

## Chapter 1. First philosophy and foundationalism

The life of an inquirer, above all, is concerned with cultivating a system of beliefs. We gather evidence, we weigh hypotheses, we evaluate arguments, and we test our theories, all for the trouble of deciding what to believe. This is how we construct our *worldviews*—which is everything we believe, from the mundane to the abstract and theoretical. But what is the ultimate value of this activity? What do we hope to achieve in developing all of our various beliefs?

One answer—which will seem like a platitude—is that *we want to know what the world is like*. We strive to *know* the truth, and we strive to know *as much of it* as we can. Notice that there are several different components to this goal. First, we want our beliefs to be *true*. (We don't want to accept any mistaken beliefs about the world.) Secondly, we want to accumulate as many true beliefs as we can. (We could avoid falsehood altogether by refraining from all belief, but that would be giving up on our intellectual ambitions.) And lastly, we want our true beliefs to be *knowledge*. (It is possible to stumble upon true beliefs by accident, by making lucky guesses, but that would be unsatisfactory since we want assurance of their truth.)

I think that this platitudinous answer is right: the aim of our cognitive lives is to acquire knowledge. But despite the triteness of this answer, it turns out that not everybody agrees with me. There are various more pragmatically-minded philosophers who will scoff at the idea that knowledge should be our ultimate cognitive aim. If you'd ask them, they will tell you that knowledge is too lofty, and too elusive, to be a worthwhile goal of intellectual inquiry. Instead, they would recommend that we set our sights on something more down to earth. Perhaps we should seek for beliefs that are *practically useful to us*, or *merely probable given our evidence*, or *well-confirmed by our experience*. But we shouldn't worry too much about whether our beliefs constitute *genuine knowledge*.

How dare they say this? In my experience, the pragmatic “who-cares-about-knowledge” attitude can usually be traced to an underlying (sometimes tacit) acceptance of *skepticism*—that is, the thesis that we can never attain true knowledge about the world around us. Whether they realise it or not, the pragmatist is often motivated by something like the following inference: *true knowledge is unattainable, therefore we ought to settle for something less ambitious* (e.g. practical utility, probable opinion, or empirically adequate belief).<sup>1</sup> If this diag-

---

<sup>1</sup>I detect a certain irony in this. Pragmatists usually see themselves as having an easy-going approach to belief, according to which you should believe whatever works best for you, given your circumstances and practical concerns, without fretting too much about ultimate reality or knowledge. However, the moment they try to justify their pragmatic approach to philosophy, they transform into extreme skeptics. They say we cannot know because they

nosis is right, then my pragmatist opponents are really skeptics in disguise. So the *real* issue that I need to grapple with, then, is skepticism. This now gives me something to worry about: how do I attain knowledge? How is knowledge even possible? And why have some philosophers been led to believe that it isn't possible? Are they right to conclude this?

Here is a picture (as old as Plato) that can make skepticism seem inevitable. According to this picture, I am an occupant of only one of two contrasting and non-overlapping worlds. The world that I inhabit is a purely individualistic and subjective one, which we may call the "world of appearances." Within this world, I only ever directly interact with purely subjective phenomena. For instance, I only get to attend to my own ideas and sensations, and I only get to experience *mere appearances* of how things are (nowadays called 'sense data'). I never directly interact with any other things *as they are in themselves*.<sup>2</sup> The rest of the objects that make up total reality (besides myself) occupy the other world, which we may call the world of "objective reality." This is the world that includes everything that is external to myself, such as material beings, other people, abstract objects, immaterial souls, and God (if such things exist).

As an inhabitant of my own subjective world, I have direct access to its contents. Knowledge of my own ideas, thoughts, sensations, and sense data is relatively easy to come by, since each of these things is essentially dependent on my perspective. But this sort of 'subjective' knowledge is not my chief interest; I also aspire to know objective reality. However, there's a great gulf between my own mind and the external world, which makes knowledge of that reality exceedingly difficult to attain. The appearances create a veil, which blocks the external objects (as they are in themselves) from laying bare to my view. If I am to ever attain knowledge of them, I must build some sort of intellectual bridge between my world and the world of objective reality. It's as if I must extend my cognitive reach outside of the confines of my own head in order to grasp the external objects and see them as they really are. But according to the two-world picture, this is patently absurd. The only basic ingredients we have to know anything are what's already inside our own heads. We cannot leave the world of our own mind, or view things from any other perspective than our own. So if that is what it takes to attain knowledge of objective reality, then such knowledge is patently impossible.

One of my overarching aims in this project is to fight against the skeptical temptation. I want to defend our claims to knowledge about ordinary external things. This commonsense view is typically known as *non-skeptical realism*. To this end, I will need to cast doubt on the two-world picture given above. We replace it, instead, with a one-world picture. If there is any theme to this project, it is this: human subjects do not stand apart from the world—we reside *within* it. (We are "beings-in-the-world", as Heidegger would say.<sup>3</sup>) As such, knowledge is not a matter of bridging some deep divide between subject and

---

advocate extremely high standards for knowledge.

<sup>2</sup>To use the famous phrase from Immanuel Kant.

<sup>3</sup>Not a philosopher that I'm usually fond of, mind you. But in this context, his phrase works fairly well.

object using only subjective materials. Whatever division there is between the subjective and objective realms, it presents no great barrier to knowledge. We should instead see knowledge as a relation borne between subjects and their environment, and it depends as much on objective features of the environment as it does on individual efforts.

Nothing about this theme is new. There has been a scattering of philosophers throughout history who have warned us against the two-world picture, including Thomas Reid, Martin Heidegger, and John McDowell (to name a few). Moreover, this theme has recently manifested in contemporary philosophy in the form of various “externalist” theses, which are widely defended today. I cannot hope, for this project, to say anything that hasn’t been said before. My only hope is to package these themes into something that’s short and comprehensive.

## 1 Knowledge and doxastic decisions

I said that my aim is to defend our ordinary knowledge of ordinary things. If I’m successful, then I hope to have shown that the following things are amongst my knowledge: that I have two hands, that I had coffee with my breakfast this morning, that the sun will rise tomorrow, and that Kilimanjaro is in Africa.<sup>4</sup> Each of these, I take it, are fairly ordinary things to believe.

But naturally one might wonder, is it *really* that important to defend our knowledge of such mundane things? At first it would seem that the only reason to defend this knowledge is so that I can quarrel with one particularly unruly philosophical character—the skeptic. But isn’t that exercise merely academic (in the pejorative sense)?

Taking this thought a bit further, one might expect that the real value of philosophical reflection on the foundations of knowledge is that we might hope to deliver some principles, or methods, or rules that would help us determine what to believe throughout the course of inquiry. After all, isn’t the point of epistemology to help us decide what to believe? Shouldn’t we instead be seeking some sort of guide to distinguish the true beliefs from false ones? Taking this thought a bit further, one might have thought that the reason epistemology *matters*—to scientific inquiry or to being an informed citizen—is because it affords a sort of practical wisdom for detecting the truth. And why would quarrelling with the skeptic give us that?

I have a lot of sympathy with this complaint. I think that it is both half-right and half-wrong. But at this stage of my project, I have not yet laid down enough theory or tools to explain the sense in which it is right and the sense in which it is wrong. For now, we just need to take care to distinguish *three* types of questions, which correspond to three tiers of epistemological inquiry. (I) In the first tier, there are the *practical* questions about what we should believe in the various evidential situations that we find ourselves in. For example, what should

---

<sup>4</sup>Each of these beliefs of mine has its source in a different epistemic modality: perception, memory, induction, and testimony (respectively). I will cover each of these epistemic modalities throughout the course of this project.

I believe, given my evidence? What rules or methods should I follow to decide what to believe? These are the concerns of *applied epistemology*. (II) In the second tier, there are philosophical questions about systems of belief *in general*, where we abstract away much of the detail about the particular kinds of evidence we have. These include fundamental normative questions about knowledge, justification, and epistemic obligation (i.e. what we ought to believe). I have in mind such questions as *can we know anything at all? Is there any foundational knowledge (i.e. knowledge that we can attain without further argument)? What is our most basic epistemic obligation?* These are the questions of *normative epistemology*. (III) Lastly, in the third tier, there are the *metaphysical* questions about the various epistemic properties—namely, knowledge, justification, and epistemic obligation. Epistemologists interested in this tier of inquiry want to know about the *nature* of these properties. For instance, some philosophers are very interested in figuring out how to analyze knowledge. (But, for better or worse, I am not one of those philosophers.<sup>5</sup>) One interested in this pursuit will also be interested in the sorts of states of affairs that determine these properties. In particular, we might be interested in whether knowledge depends entirely on the perspective of the subject, or whether it involves external factors as well (we will revisit this question later). We can call this third tier of inquiry *meta-epistemology* since it is the metaphysics of epistemology.

I hope that by the end of this piece I will have displayed some of the subtle ways in which these three tiers of inquiry are interconnected. But for now let me reiterate: these are distinct questions, and so we should start by keeping them separate. Ultimately I would love to gain some practical wisdom for choosing what to believe. It sure would be nice if our exercise in foundational philosophy could offer some of this practical wisdom. But let's not get too ahead of ourselves. For now we must deal with the second-tier question: *can I know anything at all?*

## 2 Agrippa's Trilemma and Foundationalism

To even begin to think about the question, *how can I know anything at all?*, it is helpful to take a sample of things that you think you already know, and then consider potential challenges to this knowledge. For example, I think that I know a lot of things about my current immediate environment: that it is

---

<sup>5</sup>The epistemological literature throughout the second half of the 20th century was replete with attempts to analyze knowledge into more basic properties. But this pursuit isn't so common nowadays. In case the reader is curious, I have no deep, informative answer to give them as to *what knowledge is*. In all likelihood, the concept of knowledge is one of those concepts that doesn't have any deep, informative analysis, in terms of necessary and sufficient conditions. But that's okay; *most concepts are like this*. There are lots of concepts that we can grasp well enough without being able to decompose into simpler concepts. And this is no obstacle to doing important philosophical work using those concepts. (There used to be a time when philosophers would get very worried about concepts that could not be defined in terms of the simpler concepts of experience and mathematics, but I think we now have good reasons to think that their worries were misguided. Unfortunately, following up on this line of thought would take me too far afield.)

raining outside, that there's a coffee cup before me, that I have hands, etc. I also think that I know many things about my past from my memories: that I had coffee with my breakfast, that I have visited Seattle recently, and that I've been in a loving relationship for over five years. I also think that I know some basic, highly predictable things about the future: that tomorrow the sun will rise, I will wake up and have breakfast, and eventually do some writing. Lastly, I think that I know a great deal about the world from the reports of other people. This includes all of the scientific knowledge that I gained through reading textbooks, and my knowledge of current affairs from trustworthy media outlets. But maybe I do not know any of these things. That is still to be seen.

How could I possibly fail to know these things? Well judging by history, there are two highly influential lines of thought that lead to the skeptical conclusion. For now we will consider one very ancient style of skeptical argument, which goes all the way back to the 1st century Greek skeptic Agrippa. It is commonly called Agrippa's Trilemma.<sup>6</sup>

When thinking about skeptical arguments, it is often helpful to imagine yourself in a dialogue with an imaginary foe, the skeptic. You wish to assert that you know a certain proposition, and the skeptic wishes to prove otherwise. You present your arguments and they will try their best to refute them; and likewise they will present their arguments which you will try to refute. Now, when it comes to Agrippa's trilemma, the skeptic of this persuasion has a very stubborn and irritating personality (one that you've probably encountered before in young children). Suppose that you claim to know something. This skeptic will challenge you—*how do you know that?* And let's say that you oblige and give them an answer. Whatever you say, they will challenge you again—*now how do you know that?* And so on. Whatever reasons you give to this skeptic, they will keep on asking you how you know them. Since this skeptic is indefatigably stubborn and persistent, there can only be three outcomes: either this game goes on forever, or you end up going in circles, or you get fed up and stop answering. Whichever the eventual outcome, the skeptic will think that they've won.

This skeptical challenge is not just a silly game that children play. The proper response to it will have deep consequences for the structure of justification. But before we get into that, we need to first make the skeptical reasoning a bit more explicit. Let's say that I have a belief *B*. In order for my belief to be knowledge, it must be (at least) *true* and *justified*.<sup>7</sup> But in order for *B* to be *justified*, (the skeptic claims) I must have some *reason* for thinking that it is true. I must provide some sort of *argument* or *proof* of *B*, which, in turn, will appeal to other beliefs of mine, as premises.<sup>8</sup> Since *B* is an arbitrary belief, the skeptic

---

<sup>6</sup>It is also sometimes called the Münchhausen trilemma.

<sup>7</sup>Please do not mistake me for saying that knowledge *is* justified, true belief. I am only reporting necessary conditions.

<sup>8</sup>It might be worth being more explicit about this terminology. An *argument* is an abstract structure that consists of a set of premises and a conclusion, where the premises purport to support the conclusion. An *inference* is a *mental act* whereby a subject attempts to justify one belief on the basis of other beliefs, according to some argument.

will say the same thing for all beliefs. That is their first premise.

(i) For all beliefs  $B$ , if  $B$  is justified, then it must be *inferred* from further beliefs.

The skeptic will try to justify this premise on the grounds that an unargued belief is an arbitrary belief. If you cannot cite a reason for believing what you do, then your belief is no better than a guess or an opinion or a hunch. Next, let's say that I have succeeded in providing an argument for my belief  $B$ , from premises  $B_1 \dots B_n$ . If my inference is to confer justification onto  $B$ , then the beliefs that I appeal to as premises must also be justified. That's premise two.

(ii) For all beliefs  $B$ , if  $B$  is justified by inference from  $B_1 \dots B_n$ , then  $B_1 \dots B_n$  are justified.

But then the skeptic will ask how those further beliefs are justified. And at this point, we only have a few options. If we grant the skeptic their premises so far, then we cannot say that  $B$  is ultimately justified by any *uninferred* belief, because (i) would entail that the uninferred belief is unjustified and (ii) would then entail that  $B$  is unjustified. So now there are only two options that remain. Either our process of justifying  $B$  must go on *ad infinitum* by making further inferences from further beliefs, or we'll end up going in a circle (that is, some of our beliefs are ultimately inferred from themselves). But according to the skeptic, neither of these options allow for *real* justification. Their argument is thus completed by two further premises.

(iii) For all beliefs  $B$ ,  $B$  cannot be justified by an *infinite* chain of inferences.

(iv) For all beliefs  $B$ ,  $B$  cannot be justified by a *circular* chain of inferences.

Each of these last two premises is fairly plausible. If a belief of mine must rest on an infinite chain of inferences, then it feels like I can never achieve real justification. After all, we human beings cannot complete infinite processes, so *we* cannot justify our beliefs in this way. On the other hand, it also doesn't seem like justification can be generated by circular reasoning. If a belief isn't justified to begin with, then it isn't going to become justified if you infer it from itself.

Now, if *all* of the skeptic's premises are true, then the skeptic has won. For (i) - (iv) entail that no beliefs are ever justified. And if no beliefs are justified, then knowledge is impossible. So given that we aim to defend our claims to knowledge, how do we deal with this argument? Since the argument is valid, the only rational way to respond is to deny one of the premises.

Throughout the long history of philosophy, each one of these premises has been denied by someone or other.<sup>9</sup> But I will not pause here to investigate the prospects for denying (ii) - (iv). To my mind, I do not need to, because the argument has a clear weakness. Its weakness is premise (i).

Premise (i) states that beliefs can only ever be justified if they are *inferred*. Let us get clearer on what is being asserted. First, what is it for a belief to be *justified*? Roughly, a belief is justified if it is the appropriate one to have, given the subject's evidential situation.<sup>10</sup> When a subject has responded responsibly, by forming the beliefs that they do, to the evidence that has been given to them, then their beliefs will count as justified. Secondly, what is it for a belief to be *inferred*? Roughly, a belief is *inferred* when the subject performs a mental act of forming a new belief on the basis of an argument from previous beliefs (the premises). So putting these two definitions together, premise (i) says one has only ever appropriately responded to their evidence if they have an argument from previous beliefs. It's as if the only evidence that we have is our other beliefs. Is that really true?

Let's consider an example. Suppose, for the sake of argument, that I perceive that I have hands. (Some philosophers have claimed that I can never perceive this directly. But let's not dispute with them now; that will have to wait for chapter 3.) And suppose that I form the belief *that I have hands* as a response to this perception. Finally, let's suppose, as a hypothesis, that this belief *is justified*—it *is* an appropriate response to my evidential situation. Must it follow that I have other beliefs that serve as premises in an argument to the conclusion *that I have hands*?

I think it's fairly obvious that the answer is *no*. If I perceive that I have hands, then the best doxastic response is to believe that I have hands; I don't need the additional support from further beliefs to infer I have hands. Indeed, I think that this is *far* more obvious than the Agrippan skeptic's premise (i), which is under dispute.

Besides, what other beliefs could I appeal to as premises for this conclusion? Perhaps the skeptic will claim that I need to believe *that I am having this perception*. After all, haven't I claimed *that my perception that I have hands* is the *basis* for my belief's justification? It is true that I have claimed this. But I don't mean to say that a *higher-order belief about my perception* must serve as a premise in any inference. In fact, it is not plausible to think that one must have beliefs *about* their perception in order to have perceptually justified beliefs. After all, small children do not have higher-order beliefs about their perception because they do not yet have the concept of perception. Yet, they are presumably right to believe that they have hands (once they have the concept

---

<sup>9</sup>Those who deny premise (iv) are called *coherentists*. This view was fashionable during the middle of the 20th century, mostly amongst American philosophers—and most notably by W.V.O. Quine. In their view, a justified belief system is built holistically, where each belief supports each other in a web-like structure (rather than a tree-like structure with a foundation). Those who deny premise (iii) are called *infinitists*, although to my knowledge they have been fairly rare. I don't know of any label for those who deny (ii).

<sup>10</sup>Let's say that a subject's 'evidential situation' comprises all of the factors that determine their evidence: e.g. their experiences, their background beliefs, their environment, etc.

of hands). What this example shows is that we must take care to distinguish two very different sorts of claims. First, there is the claim that I endorse, which is that one can have beliefs that are justified on the basis of perception. Or more generally, there are the claims of the form *S's belief that p is justified on the basis of F*. And secondly, there is the claim that I reject, which is that one justifies their perceptual beliefs by inferring them from beliefs about perception. Generally, such claims have the form *S's belief that p is justified by inference from S's belief that F*. Such claims are logically unrelated to each other, and hence they mustn't be conflated.

Perhaps the skeptic will try another rejoinder. They could say (correctly) that in order for my belief to be justified by perception, my perceptual processes must be working reliably. I couldn't be justified if my eyes typically deceive me, even if they happen to be right this one time. The skeptic then goes on to say that I must therefore *believe* that *my perception is reliable*, or else I couldn't be justified. Moreover, my belief that I have hands must be grounded in this further belief.

In other words, the Agrippan skeptic is suggesting that my belief must proceed from something like the following inference: (1) I perceive that I have hands, (2) my perception is reliable, therefore (3), probably, I have hands. But again, is it plausible that I *must* go through these steps in reasoning in order to be justified in believing that I have hands? *Not very*. Again, the fact that children do not perform these inferences and are still justified shows that these higher-order inferences are superfluous for justification. It is true that my perception must be working reliably in order for it to confer justification. But this doesn't mean that my justification must flow from *beliefs about my perception's reliability*.

Once again, the Agrippan skeptic's premise (i) seems to gain some superficial plausibility by conflating two logically-distinct claims. On the one hand, we have a *necessary condition* for justification: S's belief that *p* is justified only if *F* (e.g. their perception is reliable). On the other hand, the skeptic claims that another *belief* is necessary for justification. S's belief that *p* is justified only if *S infers p from their belief that F*. But again, these two kinds of claims are logically unrelated, and hence they need to be kept separate.

I have not yet said everything that I need to say in order to fend off the Agrippan skeptic. But from what I have said, I can now give a name for my response to Agrippa's Trilemma and my theory of the structure of knowledge and justification. The theory is classically known as *foundationalism*. In a word, foundationalism is the view that the Agrippan skeptic's premise (i) is false: some beliefs can be justified without inference. Which beliefs? Well, different foundationalists have different ideas as to which beliefs are candidates for non-inferential justification. I haven't yet explicitly stated my view on the matter (although my discussion should give a pretty obvious clue as to how I feel about perceptual beliefs). At any rate, a belief that can be justified on the basis of something other than inference is called a *foundational belief*. The beliefs that must be justified by inference from the foundational ones are



called *non-foundational beliefs*.<sup>11</sup> For the foundationalist, the structure of justification is like a tree. The foundations are like the roots and stem, and the non-foundational beliefs rest upon them like branches.

Earlier I alluded to another Agrippan skeptical complaint towards foundationalism: that without an argument for a belief, that belief is *arbitrary*—no better than a hunch or an opinion. This is tantamount to saying that foundational beliefs are inherently arbitrary. But from what we've seen above, we can now find this to be emphatically wrongheaded. Foundational beliefs need not be arbitrary because they may be *based on something*. Perhaps my belief *that I have hands* is foundational and is *based on perception* (not my beliefs about perception, but perception itself). In that case it would clearly not be an arbitrary guess—it is grounded in my perceptual state. We just have to keep in mind that *basin* *one's beliefs on certain states* (not beliefs) is one thing, and *inferring them from other beliefs* is another. As long as we heed this distinction between *basin* and *inferring*, foundationalism will remain highly attractive.

Before we leave this topic, I would like to address a completely different kind of worry that might be voiced by the Agrippan skeptic. This skeptic would protest that without an argument, we may be unable to defend our foundational beliefs.<sup>12</sup> So when the skeptic challenges our beliefs, we may have nothing to say that could convince them that our beliefs are true. Perhaps we would plead that these beliefs are foundational. But if that is all we can say, then the skeptic will accuse us of 'begging the question.' And if we can't defend these beliefs against naysayers without begging the question, then perhaps they are not on epistemically solid ground.

Our skeptical interlocutor is probably right that we could never rationally convince them of the truth of our foundational beliefs. But they would be wrong to infer that our foundational beliefs are thereby unjustified. This skeptical complaint appears to rest on an unstated premise: that a subject's belief is justified only if they can convince others of its truth. And this tacit premise is flagrantly false.<sup>13</sup>

To see that it is false, imagine that you are the sole first-hand witness of an otherwise incredible event. Picture yourself looking outside your window and you see what is unmistakably a lynx giving a badger a high five. Or on a more serious note, imagine that there's a very powerful and beloved local politician with a spotless reputation, and you happen to witness him committing a crime behind closed doors. When you try to tell others of what you saw, nobody will ever believe you, especially if there is no corroborating evidence. But it doesn't follow that *you* can't know what you saw. You perceived it with your own two

---

<sup>11</sup>Technically I have left out a third option. There could be beliefs that are justified *on the basis of other beliefs*, but yet the role that those other beliefs play in justification is not the simple one of serving as premises in an argument. I think it's likely that justification sometimes works this way, but I have no reason to discuss those cases here.

<sup>12</sup>This isn't to say that we *can never* defend our foundational beliefs. A belief's status as foundational does not preclude the possibility of corroborating it with inference.

<sup>13</sup>"So what is really wanted... Is it that we must find something to say that would, of necessity, make anyone who heard it become reasonable forthwith? That would be a spell, not an argument." David Lewis, *On the Plurality of Worlds*, 115.

eyes!

As for the allegation that foundational beliefs are question-begging, we have to distinguish two senses of ‘question-begging.’ According to the first sense, an *inference* is question-begging if it assumes its conclusion as a premise. Clearly a foundational belief cannot be question-begging *in this sense* because their justification does not rest on any inference at all. (Remember, the foundationalist does not endorse circular justificatory structures in a system of justified belief.) According to the other sense, it is question-begging to *assert* a belief of yours within the context of a conversation when that very belief is up for dispute. Once your interlocutor has challenged your belief, it becomes conversationally inappropriate to assert it without further evidence. But that just reflects the conventions of interpersonal conversation; it has no bearing on whether your belief is actually justified. It is quite possible to have an epistemically justified belief that is inappropriate to assert in certain conversational settings. (Here is an example that will mesh well with the theme of this manuscript later on. Imagine that you are debating a climate change denier over whether climate change is real. Within the context of this specific debate, you cannot simply assert the reality of climate change and cite expert opinion. That’s not going to fly, because your interlocutor is already committed to disbelieving the experts. But it doesn’t follow that your belief or your trust in the experts is unjustified.)

That is all that I will say in defence of foundationalism for now. Let me briefly sum up the main ideas of this section. In the first half, we saw one highly influential route to skepticism called Agrippa’s Trilemma. This skeptical challenge rested on the idea that every justified belief requires an inference from other beliefs. My central contention of this chapter is that this claim is false. I thus endorse the defining thesis of *foundationalism*: that some beliefs can be justified and known without resting on any argument or inference. I claim that a belief can be an instance of foundational knowledge provided that it is *based on the right sort of state or process* (which need not involve inference). Our foundational knowledge is based on the states and processes that confer non-inferential justification. But what is “the right sort of state and process”? And what sorts of beliefs are suited to be foundational knowledge? I have not yet “officially” given any answer to this. My only purpose so far has been to defend the bare structure of the foundationalist account of knowledge. (However, I have indicated where my sympathies lie: I believe that perception is a fairly central source of foundational knowledge. I believe that there are other sources as well. But it will take the duration of this project to defend my claims about the sources of foundational knowledge.) So for the remainder of this chapter and the next, we should investigate some of the candidates for foundational knowledge.

### 3 Presuppositions, the laws of logic, and the reliability of reason

When some philosophers start thinking about the foundations for knowledge, they are instinctively led to muse about our beliefs about highly abstract and *recherché* subject matters. One who exhibits this tendency might suppose that our ultimate foundational beliefs ought to concern the laws of logic or the reliability of our own reason, or some other such baffling thing. This kind of philosopher might even claim that these beliefs are the *only* ones that are fit to be the foundations for a system of knowledge. They would probably motivate their position like this. First they would say that one's belief in the laws of logic and the reliability of one's own reason must be the *ultimate presuppositions* for one's worldview. In all of our thought, and all of our reasoning, we must presuppose that the laws of logic are true and that our reasoning is reliable. For if these fail, then we cannot have any knowledge at all. Hence, we must *know* that the laws of logic are true and that our reasoning is reliable before we can know anything else. Therefore these beliefs ought to be our starting point, and therefore they must constitute our foundational knowledge.

Here we have a candidate proposal for the scope of foundational knowledge: it claims that our foundational knowledge is coextensive with our "ultimate presuppositions" (whatever those are). Now how plausible is this claim? I would like to suggest that, despite its loftiness, this claim is nowhere near as plausible as it might appear at first. In fact, it will not be figuring into my own foundationalist conception of knowledge. But since this idea has enticed many people, it is worth explaining why I reject it. Following this, I will then attempt to make the briefest remarks possible on the subjects of logic and the reliability of reason.

The best way to unravel a sophisticated position in philosophy is to become clear on the terms that define it. In this case, the view under consideration holds that one's foundational knowledge must consist of ultimate presuppositions. Now, we have already defined what it is for a belief to be foundational knowledge: it is knowledge that can be justified without inference from other beliefs. But what about this other concept, "ultimate presupposition"? What is a "presupposition" and what makes a presupposition "ultimate"?

At this point I must invoke my interpretational liberties, since the concept of 'presupposition' does not have any widespread definition in epistemology.<sup>14</sup> But given the common motivations for their view, I think it is fairly clear what its advocates have in mind. When they say that one belief presupposes another, they mean to say that a subject must *know* or *justify* the latter, if they are to know or justify the former. Let's make this a bit more explicit.

---

<sup>14</sup>It is clear that the proponents of this view are not using the term "presupposition" with the usual meaning that it carries in academic philosophy and linguistics. Typically, the word "presupposition" is reserved for a certain *semantic* phenomenon (the truth-aptness of a sentence requiring certain background facts, which may or may not be known to speakers). There is no distinctly epistemic notion of "presupposition" that has widespread use in contemporary epistemology.

*Definition.* A belief  $B$  presupposes proposition  $P$  just in case, in order for a subject to be justified in believing  $B$ , they must also have the justified belief that  $P$ .

With this definition, we can see that in some circumstances, some of the logical consequences of a belief will also be its presuppositions.<sup>15</sup> So for instance, to be justified in believing that there is a coffee cup before me, I must presuppose that coffee cups exist, and (perhaps, in general) that material objects exist. But other than that, how do we determine *which* propositions are the presuppositions of a given belief? That is yet to be seen. Now, what is it for a presupposition to be “ultimate”? Again, we must extract a definition out of what these philosophers say. But once again this is fairly straightforward. A presupposition is *ultimate* if it is a presupposition for all other beliefs in a subject’s worldview, and it does not itself require any further presuppositions.<sup>16</sup>

Now let’s return to foundationalism. Having clarified its terms, we can now see what the proposal under consideration amounts to. Basically, it is claiming that the *only* beliefs that can be justified without inference are the ones that we *must* hold (and justify) if we are to have any justified beliefs at all. Or to put it another way, if a belief isn’t necessary for the justification of *all* other beliefs, then it can *only* be justified by inference from the ones that are.

Now that the view is unraveled and exposed, how plausible is it? My own opinion is that it is highly *implausible*. For there appears to be many beliefs that don’t qualify as ‘ultimate presuppositions’, yet are justifiable without inference nonetheless. Consider, once again, my belief that I have hands. In the last subsection I argued that *if* this belief is justified, then its justification is (plausibly) based on perception, and not based on inference from other beliefs. Yet we would never count this belief as an ‘ultimate presupposition.’ I do not, strictly speaking, *need* to believe that I have hands in order to uphold every other belief in my worldview. (Suppose that I lose my hands in an accident, and therefore stop believing that I have hands. It’s not as if the rest of my beliefs will lose any justification.) So a foundational belief does not need to be an ultimate presupposition.

No doubt, an advocate of this presuppositional version of foundationalism is unlikely to be moved by my counterexample. They will insist that there are certain propositions that I *need* to know, in order to have any other justified beliefs. At this point they will probably gesture towards the laws of logic or the reliability of my own reason or something like that. But what reason can they offer as to why we *must* know these things in order to have any justified beliefs? In the brief vignette I gave earlier, they argued that *we need to know* these

---

<sup>15</sup>But not always! Each belief of mine will have, as logical consequences, infinitely many propositions of arbitrary complexity that result from compounding the logical operations indefinitely. And presumably I do not need to form beliefs about those highly complicated things in order to justify my original beliefs.

<sup>16</sup>This definition may not be perfectly precise, but for our purposes its unclarity will not matter.

things, *because if they weren't true*—that is, if we used invalid logical principles or our rational faculties were otherwise unreliable—then the rest of our beliefs will likely be untrue and hence unjustified.<sup>17</sup> But is this a valid reason for their view? In the previous section I already gave some reason to think that it isn't. To see this, let's grant, for the sake of argument, that in order for me to be justified in believing that I have hands, there are certain necessary conditions that must be met. And let's even grant to the presuppositional foundationalist that they've correctly identified those conditions: I must be reasoning according to a valid logic and I must have reliable rational faculties.<sup>18</sup> Does it follow that I must *know* or even *believe* that these conditions are met, in order for my belief to be justified? Not at all. Once again, we must not confuse a *necessary condition* for justification with *what a subject must believe in order to be justified*. It may be true that a subject *S*'s belief that *p* is justified only if some condition *F* obtains (e.g. *S reasons according to valid rules of logic or has reliable rational faculties*), but it doesn't follow that *S* needs to *know* (or even believe) that *F* obtains in order to be justified in believing *p*—*F just has to obtain* (whether *S* knows it or not). Or to put this point another way, a necessary condition for justification does not automatically entail a presupposition.

Once we have taken this lesson to heart, we will find that it isn't so straightforward to name the *real* presuppositions of any given belief. For we cannot simply identify them by looking at what has to be true for a belief to be justified. The genuine presuppositions are something else: they are the beliefs that are required for rational coherence. And it's not so easy to give a general account of what those must be. Indeed, I will not even try to give such an account here.<sup>19</sup> I do not need to. After all, my real concern here is foundationalism, and even if we could give an account of epistemic presupposition, there would still be another problem with the presuppositional approach to foundationalism. Namely, that even *if* there are ultimate presuppositions, they need not coincide with our foundational knowledge. There can be ultimate presuppositions that are not foundational beliefs, and there can be foundational beliefs that are not ultimate presuppositions.<sup>20</sup>

---

<sup>17</sup>This reasoning supposes that the reliability is necessary for justification, but I'm content to grant this.

<sup>18</sup>In actual fact, I have doubts that either of these are necessary conditions for knowing that I have hands. My reason for doubting this with respect to the reliability of *all* of my rational faculties will come out later in this section. As for the laws of logic, suppose that my belief that I have hands really is an *uninferred belief*. Well if it is *uninferred*, then my belief does not rest on any argument, and hence it does not rest on any use of the rules that are under the purview of logic. So even if we suppose that I generally reason according to *invalid* rules of inference, it's unclear how this would affect the justification of my *uninferred* beliefs.

<sup>19</sup>Here is my best guess. A belief *b*'s presuppositions would include (i) the obvious logical consequences of *b*, and (ii) on some occasions where *b* is a non-foundational belief, it could turn out that *b* can only be justified by inference from a specific belief *p* (e.g. the belief *that I do not not have hands* can only be justified by inference from my belief *that I have hands* using double-negation).

<sup>20</sup>The epistemological relationship between our inferred beliefs and our foundational beliefs is not the same as the relation between a belief and its presuppositions. But perhaps those who are enticed by the presuppositional approach to foundationalism are guilty of confusing the one relationship with the other.

Once again, let's look at my belief that I have hands. I have already remarked on how, if this belief is justified, then it can be justified without inference; hence it is a candidate for foundational knowledge, and yet it is not an ultimate presupposition. Now let's suppose (hypothetically) that this belief carries with it some presuppositions. Let's say that the belief that I have hands presupposes that material objects exist. So in order to be justified to believe that I have hands, I must also have the justified belief that material objects exist. Would this preclude my belief that I have hands from being foundational? No, it wouldn't. That is because, for all I've said, I could justify my belief that I have hands without inference (maybe by perception) and then justify its presupposition in some other way.<sup>21</sup> It wouldn't follow that I would have to justify my belief *that I have hands* by inferring it from the premise *that material objects exist*. Just because a belief carries a presupposition, it doesn't follow that I can only justify that belief via *inference from that presupposition*. All that follows is that the presupposition must be justified in some way or another. This doesn't impose any sort of inferential relationship between a belief and its presuppositions. So with all of that said, a belief may carry a presupposition, but that doesn't disqualify it from being a foundational belief.

By now I hope that the lesson will be clear: this whole issue of presupposition has no obvious connection to foundationalism. The theory that I develop here does not accept the claim that our foundational beliefs are our ultimate presuppositions. So for the remainder of this project, I will leave the whole idea of presupposition behind. But having now swept this idea into the dust bin, I would still like to say a few words about the two beliefs that are commonly alleged to be our ultimate presuppositions: the laws of logic and the reliability of our reasoning. My purpose is just to clarify where these things stand with respect to my own foundationalist project.

### 3.1 Foundationalism and the laws of logic

Let's start with the laws of logic. First of all, what exactly is logic, anyway? Although I cannot hope to give any deep answer to this question here (this isn't a book on the metaphysics of logic, afterall), we can make a few casual observations. For one, we should observe how the contemporary practitioners of the discipline of logic understand their own subject matter. According to them, the subject of logic has two overarching objects of study: the phenomena of logical validity and logical truth. Roughly speaking, *logical validity* has to do with the facts of *valid argument* and *inference*—of *what follows from what*. For example, it is a fact that every argument that instantiates the pattern  $\langle P, \text{if } P \text{ then } Q; \text{ therefore } Q \rangle$  is a valid argument; the truth of the premises will guarantee the truth of the conclusion. On the other hand, the phenomenon of *logical truth* has to do with the range of sentences that are *true in virtue of*

---

<sup>21</sup>I'm inclined to think that the justification for my belief *that material objects exist* will involve an inference from mundane beliefs about the existence of hands, tables, chairs, etc., and some philosophical thesis about the nature of material objects.

*their logical form.*<sup>22</sup> For an example of such a sentence, take any instance of the schema  $\langle P \text{ or it is not the case that } P \rangle$ ; no matter what declarative sentence we substitute for  $P$ , the result will yield a truth.

So what, then, are the *laws* of logic? Presumably they are just the general facts that explain the logically true sentences and the logically valid arguments. If that's right, then I have just mentioned two of them. One law of logic is that *modus ponens* is valid: that every argument of the form  $\langle P, \text{ if } P \text{ then } Q; \text{ therefore } Q \rangle$  is truth-preserving. And another law is *the law of the excluded middle*: that every sentence of the form  $\langle P \text{ or it is not the case that } P \rangle$  is true. And there are *many* others.<sup>23</sup> We will not bother to catalogue them all here, for that would require me to write an entire treatise on logic.

The real question for us, as we think about foundationalism and systems of knowledge, is the extent to which one must *know* about the laws of logic in order to have a system of knowledge. Now in order to answer this question, we need to make a crucial distinction. We need to distinguish between *having knowledge about the logical laws*, on the one hand, and *reasoning according to valid logical principles*, on the other. In order to possess the former kind of knowledge, one must be capable of representing the logical laws in thought. So one would at least have to possess such concepts as *logical form*, *validity*, *logical truth*, *modus ponens*, *the law of the excluded middle*, and the other concepts that we use to talk *about* logic. That way, one is able to form *beliefs* about the logical laws. But in order to possess the skill of *reasoning according to valid logical principles*, one does not need to have any sophisticated knowledge that explicitly represents the facts *about* logic. Rather, all that one needs is the capacity or 'know-how' for recognizing when one thing follows from another. One needs to be able to see, for example, that if *it is either raining or snowing*, and *it is not snowing*, then *it must be raining*. One must be able to recognize that *either it is raining or not raining* must be true, no matter the state of the weather. One must have the tendency to reason in accordance with *modus ponens*, and a tendency to avoid fallacious patterns of reasoning. And likewise, one must possess other similar skills, corresponding to the other logical principles. All of these skills could (in principle) be possessed by someone who lacks the higher-order reflective knowledge that explicitly represents the principles that they use in reasoning.<sup>24</sup>

---

<sup>22</sup>What is "*logical form*"? It's not easy to say. For a more precise definition of logical truth, the classical semantic conception holds that a sentence is logically true if it is true in virtue of the semantic contributions of its logical vocabulary ("if . . . then", "not", "and", "or", "every", "some", etc.) and not in virtue of its non-logical vocabulary.

<sup>23</sup>At some time during the middle ages and up until the 20th century, popularizers of the subject were prone to declare *three* laws as being *the* laws of logic. Particularly, they would name *the law of the excluded middle*, *the law of non-contradiction*, and *the law of identity*. To this day, it is still widely thought among amateur philosophers that these three laws have a special pride of place in logic. However, from the standpoint of modern logic, there is nothing particularly special about these three laws. They have no more claim to importance than, say, *modus ponens* or deMorgan's laws. They are just three amongst many other important laws in logic.

<sup>24</sup>To see this, it might be helpful to consider an analogy between the rules of logical inference and the rules of grammar. Each of us possesses an ability to tell when a sentence is grammatically correct according to the rules of grammar for our native natural language. It's

Bearing this distinction in mind, we can now inquire as to what kind of logical knowledge is required for upholding a system of knowledge. Here I would like to propose that knowledge that is explicitly *about* the laws of logic is not a requirement for a subject to have knowledge of other subject matters. (I am thus opposed to one of the central pillars of the presuppositional approach to foundationalism.) After all, most human beings never learn much about logic, and the few of us who do generally don't begin until we attend university. Yet it is absurd to think that the logically-uneducated and our pre-university selves could never have knowledge of other ordinary things. Not only that, but the historical development of explicit theories of logical truth and logical validity turned out to be hard-won intellectual achievements. Throughout ancient and early modern times, the specialists in logic adhered to theories that were ultimately defective (albeit, still highly serviceable). We did not achieve our modern understanding of logic until the turn of the 20th century, and even then, the classical systems that are now widely endorsed have well-known shortcomings.<sup>25</sup> The point is, theorizing *about* logic is *hard*. To attain knowledge of the laws that underwrite the plurality of logical truths and valid inferences is hard. Such a demanding task is not a requirement for knowing the whole lot of our pedestrian beliefs.

With that said, we should still acknowledge how the skills involved in logical reasoning are integral to a system of knowledge. These skills take on a particular importance in the context of foundationalism. For according to foundationalism, many of our beliefs are non-foundational, and so they inherit their justification from inference from other beliefs. And whenever we infer a non-foundational belief on the basis of a deductive argument, we make use of our skills in logical reasoning. In fact, I take it that these skills are necessary for knowledge when it comes to this kind of belief. For if a belief is based on an invalid attempt at deductive inference, then it isn't justified or known. Or if a belief of mine happens to be entailed by my other beliefs, but I lack the capacity to recognize this (and I have no other justification for it), then that belief isn't knowledge. So in short, I must be able to *use logical reasoning* in order to obtain non-foundational knowledge based on deductive argumentation. If, say, I am to know that  $Q$  by an inference from my knowledge *that if  $P$  then  $Q$  and  $P$* , then I must have some instinctive ability to recognize this instance of *modus ponens* as valid. However, as I argued in the last paragraph, this ability does not need to be accompanied by any explicit belief or knowledge *about* modus ponens.

So we see that there is at least one way in which the rules and principles of

---

easy to *use* the rules of grammar; even a child can do it. However, as any linguist will tell you, it is an entirely different matter to give an explicit representation of what those grammatical rules are. The task of explicitly stating those rules remains a difficult area of research for modern syntacticians.

<sup>25</sup>Our modern understanding is indebted to the revolutionary work of Gottlob Frege in the late 19th century. The system he developed is now known as 'classical' logic, even though it is less than a few centuries old. As for the 'well-known shortcomings', the classical systems that Frege produced do not sit well with the phenomena of vagueness, empty-names, and paradox. What to make of these shortcomings is still an active topic of dispute for contemporary philosophers of logic.



logic are relevant to foundationalist systems of knowledge—albeit, it’s logical *skill*, not *knowledge of logic as a subject matter*, that has primary importance. I would now like to close this section by mentioning another respect in which logic is relevant. It seems highly plausible that a logically adept subject could come to know some logical truths without inferring them from other beliefs. For example, I could use my skills in logical reasoning to recognize that *either it is snowing or it is not snowing* must be true, no matter the current state of the weather. Once I recognize this, this belief of mine would not rest on any other belief for its justification. It would thereby count as foundational knowledge.<sup>26</sup> Thus my skill in logical reasoning can also serve as a source for some kinds of foundational knowledge when the belief in question is a logical truth.<sup>27</sup>

### 3.2 The reliability of reason and our rational faculties

We turn now to the other proposal for our ultimate presupposition and foundational belief: the belief in the reliability of one’s own rational faculties. It is common to find this belief being touted as an ultimate presupposition, and it is not hard to see why. For if our rational faculties weren’t reliable, then we couldn’t trust them. And if we couldn’t trust our rational faculties, then since *every* belief is a product of these faculties, it follows that we couldn’t trust *any* of our beliefs. So if our rational faculties weren’t reliable, then we couldn’t have any justified beliefs or knowledge. And once again, this thought is often taken to support the idea that our intellectual starting point should concern the reliability of our rational faculties. It’s then claimed that this *belief*, that our rational faculties are reliable, should serve as our foremost foundational belief.

I hope that, by now, I no longer need to explain the fallacy in this line of reasoning. It is the same sort of fallacy as I’ve pointed out twice before. We can accept that *having* reliable rational faculties is a requirement for justified beliefs; but it does not follow that a *belief in the reliability of one’s own faculties* is a requirement for justified beliefs. So these considerations simply do not support any thesis about the content of our foundational knowledge. Anyway, we do not need to dwell on this by-now-familiar point for too long. Suffice to say that our discussion of this belief will not have any straightforward connection to foundationalism.

Nonetheless, I would still like to say a few words about this belief, *that our rational faculties are reliable*, since it is often taken to be of *utmost importance*

<sup>26</sup>This claim calls for a point of clarification. My coming to know that *either it is snowing or it is not snowing* will require some *inferential activity* on my part. I may have to reason “well let  $S =$  it is snowing, and suppose not- $(S$  or not- $S)$ . Then it follows that not- $S$  and not-not- $S$  (by de Morgan’s), which is self-contradictory. So it must be that  $S$  or not- $S$ .” By the end of all of this reasoning, I have discharged all of my assumptions, so my conclusion does not rest on any further *belief* of mine. This calls for another point of clarification. A foundational belief may be arrived at by *inferential activity*. But to be foundational, its justification cannot rest on *inference from other beliefs*.

<sup>27</sup>The primary example of this will be beliefs like *either it is snowing or not snowing*, where they are *instances* of logical laws, but not *about* logical laws. How we come to justify our beliefs *about* the logical laws is an entirely different matter, which I will not explore here, because it is complicated.

for a rational being to justify this belief. Well, can I justify it? Or much less, can I even explain what it means?

Well first things first, let's try to get clear on what this belief amounts to. On the one hand, this belief invokes the notion of *reliability*. This concept has a standard definition. In philosophical parlance, we say that a faculty or process is *reliable* if it has a tendency to produce mostly true beliefs; the ratio of true outputs over false outputs is relatively high.<sup>28</sup>

The other component of this belief invokes "our rational faculties" or "reason." Now what are these things? Well, we can describe these things well enough: our rational faculties are *whatever faculties are responsible for forming our beliefs*; and *reason* simply refers to *whatever processes, rules, or methods we use to form our beliefs*. But these roundabout descriptions don't really tell us anything about what these faculties, processes, rules and methods actually are. And what are they, exactly? Is there a simple, unifying definition that encompasses all of our rational faculties, processes, rules and methods? Can we identify a sort of *ur*-faculty or *ur*-method?

I would like to suggest that the search for a monolithic, all-encompassing, "umbrella" understanding of *reason* or *our rational faculties* is a search that will end in vain. *Reason* is not one big monolithic thing, and we have many diverse rational faculties. Just to give a potted survey, our rational faculties include our sensory perception, our memory, our logical abilities, our mathematical abilities, our linguistic abilities, our faculty for imagining possible scenarios, and our ability to receive and dispense information through testimony. Each of these can be subdivided into more specialized skills. As for reason, this would include all the various rules of deductive inference, all the various methods for extending our knowledge through inductive inference, all of the different methods for developing hypotheses and the procedures for testing them, and all of the various heuristics we rely on to decide whom to trust and whom not to trust. The point is, reasoning is not a uniform thing. It encompasses a plurality of processes, rules, and methods.

In light of this fact, how should we approach the question of whether our rational faculties and reasoning are reliable? I propose that we should resist the urge to try to answer it once and for all, in one fell swoop, without heeding the details of each faculty and process. For it is inevitable that each specific faculty and process will have varying levels of reliability. I can be a highly reliable mathematical reasoner while being a highly unreliable consumer of the news. What we should do instead is to investigate the reliability and justification of each of the various faculties and types of reasoning on their own, before we make any sweeping claims about the reliability of our faculties *in general*. And because there are so many different highly specific skills and reasoning processes that are involved in each of our various intellectual pursuits, this isn't a job that

---

<sup>28</sup>Actually, this definition of 'reliable' is ambiguous and vague in many respects. What is the minimum ratio of true beliefs over false beliefs that a faculty must meet to count as reliable? What are the circumstances that are relevant for assessing reliability? Must a faculty be reliable in all possible environments, or just the ones that we were evolved or created to inhabit?

can be done quickly. The self-reflective investigation of the methodology of each of the sciences, of mathematics, of logic, of history, and of philosophy are each incredibly large topics on their own. Each discipline spends an enormous amount of time and resources to study their own methodologies.

So where does this leave us? At best, the belief *in the reliability of our own rational faculties and reason* should function as a high level generalization over a multitude of assessments of more specific faculties, processes, rules, and methods. As such, it is not suited to be a foundational belief. What we should do, instead, is assess each of our faculties, rules, and methods individually, and treat their justification on a case-by-case basis. There is a sense in which I will be engaged in a small part of this project throughout the course of this book. For in the following chapters, I will investigate the justification and knowledge that comes from *sense perception, induction, memory, and testimony*. Obviously I cannot cover *all* (or even *most*) facets and types of reason, because much of that work is better left to more specialized fields. But that still leaves room for a number of foundational questions for philosophy, and that is where I hope to conduct my enquiry.