasons, i.e., does not depend upon being supported by other beliefs. (§4.5.1)

Foundational justification: Epistemic justification that does not depend upon reasons. (§4.5.1)

Foundational knowledge: Non-inferential knowledge. (§7.1.1)

Foundationalism: The view that some items of knowledge or justified belief do not require reasons, and that all other knowledge or justified belief is based upon those things. (§4.5.1)

Foundherentism: The view that our belief systems are justified by virtue of a small amount of foundational justification for certain beliefs, combined with coherence relations among the beliefs. (§5.4)

Frege-Geach problem: The problem for expressivists of explaining what moral statements mean when they are embedded within larger statements, in contexts where a proposition-expressing clause is normally required. (§15.2.1)

Fully grounded: Having no false beliefs in its evidential ancestry. Also known as having “no false lemmas”. (§2.5.1)

Generality problem: The problem for reliabilists of specifying what counts as “the method” by which a belief was formed, esp. deciding how specific one’s description of the method should be. (§2.5.2)

Genuine defeaters: In the defeasibility theory of knowledge: The sort of defeaters that prevent you from having knowledge; defeaters that aren't misleading. E.g., the fact that Tom has an identical twin genuinely defeats your belief that you saw Tom steal a book from the library. (§2.5.7)

G.E. Moore shift: See Moorean response.

Gestalt perception: A perception in which one perceives a complex scene or object as a unified whole, rather than a set of individual details. (§19.6.2)

Global skepticism: The view that we don't know, or aren't justified in believing, anything whatsoever. (§4.4.1)

Grue: The property of being green if observed before 2100 A.D. and blue otherwise. (§12.8.1)

Hallucination: A sensory experience that fails to represent any real object. (§9.3.2)

Hume's Law: See is/ought gap.

Idealism: The view that there are no material things existing independent of the mind; there are only minds and "ideas" in the mind. (§9.1) Contrasted with realism.

Illusion: A sensory experience that represents a real object but misrepresents one or more of its characteristics. (§9.3.2)

Inclusive fitness: In evolutionary biology: An organism's tendency to cause more copies of its genes to exist in the next generation. Includes both its tendency to directly reproduce and its tendency to promote the survival and reproduction of kin who are likely to carry the same genes. (§15.8.1)

Indirect awareness: A state of awareness of something that depends upon awareness of something else. (§9.2.2) Contrasted with direct awareness.

Indirect realism: The view that perception gives us direct awareness and non-inferential knowledge only of something dependent on our own minds (such as mental images), but that it also enables us to have *indirect* awareness and *inferential* knowledge of external objects. (§8.3.5) Contrasted with direct realism.

Induction: A type of non-deductive reasoning in which one generalizes from particular cases, i.e., the premises say something about certain objects, and the conclusion extends what was true of those objects to a wider class of objects. (§12.1.1)

Inductive skepticism: The view that inductive reasoning never provides any justification at all for its conclusions. (§12.1.2)

Inference to the best explanation: A type of non-deductive inference in which one infers that some theory is likely true because it provides the best

explanation for some evidence that would be improbable if there were no explanation. (§12.7.1)

Inferential knowledge: Knowledge that is justified on the basis of other beliefs. (§7.1.1)

Infinitism: The view that justified beliefs are justified because they have an infinite series of supporting reasons available. (§4.2.1)

Initial plausibility: The degree to which a proposition seems correct, prior to argument. (§4.4.3)

Instrumental reasons: A species of practical reasons (reasons for action) that are based on the agent's actual desires or goals. (§15.1.2)

Intellectual appearances: Mental states in which something seems true to one on the basis of intellectual reflection, as opposed to observation. (§15.7.2)

Intentionality: The property of representing something; being *of* or *about* something. (§8.3.3)

Internalism: See access internalism; internal-state internalism; semantic internalism.

Internal-state internalism: The view that knowledge requires justification, and justification is entirely determined by one's internal mental states. (§3.3.1)

Intuition: A mental state in which something seems correct upon direct, intellectual reflection, as opposed to observation or reasoning. (§7.1.3)

Intuitionism: (a) In epistemology: The view that our synthetic, a priori knowledge is not innate but is acquired by a non-observational faculty which gives us intellectual insights. (§10.5.1) (b) In metaethics: ethical intuitionism.

Is/ought gap: The thesis that one cannot validly infer an evaluative conclusion from descriptive premises. (§15.3.1)

JTB: Short for "justified, true belief", which is the traditional definition of knowledge. (§2.4)

Justification skepticism: A form of skepticism that says we lack knowledge because our beliefs are not even justified. (§8.2.4) Contrasted with certainty skepticism.

Kantianism: The philosophical views of Immanuel Kant, which prominently include that there is synthetic, a priori knowledge; that this knowledge is explained by the fact that our mind imposes a certain structure on everything that we are aware of; and that as a result, we can only know things as they appear, not things in themselves. (§10.6.1) Contrasted with rationalism and empiricism.

KK Thesis: The thesis that if one knows that *p*, then one knows that one knows that *p*. (§6.1)

Knowledge by acquaintance: Being in a position to refer to something in virtue of having direct awareness of that thing. (§5.1.1) Contrasted with knowledge by description.

Knowledge by description: Being in a position to refer to something in virtue of understanding a description that uniquely applies to that thing. (§5.1.1) Contrasted with knowledge by acquaintance.

Kolmogorov axioms: The four basic axioms of probability theory, namely, (i) that the probability of anything is greater than or equal to 0, (ii) that the probability of any tautology is 1, (iii) that the probability of $(a \vee b) = P(a) + P(b)$ if a and b are mutually exclusive, and (iv) that the probability of $(a \& b) =$ the probability of a times the probability of b given a . (§12.3.2)

Law of large numbers: The principle of probability theory that if an outcome O has a probability p of occurring in circumstance C , then if C is repeated many times, the frequency with which O happens will tend to approximate p (with increasing accuracy and certainty as the number of trials increases). (§12.3.2)

Level confusion: A confusion between knowing p and knowing that one knows p , or between justifiedly believing p and justifiedly believing that one justifiedly believes p , or something like that. (§4.5.4)

Likelihood (of a hypothesis, b): The probability of e given b , $P(e|b)$. (Yes, this is a weird name for it.) (§12.5.4)

Likelihood ratio (of a hypothesis, b): The ratio of the likelihood of b to that of $\sim b$, i.e., $P(e|b)/P(e|\sim b)$. (§12.5.4)

Logical positivism: The view that there is no synthetic, a priori knowledge and that the meaning of a sentence is given by its verification conditions; the conjunction of empiricism and verificationism. (§10.3.1)

Logical probability: Probability construed as a logical property of a proposition or a logical relation between propositions; often glossed as the proportion of possible worlds in which a given proposition is true. (§12.3.3)

Matters of fact: In David Hume's philosophy: Propositions that are made true by mind-independent facts and therefore can only be known empirically; propositions whose expression is synthetic and that can be known only empirically. (§12.1.2) Contrasted with relations of ideas.

Meta-belief: A belief that is about one or more beliefs. (§4.5.4)

Meta-justification (for a belief): An argument that shows that the characteristic that (allegedly) makes a given belief justified is one that in general renders beliefs highly likely to be true. (§4.5.4)

Meta-knowledge: Knowledge about one's own knowledge. (Ch. 6)

Misleading defeaters: In the defeasibility theory of knowledge: The sort of defeaters that don't prevent you from having knowledge. E.g., in the case where Tom does not actually have a twin, the fact that Tom's mother *says* that Tom has an identical twin is a misleading defeater for your belief that you saw Tom steal a book from the library. There is no consensus on exactly what differentiates genuine and misleading defeaters. (§2.5.7)

Moorean response: A type of response to skepticism in which one argues that one of the skeptic's premises must be false since the skeptic's conclusion is so implausible. (§4.4.3)

Moral evaluation: Evaluation of an action in terms of whether it is morally right or wrong. (§15.1.2) Contrasted with epistemic and aesthetic evaluation.

Moral evidentialism: The view that it is morally wrong to hold epistemically unjustified attitudes. (§16.1.2)

Moral sense theory: The view that we have a faculty specifically designed for apprehending moral truths. (§15.8.3)

Naïve Comprehension Axiom: In set theory, the principle that for any meaningful predicate, there is a set containing all and only the things satisfying that predicate. Now known to be false since it generates Russell's Paradox. (§5.1.3)

Naïve realism: See direct realism.

Nativism: A form of rationalism that holds that some synthetic, a priori knowledge is innate. (§10.5.1) Contrasted with intuitionism.

Naturalistic Fallacy: (a) The mistake of trying to deduce an evaluative conclusion from purely descriptive premises. (b) The mistake of trying to define an evaluative term using only descriptive terms. (§§15.3.1-2)

Nearby possible worlds: Possible worlds (or ways that the world could have been) that are similar to the actual world (the way the world in fact is). (§2.5.4)

Necessary: Said of a proposition that could not have been false; the negation of it is impossible. (§8.1)

Negation: A sentence/proposition of the form "It's not the case that A ". (§1.3.2)

Negative claim: A claim that something does not exist or does not have some property. (§14.3.1) Contrasted with positive claims.

Nicod's Criterion: The three-part thesis that (i) observation of an A that is B confirms [All A 's are B], (ii) observation of an A that is non- B disconfirms [All A 's are B], and (iii) observation of a non- A is irrelevant to [All A 's are B]. (§14.1.3)

Nililism: In metaethics: The view that nothing has any moral properties and thus all (positive) moral statements are false. (§15.2.2)

Non-cognitivism: See expressivism.

Non-deductive reasoning: A form of reasoning in which the premises are supposed to render the conclusion more likely but not to entail the conclusion. (§12.1.1)

Non-epistemic belief preferences: Desires to believe or not believe certain propositions for reasons independent of their truth or degree of epistemic justification. (§17.3.4)

Non-inferential knowledge: Knowledge that does not depend on other beliefs for its justification. (§7.1.1)

Null hypothesis: In a statistical study: The hypothesis that the variables being studied have no causal connection, and thus that any observed correlations are due to chance. (§18.3.1)

Objective Bayesianism: A form of Bayesianism that holds that there are significant constraints on rational initial credences beyond the Kolmogorov Axioms. (§12.5.1) Contrasted with subjective Bayesianism.

Observational study: A type of study in which the scientists do not have control over who does and does not receive the treatment whose effects are being evaluated. (§18.3.1)

Occam's Razor: The thesis that, other things being equal, the simplest explanation of some evidence is most likely to be correct. (§14.3.1)

Overconfidence: The common tendency to hold credences that are more extreme (closer to 0 or 1) than the evidence justifies. (§17.4.2)

Peer disagreement: Disagreement between epistemic peers. (§19.1)

Perceptual experience: See sensory experience.

Perspectival variation: The phenomenon whereby sensory appearances vary depending on the observer's relationship (esp. spatial relationship) to the object, as opposed to the object's intrinsic properties. (§9.3.1)

p-hacking: The practice of performing multiple statistical tests using different pairs of variables in order to find one that passes a test for statistical significance. (§18.3.3)

Phenomenal Conservatism: The theory that if it seems to one that p , and one has no reason for doubting that appearance, then one thereby has at least some justification for believing p . (§5.2.1)

Platonism: The view that universals exist necessarily. (§10.5.4)

Positive claim: A claim that something exists or has some property. (§14.3.1) Contrasted with negative claims.

Practical evaluation: Evaluation based on reasons for and against actions. (§15.1.2) Contrasted with epistemic and aesthetic evaluation.

Practical rationality: The property of being best supported by an agent's reasons for action. (§15.1.2)

Practical reason: A reason for performing or not performing an action. Includes prudential, instrumental, and moral reasons. (§2.3.2) Contrasted with epistemic reasons.

Predicate (of a proposition): The thing that a proposition attributes to its subject. E.g., in the proposition [John is bald], the predicate is baldness. Also used for a linguistic expression that is used to ascribe a predicate, e.g., the phrase "is bald". (§1.3.2)

Premise circularity: The fallacy that one commits when one infers a conclusion from itself. (§6.3.2)

Premises: The starting points of an argument; the statements that are used to support the rest. (§1.3.3)

Prima facie justification: Justification that is foundational and defeasible. (§5.2.1)

Primary qualities: Qualities of external objects, such as shape, size, mass, and number, that, according to some philosophers, are objective, in contrast to the "secondary qualities". (§10.6.3)

Principle of Charity: The idea that one should generally try to interpret other people as having mostly true, or at least reasonable, beliefs. (§13.3.3)

Principle of Indifference: The thesis that if there are no reasons for favoring either of two alternatives over the other, then they have equal epistemic probabilities. (§12.6.1)

Prior probability ("prior" for short): The probability that a proposition has prior to gathering evidence about it. (§12.5.4)

Propensity: A type of probability that indicates the strength of a causal tendency of some circumstance to produce some outcome. Used in indeterministic interpretations of quantum mechanics. (§12.3.3)

Proper function analysis: The theory that knowledge is true belief formed by a properly functioning, reliable, truth-directed faculty operating in the conditions it was designed for. (§2.5.3)

Proportional syllogism: A type of non-deductive reasoning in which one reasons from the fact that some event occurs in $x\%$ of cases to the conclusion that the event will occur in a given case (with $x\%$ confidence). E.g., if 90% of A 's are B , and you know only that y is an A , you might infer (with 90% confidence) that y is B . (§12.4.1)

Propositional content: The proposition that a belief or other mental state is about. (§5.1.2)

Propositional justification: Epistemic justification that one has available for a given proposition. (§2.3.3)

- Propositions*: The sort of things that can be true or false, that can be asserted, believed, doubted, etc. (§1.3.1)
- Prudential reasons*: A species of practical reasons (reasons for action) that are based on self-interest. (§15.1.2)
- Publication bias*: The phenomenon whereby scientific studies are more likely to be published if they report statistically significant results. (§18.3.3)
- p-value*: In a statistical study: The probability of obtaining a result of the size observed, given the null hypothesis. (§18.3.1)
- Randomized, controlled trial*: A kind of experiment in which there is a control group and an experimental group, and the people/things being tested are randomly assigned to either the control or the experimental group. (§18.3.1)
- Rational disagreement*: A disagreement in which both parties are rational in sticking to their views. (§19.3.1)
- Rationalism*: The view that there is synthetic, a priori knowledge, or: that there is substantive knowledge of the world that does not depend on observation for its justification. (§10.5.1) Contrasted with empiricism and Kantianism.
- Real World Hypothesis*: The “hypothesis” that one is perceiving the world normally. (§8.3.4) Contrasted with the brain-in-a-vat hypothesis.
- Realism*: (a) In general: The view that some class of phenomena of philosophical interest exists objectively and is knowable. (b) In philosophy of perception: The view that external objects exist objectively and can be known through perception. (§8.3.5) *See also*: direct realism, indirect realism.
- Rebutting defeater*: A defeater for a proposition, p , which works by supporting $\sim p$. (§4.5.1)
- Regress argument*: (a) For skepticism: The argument that we can’t know (or have justification for believing) anything because we cannot base knowledge/justified belief on something that we have no reason for, nor on circular reasoning, nor on an infinite regress of reasons. *See also*: Agrippa’s trilemma. (§4.4.1) (b) For foundationalism: The argument that our knowledge (or justified beliefs) must be based on things that are self-evident, or that we do not need reasons for, since we can’t have an infinite regress and we can’t rely on circular reasoning. (§4.5.2)
- Relations of ideas*: In David Hume’s philosophy: Propositions that are made true by the meanings of words or the relationships of concepts and therefore can be known a priori, e.g., [All squares have four sides]; propositions whose expression is analytic and that can be known a priori. (§12.1.2) Contrasted with matters of fact.

Relevant alternatives (to a proposition): Alternatives that need to be ruled out in order for one to know a given proposition; commonly understood to be the alternatives that could fairly easily have been realized. The Relevant Alternatives Analysis holds that knowledge is true belief supported by evidence that rules out all the relevant alternatives, where this is a proper subset of the logically possible alternatives. (§2.5.6)

Reliabilism: Roughly, the view that knowledge is true belief formed by a reliable method, i.e., a method that if used many times would tend to generate true beliefs a great majority of the time. (§2.5.2)

Representationalism: See indirect realism.

Right Reasons View: The view that the rational response to a case of peer disagreement is for the party who *in fact* made a mistake to change his view and the party with the correct view to remain steadfast. (§19.1, §19.4.1)

Rule circularity: The alleged mistake of using a form of inference in reasoning to the conclusion that that form of inference is good. (§6.3.2)

Russell's Paradox: The paradox that results from the notion of the set of all sets that don't contain themselves: This set must contain itself if and only if it doesn't contain itself. (§5.1.3)

Safety: Where S believes that p , this belief is "safe" if and only if p is true in all the nearby worlds in which S believes p ; that is, S would not easily have been wrong in believing p . The safety condition on knowledge says that S knows p only if S 's belief that p is safe. (§2.5.5)

Scientific realism: The view that science reveals to us important truths about mind-independent reality. (§14.4)

Secondary qualities: Observable qualities of external objects (such as color, taste, smell, sound, and temperature) that, according to some philosophers, are not entirely objective but in some way depend upon our sensory faculties; sometimes thought to be illusory, or to be dispositions to produce certain sensations in us, or to be constituted by collections of primary qualities that give rise to said dispositions. (§10.6.3) Contrasted with primary qualities.

Second-order evidence: Discussed especially in cases of disagreement: Evidence about who is more reliable, which assessment of the evidence is more likely to be correct, and the like, as opposed to evidence directly about the matter in controversy. (§19.1) Contrasted with first-order evidence.

Second-order knowledge: See meta-knowledge.

Seeming: The mental state one is in when it seems to one that something is the case. Also called an "appearance". (§5.2.1)

- Semantic externalism*: The view that what one's mental states represent depends in part on factors external to one's mind. (§8.3.3) Contrasted with semantic internalism.
- Semantic internalism*: The view that what one's mental states represent depends only on purely internal features of one's mental states. (§8.3.3) Contrasted with semantic externalism.
- Sensitivity*: Where S believes that p , this belief is "sensitive" if and only if, if p were false, S would not believe p . The sensitivity condition on knowledge says that you know p only if your belief that p is sensitive. (§2.5.4)
- Sensory experience*: The purely internal mental state that one has during perception or hallucination; a state of seemingly perceiving something. (§9.2.1)
- Skeptical scenario*: A scenario in which everything would appear to you as it presently does but most of your beliefs would be mistaken. (§8.2)
- Skepticism*: A philosophical theory according to which we don't know many of the things we normally think we know. (§4.4.1, §9.1) *See also* external world skepticism, global skepticism, certainty skepticism, justification skepticism.
- Sound*: Of an argument: Valid *and* having all true premises. (§1.3.3)
- Statistically significant*: Of an experimental result: Having a low p -value, typically either <0.05 or <0.01 . (§18.3.1)
- Steadfast View*: The view that it is epistemically permissible to stick with your original view when you learn that an epistemic peer disagrees with you. (§19.1, §19.3.1)
- Subject (of a mental state)*: The being who has a given mental state. (§1.4)
- Subject (of a proposition)*: The thing that a proposition is about. E.g., in the proposition [John is bald], the subject is John. (§1.3.2)
- Subject factors*: In the contextualist view of knowledge: Features of a potential knower's situation that might affect whether that person counts as "knowing", including the truth of their belief, their justification, etc. (§3.1.2) Contrasted with attributor factors.
- Subjective Bayesianism*: A form of Bayesianism that holds that any set of initial credences that satisfies the Kolmogorov Axioms (with at most minor additions) is rational. (§12.5.1) Contrasted with objective Bayesianism.
- Subjective probability*: A type of probability that indicates a person's degree of confidence in a proposition; credence. (§12.3.3)
- Synthetic*: Of a sentence: Not true by definition; the negation of a synthetic sentence is consistent. (§10.1.2) Contrasted with analytic statements.
- Testimony*: The act of one person's telling someone else something. (Ch. 13)

Total Evidence View: The view that the rational response to peer disagreement is the outcome of a weighing of both first-order and second-order evidence. (§19.1, §19.5.1)

Tracking analysis: The theory that knowledge is true belief formed in such a way that if the proposition were false, one would not have believed it, and if it were true, one would have believed it. (§2.5.4)

Track-record argument: A type of epistemic circularity in which one forms a number of beliefs using some method *M*, introspectively observes the content of those beliefs, then concludes that method *M* has gotten one true beliefs on all these occasions and therefore is likely reliable. (§6.3.3)

Undercutting defeater: A defeater for a proposition, *p*, which works by casting doubt on the reliability of one's method of forming the belief that *p*. (§4.5.1)

Underdetermination: In science: The phenomenon whereby the empirical data can be explained by more than one theory. (§14.4.1)

Uniformity Principle: The principle that unobserved things tend to resemble observed things. (§12.1.2)

Valid: Of an argument: Such that the premises could not all be true while the conclusion was false. (§1.3.3)

Verification criterion of meaning: See verificationism.

Verificationism: The view that the meaning of a sentence is given by its verification conditions; hence, if there cannot be evidence for or against a particular statement, then that statement is meaningless. (§10.3.1)

Weak foundationalism: See foundherentism.