THEORIES of EPISTEMIC JUSTIFICATION

There are several different views as to what entails justification, mostly focusing on the question "How sure do we need to be that our beliefs correspond to the actual world?" Different theories of justification require different amounts and types of evidence before a belief can be considered justified. Popular theories of justification include:

- **Foundationalism** – Basic beliefs justify other, non-basic beliefs.
- **Coherentism** – Beliefs are justified if they cohere with other beliefs a person holds, each belief is justified if it coheres with the overall system of beliefs.
-
- **Internalism** – The believer must be able to justify a belief through internal knowledge. [Related: Evidentialism – Beliefs depend solely on the evidence for them, so the believer has to have access to this evidence to be justified in believing that something is the case.]
- **Externalism** – Outside sources of knowledge can be used to justify a belief. [Related: **Reliabilism** or 'Proper-functionalism' [= 'Reformed epistemology']: Beliefs are 'warranted' by proper cognitive function, so the believer does NOT need to be able to to justify a belief through internal knowledge.]

2. What is Justification?

When we discuss the nature of justification, we must distinguish between two different issues: First, what do we *mean* when we use the word 'justification'? Second, what *makes* beliefs justified? It is important to keep these issues apart because a disagreement on how to answer the second question will be a mere verbal dispute, if the disagreeing parties have different concepts of justification in mind. So let us first consider what we might mean by 'justification' and then move on to the non-definitional issues.[9]

2.1 Deontological and Non-Deontological Justification

How is the term 'justification' used in ordinary language? Here is an example: Tom asked Martha a question, and Martha responded with a lie. Was she justified in lying? Jane thinks she was, for Tom's question was an inappropriate one, the answer to which was none of Tom's business. What might Jane mean when she thinks that Martha was justified in responding with a lie? A natural answer is this: She means that Martha was *under no obligation* to refrain from lying. Due the inappropriateness of Tom's question, it wasn't Martha's *duty* to tell the truth. This understanding of justification, commonly labeled *deontological*, may be defined as follows: *S* is justified in doing *x* if and only if *S* is not obliged to refrain from doing *x*.[10]

Suppose, when we apply the word justification not to actions but to beliefs, we mean something analogous. In that case, the term 'justification' as used in epistemology would have to be defined this way:

Deontological Justification (DJ)
S is justified in believing that *p* if and only if *S* believes that *p* while it is not the case that *S* is obliged to refrain from believing that *p*.[11]

What kind of obligations are relevant when we wish to assess whether a *belief*, rather than an action, is justified or unjustified? Whereas when we evaluate an action, we are interested in assessing the action from either a moral or a prudential point of view, when it comes to beliefs, what matters is the pursuit of *truth*. The relevant kinds of obligations, then, are those that arise when we aim at having true beliefs. Exactly what, though, must we do in the pursuit of this aim? According to one answer, the one favored by evidentialists, we ought to believe in accord with our evidence. For this answer to be helpful, we need an account of what our evidence consists of. According to another answer, we ought to follow the correct epistemic norms. If this answer is going to help us figure out what obligations the truth-aim imposes on us, we need to be given an account of what the correct epistemic norms are.[12]

The deontological understanding of the concept of justification is common to the way philosophers such as Descartes, Locke, Moore and Chisholm have thought about justification. Today, however, the dominant view is that the deontological understanding of justification is unsuitable for the purposes of epistemology. Two chief objections have been raised against conceiving of justification deontologically. First, it has been argued that DJ presupposes that we can have a sufficiently high degree of control over our beliefs. But beliefs are akin not to actions but rather things such as digestive processes, sneezes, or involuntary blinkings of the eye. The idea is that beliefs simply arise in or happen to us. Therefore, beliefs are not suitable for deontological evaluation.[13] To this objection, some advocates of DJ have replied that lack of control over our beliefs is no obstacle to using the term ‘justification’ in its deontological sense.[14] Others have argued that it's a mistake to think that we can control our beliefs any less than our actions.[15]

Non-Deontological Justification (NDJ)
S is justified in believing that *p* if and only if *S* believes that *p* on a basis that properly probabilifies *S*'s belief that *p*.

If we wish to pin down exactly what probabilification amounts to, we will have to deal with a variety of tricky issues.[18] For now, let us just focus on the main point. Those who prefer NDJ to DJ would say that probabilification and deontological justification can diverge: it's possible for a belief to be deontologically justified without being properly probabilified. This is just what cases involving benighted cultures or cognitively deficient subjects are supposed to show.[19]

2.2 Evidence vs. Reliability

What *makes* justified beliefs justified? According to evidentialists, it is the possession of evidence. What is it, though, to possess evidence for believing that *p*? Some evidentialists would say it is to be in a mental state that represents *p* as being true. For example, if the coffee in your cup tastes sweet to you, then you have evidence for believing that the coffee is sweet. If you feel a throbbing pain in your head, you have evidence for believing that you have a headache. If you have a memory of having had cereal for breakfast, then you have evidence for a belief about the past: a belief about what you ate when you had breakfast. And when you clearly "see" or "intuit" that the proposition

"If Jack had more than four cups of coffee, then Jack had more than three cups of coffee" is true, then you have evidence for believing that proposition. In this view, evidence consists of perceptual, introspective, memorial, and intuitional experiences, and to possess evidence is to have an experience of that kind. So according to this evidentialism, what makes you justified in believing that *p* is your having an experience that represents *p* as being true.

Many reliabilists, too, would say that the experiences mentioned in the previous paragraph matter. However, they would deny that justification is solely a matter of having suitable experiences. Rather, they hold that a belief is justified if, and only if, it results from cognitive origin that is reliable: an origin that tends to produce true beliefs and therefore properly probabilifies the belief. Reliabilists, then, would agree that the beliefs mentioned in the previous paragraph are justified. But according to a standard form of reliabilism, what makes them justified is not the possession of evidence, but the fact that the types of processes in which they originate — perception, introspection, memory, and rational intuition — are reliable.

2.3 Internal vs. External

In contemporary epistemology, there has been an extensive debate on whether justification is internal or external. Internalists claim that it is internal; externalists deny it. How are we to understand these claims?

To understand what the internal-external distinction amounts to, we need to bear in mind that, when a belief is justified, there is something that *makes* it justified. Likewise, if a belief is unjustified, there is something that *makes* it unjustified. Let's call the things that make a belief justified or unjustified J-factors. The dispute over whether justification is internal or external is a dispute about what the J-factors are.

Among those who think that justification is internal, there is no unanimity on how to understand the concept of internality. We can distinguish between two approaches. According to the first, justification is internal because we enjoy a special kind of access to J-factors: they are *always* recognizable on reflection.[20] Hence, assuming certain further premises (which will be mentioned momentarily), justification itself is always recognizable on reflection.[21] According to the second approach, justification is internal because J-factors are always mental states.[22] Let's call the former *accessibility internalism* and the latter *mentalist internalism*. Externalists deny that J-factors meet either one of these conditions.

Evidentialism is typically associated with internalism, and reliabilism with externalism.[23] Let us see why. Evidentialism says, at a minimum, two things:

E1 Whether one is justified in believing *p* depends on one's evidence regarding *p*.

E2 One's evidence consists of one's mental states.

By virtue of E2, evidentialism is obviously an instance of mentalist internalism.

Whether evidentialism is also an instance of accessibility internalism is a more complicated issue. The conjunction of E1 and E2 by itself implies nothing about the recognizability of justification.

Recall, however, that in Section 1.1 we distinguished between TK and NTK: the traditional and the nontraditional approach to the analysis of knowledge and justification. TK advocates, among which evidentialism enjoys widespread sympathy, tend to endorse the following two claims:

Luminosity
One's own mind is cognitively luminous: Relying on introspection, one can always recognize on reflection what mental states one is in.[24]

Necessity
a priori recognizable, necessary principles say what is evidence for what.[25] Relying on *a priori* insight, one can therefore always recognize on reflection whether one's mental states are evidence for *p*.[26]

2.4 Why Internalism?

Why think that justification is internal? One argument for the internality of justification goes as follows: "Justification is deontological: it is a matter of duty-fulfillment. But duty-fulfillment is internal. Therefore, justification is internal." Another argument appeals to the brain-in-the-vat scenario we considered above: "Tim*'s belief that he has hands is justified in the way that Tim's is justifed. Tim* is internally the same as Tim and externally quite different. Therefore, internal factors are what justify beliefs." Finally, since justification resulting from the possession of evidence is internal justification, internalism can be supported by way of making a case for evidentialism. What, then, can be said in support of evidentialism? Evidentialists would appeal to cases in which a belief is reliably formed but not accompanied by any experiences that would qualify as evidence. They would say that it's not plausible to claim that, in cases like that, the subject's belief is justified. Hence such cases show, according to evidentialists, that a belief can't be justified unless it's supported by evidence.[31]

2.5 Why Externalism?

Why think that justification is external? To begin with, externalists about justification would point to the fact that animals and small children have knowledge and thus have justified beliefs. But their beliefs can't be justified in the way evidentialists conceive of justification. Therefore, we must conclude that the justification their beliefs enjoy is external: resulting not from the possession of evidence but from origination in reliable processes. And second, externalists would say that what we want from justification is the kind of objective probability needed for knowledge, and only external conditions on justification imply this probability. So justification has external conditions.[32]

3. The Structure of Knowledge and Justification

The debate over the structure of knowledge and justification is primarily one among those who hold that knowledge requires justification. From this point of view, the structure of knowledge derives from the structure of justification. We will, therefore, focus on the latter.

3.1 Foundationalism

According to foundationalism, our justified beliefs are structured like a building: they are divided into a foundation and a superstructure, the latter resting upon the former. Beliefs belonging to the foundation are *basic*. Beliefs belonging to the superstructure are *nonbasic* and receive justification from the justified beliefs in the foundation.[33]

For a foundationalist account of justification to be plausible, it must solve two problems. First, by virtue of exactly what are basic beliefs justified? Second, how do basic beliefs justify nonbasic beliefs? Before we address these questions, let us first consider the question of what it is that makes a justified belief basic in the first place. Once we have done that, we can then move on to discuss by virtue of what a basic belief might be justified, and how such a belief might justify a nonbasic belief.

According to one approach, what makes a justified belief basic is that it doesn't receive its justification from any other beliefs. The following definition captures this thought:

Doxastic Basicality (DB)
S's justified belief that *p* is basic if and only if *S*'s belief that *p* is justified without owing its justification to any of *S*'s other beliefs.

Let's consider what would, according to DB, qualify as an example of a basic belief. Suppose you notice (for whatever reason) someone's hat, and you also notice that that hat looks blue to you. So you believe

(B) It appears to me that that hat is blue.

Unless something very strange is going on, (B) is an example of a justified belief. DB tells us that (B) is basic if and only if it does not owe its justification to any other beliefs of yours. So if (B) is indeed basic, there might be some item or other to which (B) owes its justification, but that item would not be another belief of yours. We call this kind of basicality 'doxastic' because it makes basicality a function of how your doxastic system (your belief system) is structured.

Let us turn to the question of where the justification that attaches to (B) might come from, if we think of basicality as defined by DB. Note that DB merely tells us how (B) is *not* justified. It says nothing about *how* (B) is justified. DB, therefore, does not answer that question. What we need, in addition to DB, is an account of *what it is* that justifies a belief such as (B). According to one strand of foundationalist thought, (B) is justified because it can't be false, doubted, or corrected by others. So (B) is justified because (B) carries with it an *epistemic privilege* such as infallibility, indubitability, or incorrigibility.[34] The idea is that (B) is justified by virtue of its intrinsic nature, which makes it possess some kind of an epistemic privilege.

Note that (B) is not a belief about the hat. Rather, it's a belief about how the hat *appears* to you. So (B) is an introspective belief about a perceptual experience of yours. According to the thought we are considering here, a subject's basic beliefs are made up of introspective beliefs about the subject's own mental states, of which perceptual experiences make up one subset. Other mental states about which a subject can have basic beliefs include such things as having a headache, being tired, feeling pleasure, or having a desire for a cup of coffee. Beliefs about external objects do not and indeed

cannot qualify as basic, for it is impossible for such beliefs to own the kind of epistemic privilege needed for the status of being basic.

3.2 Coherentism

Foundationalism says that knowledge and justification are structured like a building, consisting of a superstructure that rests upon a foundation. According to coherentism, this metaphor gets things wrong. Knowledge and justification are structured like a *web* where the strength of any given area depends on the strength of the surrounding areas. Coherentists, then, deny that there are any basic beliefs. As we saw in the previous section, there are two different ways of conceiving of basicality. Consequently, there are two corresponding ways of construing coherentism: as the denial of doxastic basicality or as the denial of epistemic basicality. Consider first coherentism as the denial of doxastic basicality:

Doxastic Coherentism
Every justified belief receives its justification from other beliefs in its epistemic neighborhood.

Let us apply this thought to the hat example we considered in Section 3.1. Suppose again you notice someone's hat and believe

(H) That hat is blue.

Let's agree that (H) is justified. According to coherentism, (H) receives its justification from other beliefs in the epistemic vicinity of (H). They constitute your evidence or your reasons for taking (H) to be true. Which beliefs might make up this set of justification-conferring neighborhood beliefs?

We will consider two approaches to answering this question. The first is known as *inference to the best explanation*. Such inferences generate what is called *explanatory coherence*.[39] According to this approach, we must suppose you form a belief about the way the hat appears to you in your perceptual experiences, and a second belief to the effect that your perceptual experience, the hat's looking blue to you, is best explained by the assumption that (H) is true. So the relevant set of beliefs is the following:

(1) I am having a visual experience (E): the hat looks blue to me.

(2) My having (E) is best explained by assuming that (H) is true.

There are of course alternative explanations of why you have (E). Perhaps you are hallucinating that the hat is blue. Perhaps an evil demon makes the hat look blue to you when in fact it is red. Perhaps you are the sort of person to whom hats always look blue. An explanatory coherentist would say that, compared with these, the hat's actual blueness is a superior explanation. That's why you are justified in believing (H). Note that an explanatory coherentist can also explain the *lack* of justification. Suppose you remember that you just took a hallucinatory drug that makes things look blue to you. That would prevent you from being justified in believing (H). The explanatory coherentist can account for this by pointing out that, in the case we are considering now, the truth

of (H) would not be the *best* explanation of why you are having experience (E). Rather, your having taken the hallucinatory drug would be an explanation at least as good as the assumption as (H) is true. That's why, according to the explanatory coherentist, in this variation of our original case you wouldn't be justified in believing (H).

3.3 Why Foundationalism?

The main argument for foundationalism is called the *regress argument*. It's an argument from elimination. With regard to every justified belief, B_1, the question arises of where B_1's justification comes from. If B_1 is not basic, it would have to come from another belief, B_2. But B_2 can justify B_1 only if B_2 is justified itself. If B_2 is basic, the justificatory chain would end with B_2. But if B_2 is not basic, we need a further belief, B_3. If B_3 is not basic, we need a fourth belief, and so forth. Unless the ensuing regress terminates in a basic belief, we get two possibilities: the regress will either *loop back* to B_1 or continue *ad infinitum*. According to the regress argument, both of these possibilities are unacceptable. Therefore, if there are justified beliefs, there must be basic beliefs.[42]

3.4 Why Coherentism?

Coherentism is typically defended by attacking foundationalism as a viable alternative. To argue against privilege foundationalism, coherentists pick an epistemic privilege they think is essential to foundationalism, and then argue that either no beliefs, or too few beliefs, enjoy such a privilege. Against experiential foundationalism, different objections have been advanced. One line of criticism is that perceptual experiences don't have propositional content. Therefore, the relation between a perceptual belief and the perceptual experience that gives rise to it can only be causal. Consider again, however, the hat example from above. When you see the hat and it looks blue to you, doesn't your visual experience — its looking blue to you — have the propositional content *that the hat is blue*? It would seem it does. If it does, there seems to be no reason to deny that your perceptual experience can play a justificatory role.[45]

4. Sources of Knowledge and Justification

Beliefs arise in people for a wide variety of causes. Among them, we must list psychological factors such as desires, emotional needs, prejudice, and biases of various kinds. Obviously, when beliefs originate in sources like these, they don't qualify as knowledge even if true. For true beliefs to count as knowledge, it is necessary that they originate in sources we have good reason to consider reliable. These are perception, introspection, memory, reason, and testimony. Let us briefly consider each of these.

4.1 Perception

Our perceptual faculties are our five senses: sight, touch, hearing, smelling, and tasting. We must distinguish between an experience that can be classified as *perceiving* that *p* (for example, seeing that there is coffee in the cup and tasting that it is sweet), which entails that *p* is true, and a perceptual experience in which it seems to us as though *p*, but where *p* might be false. Let us refer to this latter kind of experience as *perceptual seemings*. The reason for making this distinction lies in the fact that perceptual experience is fallible. The world is not always as it appears to us in our

perceptual experiences. We need, therefore, a way of referring to perceptual experiences in which *p* seems to be the case that allows for the possibility of *p* being false. That's the role assigned to perceptual seemings. So some perceptual seemings that *p* are cases of perceiving that *p*, others are not. When it looks to you as though there is a cup of coffee on the table and in fact there is, the two states coincide. If, however, you hallucinate that there is a cup on the table, you have perceptual seeming that *p* without perceiving that *p*.

4.2 Introspection

Introspection is the capacity to inspect the, metaphorically speaking, "inside" of one's mind. Through introspection, one knows what mental states one is in: whether one is thirsty, tired, excited, or depressed. Compared with perception, introspection appears to have a special status. It is easy to see how a perceptual seeming can go wrong: what looks like a cup of coffee on the table might be just be a clever hologram that's visually indistinguishable from an actual cup of coffee. But could it be possible that it introspectively seems to me that I have a headache when in fact I do not? It is not easy to see how it could be. Thus we come to think that introspection has a special status. Compared with perception, introspection seems to be privileged by virtue of being less error prone. How can we account for the special status of introspection?

First, it could be argued that, when it comes to introspection, there is no difference between appearance and reality; therefore, introspective seemings are necessarily successful introspections. According to this approach, introspection is infallible. Alternatively, one could view introspection as a source of certainty. Here the idea is that an introspective experience of *p* eliminates all possible doubt as to whether *p* is true. Finally, one could attempt to explain the specialness of introspection by examining the way we respond to first-person reports: typically, we attribute a special authority to such reports. According to this approach, introspection is incorrigible. Others are not, or at least not typically, in a position to correct first-person reports of one's own mental states.

Introspection reveals how the world appears to us in our perceptual experiences. For that reason, introspection has been of special interest to foundationalists. Perception is not immune to error. If certainty consists in the absence of all possible doubt, perception fails to yield certainty. Hence beliefs based on perceptual experiences cannot be foundational. Introspection, however, might deliver what we need to find a firm foundation for our beliefs about external objects: at best outright immunity to error or all possible doubt, or perhaps more modestly, an epistemic kind of directness that cannot be found in perception.

Is it really true, however, that, compared with perception, introspection is in some way special? Critics of foundationalism have argued that introspection is certainly not infallible. Might one not confuse an unpleasant itch for a pain? Might I not think that the shape before me appears circular to me when in fact it appears slightly elliptical to me? If it is indeed possible for introspection to mislead, then it is hard to see why introspection should eliminate all possible doubt. Yet it isn't easy to see either how, if one clearly and distinctly feels a throbbing headache, one could be mistaken about that. Introspection, then, turns out to be a mysterious faculty. On the one hand, it does not seem to be in general an infallible faculty; on the other hand, when looking at appropriately described specific cases, error does seem impossible.[48]

4.3 Memory

Memory is the capacity to retain knowledge acquired in the past. What one remembers, though, need not be a past event. It may be a present fact, such as one's telephone number, or a future event, such as the date of the next elections. Memory is, of course, fallible. Not every instance of taking oneself to remember that *p* is an instance of actually remembering that *p*. We should distinguish, therefore, between remembering that *p* (which entails the truth of *p*) and *seeming* to remember that *p* (which does not entail the truth of *p*).

One issue about memory concerns the question of what distinguishes memorial seemings from perceptual seemings or mere imagination. Some philosophers have thought that having an image in one's mind is essential to memory, but that would appear to be mistaken. When one remembers one's telephone number, one is unlikely to have an image of one's number in one's mind. The distinctively epistemological questions about memory are these: First, what makes memorial seemings a source of justification? Is it a necessary truth that, if one has a memorial seeming that *p*, one has thereby prima facie justification for *p*? Or is memory a source of justification only if, as coherentists might say, one has reason to think that one's memory is reliable? Or is memory a source of justification only if, as externalists would say, it is in fact reliable? Second, how can we respond to skepticism about knowledge of the past? Memorial seemings of the past do not guarantee that the past is what we take it to be. We think that we are a bit older than just five minutes, but it is logically possible that the world sprang into existence just five minutes ago, complete with our dispositions to have memorial seemings of a more distant past and items such as apparent fossils that suggest a past going back millions of years. Our seeming to remember that the world is older than a mere five minutes does not entail, therefore, that it really is. Why, then, should we think that memory is a source of knowledge about the past?[49]

4.4 Reason

Some beliefs would appear to be justified solely by the use of reason. Justification of that kind is said to be *a priori*: prior to any kind of experience. A standard way of defining *a priori* justification goes as follows:

***A Priori* Justification**
S is justified *a priori* in believing that *p* if and only if *S*'s justification for believing that *p* does not depend on any experience.

Beliefs that are true and justified in this way (and not somehow "gettiered") would count as instances of *a priori* knowledge.[50]

What exactly counts as experience? If by ‘experience’ we mean just *perceptual* experiences, justification deriving from introspective or memorial experiences would count as *a priori*. For example, I could then know *a priori* that I'm thirsty, or what I ate for breakfast this morning. While the term ‘*a priori*’ is sometimes used in this way, the strict use of the term restricts *a priori* justification to justification derived *solely* from the use of reason. According to this usage, the word ‘experiences' in the definition above includes perceptual, introspective, and memorial experiences alike. On this narrower understanding, paradigm examples of what I can know on the basis of *a*

priori justification are conceptual truths (such as "All bachelors are unmarried"), and truths of mathematics, geometry and logic.

Justification and knowledge that is not *a priori* is called ‘a posteriori’ or ‘empirical’. For example, in the narrow sense of ‘*a priori*’, whether I'm thirsty or not is something I know empirically (on the basis of introspective experiences), whereas I know *a priori* that 12 divided by 3 is 4.

Several important issues arise about *a priori* knowledge. First, does it exist at all? Skeptics about apriority deny its existence. They don't mean to say that we have no knowledge of mathematics, geometry, logic, and conceptual truths. Rather, what they claim is that all such knowledge is empirical.

Second, if *a priori* justification is possible, exactly how does it come about? What *makes* a belief such as "All bachelors are unmarried" justified solely on the basis of reason? Is it an unmediated grasp of the truth of this proposition? Or does it consist of grasping that the proposition is *necessarily* true? Or is it the purely intellectual experience of "seeing" (with they "eye of reason") or "intuiting" that this proposition is true (or necessarily true)? Or is it, as externalists would suggest, the reliability of the cognitive process by which we come to recognize the truth of such a proposition?

Third, if *a priori* knowledge exists, what is its extent? *Empiricists* have argued that *a priori* knowledge is limited to the realm of the *analytic*, consisting of propositions of a somehow inferior status because they are not really "about the world". Propositions of a superior status, which convey genuine information about world, are labeled *synthetic*. *a priori* knowledge of synthetic propositions, empiricists would say, is not possible. *Rationalists* deny this. They would say that a proposition such as "If a ball is green all over, then it doesn't have black spots" is synthetic and knowable *a priori*.

A fourth question about the nature of *a priori* knowledge concerns the distinction between necessary and contingent truths. The received view is that whatever is known *a priori* is necessarily true, but there are epistemologists who disagree with that.[51]

4.5 Testimony

Testimony differs from the sources we considered above because it isn't distinguished by having its own cognitive faculty. Rather, to acquire knowledge of *p* through testimony is to come to know that *p* on the basis of someone's saying that *p*. "Saying that *p*" must be understood broadly, as including ordinary utterances in daily life, postings by bloggers on their web-logs, articles by journalists, delivery of information on television, radio, tapes, books, and other media. So, when you ask the person next to you what time it is, and she tells you, and you thereby come to know what time it is, that's an example of coming to know something on the basis of testimony. And when you learn by reading the *Washington Post* that the terrorist attack in Sharm el-Sheikh of July 22, 2005 killed at least 88 people, that, too, is an example of acquiring knowledge on the basis of testimony.

The epistemological puzzle testimony raises is this: Why is testimony a source of knowledge? An externalist might say that testimony is a source of knowledge if and only if it comes from a reliable source. But here, even more so than in the case of our faculties, internalists will not find that answer satisfactory. Suppose you hear someone saying '*p*'. Suppose further that person is in fact utterly reliable with regard to the question of whether *p* is the case or not. Finally, suppose you have no evidential clue whatever as to that person's reliability. Wouldn't it be plausible to conclude that, since that person's reliability is unknown to you, that person's saying '*p*' does not put you in a position to know that *p*? But if the reliability of a testimonial source is not sufficient for making it a source of knowledge, what else is needed? Thomas Reid suggested that, by our very nature, we accept testimonial sources as reliable and tend to attribute credibility to them unless we encounter special contrary reasons. But that's merely a statement of the attitude we in fact take toward testimony. What is it that makes that attitude reasonable? It could be argued that, in one's own personal experiences with testimonial sources, one has accumulated a long track record that can be taken as a sign of reliability. However, when we think of the sheer breadth of the knowledge we derive from testimony, one wonders whether one's personal experiences constitute an evidence base rich enough to justify the attribution of reliability to the totality of the testimonial sources one tends to trust. An alternative to the track record approach would be to declare it a necessary truth that trust in testimonial sources is justified. This suggestion, alas, encounters the same difficulty as the externalist approach to testimony: it does not seem we can acquire knowledge from sources the reliability of which is utterly unknown to us.[52]

6

How Can We Know about the External World?

You know that the earth is round, that penguins inhabit Antarctica, that trees shed their leaves in the fall, that you have a heart, and so on and so on. In other words, you know a lot about the "external world," including your own body.[1] That much is obvious.

Or is it? Consider the hypothesis that your entire life has been a remarkably vivid dream. Not only have you been dreaming the whole time, but the earth never existed. No penguins, trees, nothing like that. In fact, you don't even have a heart. You are a heartless android, lying comatose in a robot junkyard on a planet orbiting the star Kepler-11. "From the inside," things seem exactly the same to you: you seem to see this page, you seem to remember that penguins inhabit Antarctica, and so on, even though there is no page, and no Antarctica. So how can you know that this "android hypothesis" is false? That question can seem very difficult to answer, which suggests that you *can't* know that the android hypothesis is false.

This claim of ignorance might not seem so bad by itself, but once it is conceded, it is difficult to stop ignorance from spreading much more widely. Take, for example, one thing that you apparently know; namely, that penguins inhabit Antarctica. Now the claim that penguins inhabit Antarctica straightforwardly **entails** that the android hypothesis is false. If penguins inhabit Antarctica, you are *not* a dreaming android who lives in a penguin-free world. So, *if* you know that penguins inhabit Antarctica, you can perform an elementary logical inference and come to know that the android hypothesis is false. So, if you *can't* know that the android hypothesis is false, you don't know that penguins inhabit Antarctica. By the same **argument**,

1. This kind of knowledge is called **propositional knowledge**, or *factual knowledge*. See the introduction to Chapter 3, "What Is Knowledge?"

neither do you know that the earth is round, that trees shed their leaves in the fall, and so on. In short, if you can't know that the android hypothesis is false, you are completely ignorant about the external world; that is, *external world* **skepticism** is true. (A *skeptic* about some subject matter *M* is someone who denies that we have knowledge about *M*.)

Still, you might wonder whether even external world skepticism is worth worrying about. Suppose you're offered a choice between going on a roller coaster ride and entering the roller coaster simulator. The simulator is perfect: as far as excitement goes, it's just as good as the real thing, although you aren't really rattling down a narrow track at 100 mph. The choice doesn't seem to matter much (actually, you might even prefer the simulator on the grounds that it's much safer). Here, virtual reality is no worse than reality itself. Isn't that true in general? Why care whether you're a dreaming android? The thrills and spills of life would be the same in any case.

But this reaction is overly complacent. The dreaming android has no friends, has no mother who loves it, and has never accomplished anything—vividly dreaming that you are acing your final exams is not a way of doing well in school. Having friends, to say nothing of a mother who loves you, is a valuable thing. (Imagine discovering that someone whom you thought a faithful friend was just pretending.) So if you are a dreaming android, you are in a very unfortunate predicament—friendless, unloved, and unaccomplished. You should want to be reassured that you are *not* in this predicament. That is, you should want to *know* that you have friends, are loved, and so forth. If external world skepticism is true, reassurance that your life is not an empty sham is forever beyond your reach.

The readings in this chapter respond to the threat of skepticism about the external world. (One exception, as we will see below, is the essay by Rae Langton.) Before getting to the many different responses, it will help to set out the argument for the skeptical conclusion more precisely. And in the course of doing that, we will see how the argument is a particular instance of a general form of skeptical argument.

A General Skeptical Argument

Let a *skeptical hypothesis* be a hypothesis according to which the world is different from how you take it to be. We have already seen one skeptical hypothesis, according to which you're a dreaming android and the earth never existed. There are other similar skeptical hypotheses, the most famous of which is René Descartes's *demon hypothesis*: "some malicious demon of the utmost power and cunning has employed all his energies in order to deceive me." The contemporary version of the demon hypothesis is the *brain in a vat hypothesis*: you are a brain kept alive in a vat by some evil scientist and stimulated so that "from the inside" things seem exactly as if you see this page, and so on.

These hypotheses are *global* skeptical hypotheses; that is, if they are true, almost *nothing* you take yourself to know about the external world is true. But skeptical hypotheses can be more modest. Indeed, in an everyday situation in which you wonder whether you really are right to think that you left your laptop at home, you are entertaining a very modest skeptical hypothesis—that the world is very similar to the way you take it to be, except that your laptop is not at home. Philosophers have devised many other skeptical hypotheses that are intermediate in strength between global skeptical hypotheses and very modest skeptical hypotheses like the one just mentioned.

For instance, there is the *no other minds hypothesis*, according to which you are the only creature with a mind—everyone else behaves just *as if* they believe, feel, and perceive but are actually entirely mindless (see Saul Kripke's "Wittgenstein and Other Minds" in Chapter 5). And there is the *unexpected future hypothesis*, according to which the future is radically different from the past—if this hypothesis is true, bread will not nourish tomorrow, the sun will not rise tomorrow, and so on (see the introduction to Chapter 4, "How Can We Know about What We Have Not Observed?").

Now take a skeptical hypothesis SH, and any claim *p* that entails that SH is false. We can argue that you don't know *p* as follows:

1. If you know *p*, you can know that SH is false.

2. You can't know that SH is false.

So:

3. You don't know *p*.

For example, suppose you think (*p*) that your bike is where you left it, padlocked to a bike rack. Let SH be the modest skeptical hypothesis that your bike has been stolen. The claim *p* (your bike is where you left it) entails that SH (your bike has been stolen) is false, in other words that your bike has not been stolen. So we can argue that you don't know that your bike is where you left it as follows:

1†. If you know that your bike is where you left it, you can know that your bike has not been stolen.

2†. You can't know that your bike has not been stolen.

So:

3†. You don't know that your bike is where you left it.

This form of argument—If *P* then *Q*, it is not the case that *Q*, so it is not the case that *P*—is called **modus tollens**.

Let us look more carefully at **premise** 1. Suppose you know that all fish have gills and that Wanda is a fish. Now the statement that all fish have gills and Wanda is a fish entails that Wanda has gills. So you are now in a position to draw the conclusion that Wanda has gills from what you already know. And if you go ahead and do

that, it seems very plausible that you will end up *knowing* that Wanda has gills. In general, one way to extend our knowledge is to trace out the **logical consequences** of what we already know: this happens whenever someone proves a theorem in mathematics, for example. Put more precisely: if *p* entails (or logically implies) *q*, and you know *p*, then you are in a position to know *q*. This is (one version of) a principle called **closure**.

Given closure, premise 1 of the skeptical argument is true. Closure is difficult to deny and the argument is **valid**, so when faced with a skeptical argument of this form you have two options: deny premise 2 or accept the conclusion.[2]

We can generate an argument for external world skepticism by letting SH be a global skeptical hypothesis, say, Descartes's demon hypothesis, and letting *p* be any claim about the external world that entails that the demon hypothesis is false, say, that the earth is round:

1*. If you know that the earth is round, you can know that the demon hypothesis is false.

2*. You can't know that the demon hypothesis is false.

So:

3*. You don't know that the earth is round.

Again, assuming closure, there are two options: deny premise 2* or accept 3*, the skeptical conclusion.

Responses to External World Skepticism

The readings from David Hume, G. E. Moore, and Jonathan Vogel offer contrasting responses to external world skepticism.[3]

According to Hume, our senses provide scant evidence for hypotheses about the external world. In a paragraph omitted from the selection he writes: "'Tis impossible . . . that from the existence or any of the qualities of [perceptions], we can ever form any conclusion concerning the existence of [objects]." So Hume is a (rare) example of a real-life skeptic: he accepts the conclusion of the skeptical argument. His main concern is not so much to defend skepticism (which he thinks is pretty much unassailable) but rather to give a psychological explanation for why we think that there is an external world of familiar tables, chairs, penguins, and so on, even though we have no good reason for doing so.

2. In fact, despite the plausibility of closure, some philosophers deny it. A notable example is Robert Nozick: see his *Philosophical Explanations* (Harvard University Press, 1981), chapter 3. But none of the contributors to this chapter deny it. See "Analyzing the Arguments" at the end of this chapter.

3. You can find another response by Stewart Cohen at http://digital.wwnorton.com/introphilosophy2.

Vogel, in effect, directly replies to Hume. While Hume thinks that our senses cannot show us that we are not brains in vats or deceived by an evil demon, Vogel thinks otherwise, and so denies the second premise. Vogel argues that the "real world hypothesis"—that the earth is round, you have a head, are reading this book, and so on—provides a much better explanation of your "sensory experiences" (or, in Hume's terminology, "perceptions") than any global skeptical hypothesis. So you have a good reason to believe the real world hypothesis by an "inference to the best explanation."[4]

Moore is principally concerned to deny the conclusion of the skeptical argument, rather than to explain which premise is false. He tries to turn the tables on the skeptic by offering what he claims is a *proof* of the existence of things like tables and books. For example: here is a book (Moore holds up a copy of this book), here is another book (Moore holds up a copy of his own famous book on ethics, *Principia Ethica*), therefore books exist. Of course, the skeptic will not grant that this is a proof, on the grounds that Moore does not know the premises. But, as Moore points out, in ordinary life we take arguments of this sort "as absolutely conclusive proofs of certain conclusions." For instance, we allow that someone can prove that there are at least three misprints on a page from the premises "There's one misprint here, another there, and another here." And if we really can prove such things, we must have knowledge about the external world. Why isn't the skeptic just being unreasonable in rejecting Moore's proof?

Kantian Skepticism

Rae Langton's essay defends a limited but nonetheless fascinating form of skepticism, which she finds suggested by the work of the German philosopher Immanuel Kant (1724–1804). This Kantian skepticism (or "Kantian humility," as Langton calls it) is not external world skepticism. As both Moore and Langton note, Kant thought we had plenty of knowledge about the external world.

However, in his *Critique of Pure Reason*, Kant argues for another kind of skepticism: as he put it, we have no knowledge of "things in themselves." What did he mean by that? Langton offers an answer, and a defense of a kind of skepticism that she thinks is at least in the spirit of Kant's actual view. According to Langton, the physical sciences can only penetrate so far into reality: there is a layer further down that is in principle beyond their reach. If that is right, then although we can know that there are books, and that we have friends, ignorance of the fundamental nature of the world is part of the human condition.

4. For more on this kind of inference, see Gilbert Harman, "The Inference to the Best Explanation," in Chapter 4 of this anthology.

NOTES AND QUESTIONS

1. Consider the following argument:

   P1. I am wide awake.

   P2. If I am wide awake, then I am not dreaming.

   C. I am not dreaming.

   Can you use this argument to prove that you are not dreaming? If not, why not? Moore says he cannot prove that he is "not now dreaming." Are his reasons persuasive?

Jonathan Vogel (b. 1954)

Vogel is Professor of Philosophy at Amherst College and has made many contributions to skepticism and related topics in epistemology.

SKEPTICISM AND INFERENCE TO THE BEST EXPLANATION

More than two thousand years ago, the philosopher Zhuangzi posed the question: How can a man know that he is a man, rather than a butterfly dreaming that he is a man? Later, René Descartes asked how he could be sure that his sensory experience wasn't caused by an evil demon, who was bent on deceiving him.[1] And, today, you might consider the possibility that, instead of you really seeing this book, your brain is being stimulated by a computer to make it appear to you, falsely, that you're seeing this book. What reason do you have for thinking that this isn't happening to you now?

Thoughts like these raise one of the oldest and deepest problems in philosophy, the problem of skepticism about the external world. Skepticism is the sweeping and unsettling doctrine that we have no knowledge of the world around us. It presents a philosophical problem because it is supported by a simple, formidable line of thought, known as the *deceiver argument*:

1. Your sensory experiences could come about through ordinary perception, so that most of what you believe about the world is true. But your sensory experiences could also be caused deceptively, so that what you believe about the world is entirely false.

1. According to tradition, Zhuangzi (Chinese, fourth century BCE) is the author of the influential treatise that bears his name. René Descartes (French, 1596–1650) is often viewed as the founder of modern philosophy in the West. His best known work is *Meditations on First Philosophy* (1641), where Descartes raises the possibility that an evil demon is deceiving him (see the selection from the *Meditations* earlier in this chapter). Other selections from the *Meditations* are in Chapter 7 of this anthology.

2. You have no reason at all to believe that your sensory experiences arise in one way rather than the other.

3. Therefore, you have no knowledge of the world around you.[2]

But, of course you do have knowledge of the world. There must be something wrong with the deceiver argument. What, though? Premise 1 seems extremely plausible.[3] So, if the deceiver argument fails, either premise 2 must be false or else the conclusion 3 doesn't really follow from premise 2.

I think the conclusion does follow from premise 2, by way of a general principle about knowledge called the underdetermination principle. This principle says that if you are faced with two (or more) mutually exclusive hypotheses, and the information available to you gives you no reason to believe one rather than the other, then you don't know that either hypothesis is the case. For example, consider the claim that there is water on Mars. It might be that water formed there long ago and remains to this day. However, Mars is smaller than Earth and has a very different surface and atmosphere. Maybe the martian water, if any, was lost over the ages. In the absence of any information one way or the other—say, evidence provided by telescopes or by probes sent to Mars—a scientist who maintained or denied the presence of water on Mars would just be guessing. She wouldn't know that her hypothesis is true, precisely because she has no information favoring that hypothesis over its competitor. This is an illustration of how the underdetermination principle governs what counts as knowledge.[4]

2. There are subtle but significant questions about exactly how the argument ought to proceed, but I'm setting those aside. See Jonathan Vogel, "Skeptical Arguments," *Philosophical Issues* 14 (2004): 426–55. [Vogel's note.]

3. Some philosophers have rejected premise 1. J. L. Austin, *Sense and Sensibilia* (Oxford University Press, 1962), may have held such a view, and nowadays John McDowell, "Criteria, Defeasibility, and Knowledge," *Proceedings of the British Academy* 68 (1982): 455–79, and Timothy Williamson, *Knowledge and Its Limits* (Oxford University Press, 2000), might be read as denying premise 1, too. Another response regards the whole argument as somehow misconceived, perhaps because it involves hidden presuppositions about knowledge we need not accept or ought to reject. See the classic writings of Ludwig Wittgenstein, *On Certainty*, ed. G. E. M. Anscombe and G. H. Von Wright (Blackwell, 1975), and J. L. Austin, "Other Minds," *Proceedings of the Aristotelian Society*, supplementary vol. 20 (1946): 148–87, and more recently Michael Williams, *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism* (Blackwell, 1992), and Alex Byrne, "How Hard Are the Skeptical Paradoxes?" *Noûs* 38 (2004): 299–325. [Vogel's note.]

4. In my opinion, the underdetermination principle is perfectly correct. However, some philosophers have developed sophisticated views about knowledge that are inconsistent with it. *Relevant alternatives theorists* hold that, to know a proposition, a person may need to have reasons to reject some competitors to it, but she doesn't need to have reasons to reject *all* the competitors, in every case. Thus, the underdetermination principle doesn't hold in full generality, and the conclusion of the deceiver argument doesn't follow from premise 2 after all. Certain *reliabilist* theories of knowledge go further than the relevant alternatives theory. They deny that there is any fundamental connection between knowing a proposition and having reasons to believe that proposition (and to reject its competitors). In particular, advocates of these accounts hold that we can know propositions about the world, whether we have reasons to reject skeptical hypotheses or not. Thus they, too, will deny that the conclusion of the deceiver argument is supported by premise 2. I've offered various criticisms of the relevant alternatives and reliabilist approaches. See Jonathan Vogel, "The New Relevant Alternatives Theory," *Philosophical Perspectives* 13 (1999): 155–80; Jonathan Vogel, "Reliabilism Leveled," *Journal of Philosophy* 97 (2000): 602–23; Jonathan Vogel, "Externalism Resisted," *Philosophical Studies* 131 (2006): 729–42; and Jonathan Vogel, "Subjunctivitis," *Philosophical Studies* 134 (2006): 73–88. [Vogel's note.]

This principle matters for our purposes as follows. The deceiver argument confronts you with two competing hypotheses. One is that the world is the way it appears to be, the other is that you are the victim of massive sensory deception. If premise 2 is true, you have no reason to favor the first over the second or vice versa. It then follows by the underdetermination principle that you don't know that either hypothesis is true. In particular, you may believe that the world is the way it appears to be, but you don't know that it is. So, premise 2 leads to the conclusion, given the underdetermination principle.

Since premise 2 does support the conclusion, the only way to escape the deceiver argument is to deny premise 2. If premise 2 is false, then skepticism is incorrect because we really do possess some reason for believing that we aren't victims of sensory deception after all. Contemporary philosophers have advanced three principal proposals as to what such a reason might be. One is the *Moorean* view.[5] The Moorean maintains that a sensory experience has a distinctive character or content, and, other things being equal, your having such an experience justifies you in holding a corresponding belief. For instance, suppose you seem to see a tree. The experience you have makes you justified in believing that there is a tree before you. But if there really is a tree before you, it can't be that you're deceived by a nefarious computer when you seem to see a tree.[6] So, according to the Moorean, your sensory experience gives you reason to believe that there is a tree before you, which in turn gives you a reason to believe that you're not deceived by a computer when you seem to see a tree before you. The same goes for any other sensory experience you may have. In general, your sensory experience gives you a reason to believe that you're not deceived by a computer, and premise 2 of the skeptical argument is false.

The Moorean approach is simple and decisive, but it strikes many as unsatisfactory. The trouble can be brought out by an analogy. Suppose you use the gas gauge in your car to tell you how much gas is in the fuel tank. You take a look and determine that the tank is half full. But surely, you can't then infer that, because the gauge says that the tank is half full, the gauge must be reading correctly! The Moorean seems to be up to something just as dubious. You can't use a particular sensory experience to establish that the experience itself isn't deceptive, any more than you can use a gauge to establish that the gauge itself isn't deceptive (i.e., to establish that it is reading correctly).

One reaction to this situation is to think that experience can't give you *any* reason for believing that you're not deceived by a computer. If you do have some basis for this belief, the basis can't be experience; your belief must be justified nonexperientially, or *a priori*. This is a second strategy for denying premise 2 of the deceiver argument. However, what is being suggested is somewhat hard to fathom. If you have nonexperiential grounds for thinking that your sensory experience isn't deceptive, it seems that

5. Named for the British philosopher G. E. Moore. A classic statement of Moore's view is his "Proof of an External World," *Proceedings of the British Academy* 25 (1939): 273–300 [see earlier in this chapter]. For a more recent version, see James Pryor, "The Skeptic and the Dogmatist," *Noûs* 34 (2000): 517–49. [Vogel's note.]

6. To be more explicit: If the computer is deceiving you, then there is no tree before you. But, then, if there is a tree before you, the computer isn't deceiving you. [Vogel's note.]

the source of those grounds would have to be reason. But how could reasoning—just thinking about things—establish whether your sensory experience is caused by a computer or not? In the face of such worries, some philosophers maintain that believing that one's sensory experience isn't generally deceptive is simply part of what it is to *be rational*.[7] But what would make that so? And how is this different from just coming up with something nice to say ("it's rational") about an assumption that we make blindly and without any reason whatsoever?

These remarks about the Moorean and a priori replies to the deceiver argument leave a great deal unsaid. But let's move on and consider a third sort of reply, which I'll call *explanationism*.[8] The idea behind this approach is that very often we are justified in adopting hypotheses because they do a good job of explaining the data we have. Here is an illustration. Suppose that a patient, Roger, goes to see Dr. G, his physician. Roger is sneezing, he has moist eyes, and his condition recurs at a certain time of the year. Roger's having an allergy explains these symptoms. There are other possible explanations. It could be that Roger has had a series of colds over the years or that he has a chronic respiratory infection that lies dormant for much of the time. But if Roger's having an allergy explains his symptoms extremely well, and his having any of these other conditions would explain his symptoms much less well, then Dr. G has good reason to reject those other diagnoses. Dr. G would be in a position to conclude that Roger's symptoms are due to an allergy. Dr. G's arriving at a diagnosis in this way is an example of what is known as *inference to the best explanation*. In general, if one hypothesis provides a significantly better explanation of the available evidence than its competitors do, that is a reason to accept the explanation and to reject the competitors.

This point carries over to the deceiver argument as follows. One explanation for the occurrence of your sensory experiences is that the world is actually the way you think it is and you're perceiving it properly. For example, we normally suppose that you have a visual experience of the ocean, because you are seeing the ocean. Similarly, you have a visual experience of sand dunes, because you are really seeing sand dunes, and so on. Call the collection of your ordinary beliefs about the world the "real world hypothesis." Skepticism emerges as a problem because there are alternative explanations of how your sensory experiences as a body come about. Call these "skeptical hypotheses." If the real world hypothesis explains the occurrence of your sensory experiences as a body better than skeptical hypotheses do, then, by inference to the best explanation, you have good reason to accept the real world hypothesis and to reject the skeptical alternatives. This outcome contradicts premise 2 of the deceiver argument, which says that you have no reason to believe one thing rather than the other. We see that premise 2 is false, and that the deceiver argument as a whole is no good.

7. Wittgenstein (1975) wrote (see footnote 3): "'The reasonable man does not have certain doubts" (Remark 220, p. 29e). [Vogel's note.]

8. Other philosophers have offered explanationist approaches to skepticism that differ from the one presented here. See, for example, Laurence BonJour, *The Structure of Empirical Knowledge* (Harvard University Press, 1985), and Christopher Peacocke, *The Realm of Reason* (Oxford University Press, 2006). [Vogel's note.]

In my view, inference to the best explanation provides a refutation of the deceiver argument along the lines just set out. But to make this response to skepticism work, a proponent has to show why and how the real world hypothesis offers better explanations than the skeptical hypotheses do. There are some significant obstacles in the way. For one thing, there is no philosophical consensus about what an explanation is or about the details of what makes one explanation superior to another. Another difficulty is that skeptical hypotheses can be formulated in importantly different ways. These issues can't be fully resolved here, but it's possible nonetheless to take some significant steps toward formulating an explanationist response to skepticism. I'll proceed by examining two diametrically different kinds of skeptical hypothesis. One is reticent and minimal, and we'll find that it is too impoverished to be acceptable. The other is more explicit and elaborate. It turns out that this fully developed version also fails to match the real world hypothesis in explanatory merit.

At bottom, the concern raised by skepticism is that our sensory experiences are caused *unveridically*; that is, in such a way that they don't correctly reflect the way the world really is. Consider some particular experience you have, such as your seeing this page now (hereafter, *P*). On the one hand, the real world hypothesis furnishes what seems like a perfectly adequate explanation of how *P* comes about (there is a page in front of you, your eyes are open, there is light shining on the page, and so you see it). Now, suppose that the skeptic offers, as an alternative, the following *minimal skeptical hypothesis*: Your experiences are caused unveridically (i.e., caused deceptively by something other than what you think). That's it. The minimal skeptical hypothesis is incompatible with the real world hypothesis. In particular, the minimal skeptical hypothesis and the real world hypothesis differ as to what the cause of your experience *P* is. But the minimal skeptical hypothesis says little or nothing about why you have *P* or how *P* comes about. In fact, it hardly explains the occurrence of *P* at all. Explanations that are defective in this way are called *ad hoc* explanations.

To get a feel for what's wrong with ad hoc explanations, you may recall the example of Dr. G diagnosing Roger's symptoms. Dr. G considers various hypotheses about the cause of Roger's symptoms (it's an allergy, a series of colds, or an ongoing respiratory infection). Roger's having an allergy explains his symptoms better than the other possibilities do, so Dr. G concludes that Roger does, indeed, have an allergy. Now imagine that someone else, Mr. S, comes along and says to Dr. G: "You don't really have a reason to think Roger suffers from an allergy. There's a competing explanation that's just as good, which you have no reason to reject. That competing explanation is: Something other than an allergy (don't ask me what) is causing Roger's symptoms." Taken verbatim, Mr. S's suggestion does little or nothing to explain why or how Roger's symptoms have come about. Dr. G would be foolish to set aside her diagnosis that Roger has an allergy, if all that Mr. S can offer as an alternative is the bare suggestion that Roger's symptoms are due to something other than an allergy. Similarly, you have no reason to set aside your belief that you are seeing the page of a book if the competing hypothesis is just the bare claim that your experience of the page is caused by something else.

The upshot is that our ordinary beliefs about the world provide a rich and comprehensive explanation of why we have the experiences we do, and the minimal skeptical hypothesis falls short by comparison. If a skeptical hypothesis is going to keep up with the real world hypothesis, it will have to go into more detail than the minimal skeptical hypothesis does as to how our experiences arise. But it's possible to overshoot in this direction, too. As a general matter, we want hypotheses that say enough to get the explanatory job done, and no more. Some hypotheses are defective because they say too little, but others are defective because they say too much.

Here is an example of this second failing. You believe that the earth is round. The fact that the earth is round explains why someone traveling in the same direction eventually gets back to the same spot, why the earth appears as a disk from the surface of the moon, and so on. However, there are people who belong to an organization called the "Flat Earth Society," and they believe that the earth isn't really round at all. To explain why travelers can get back to the same spot by apparently going in the same direction, Flat Earthers have to claim that compasses and other navigational devices systematically mislead us. We think we're going in the same direction, but we're really not. To explain why pictures from the moon show the earth as a disk, Flat Earthers say that there was a conspiracy to preserve the conviction that the earth is round, and the moon landings and pictures were all faked. The Flat Earth story is just too complicated to be believable. A much simpler and better explanation of everything is that the earth really is round. This point generalizes. Other things being equal, a simpler, more economical explanatory hypothesis is better than a less simple, unnecessarily complicated one.

Now, we've seen that some hypotheses of massive sensory deception, like the minimal skeptical hypothesis, don't serve the skeptic's purposes. The skeptic needs to advance a hypothesis of massive sensory deception that will match exactly the explanatory success of the real world hypothesis. Could there be such a skeptical hypothesis, and if so, what would it look like? Consider what I'll call the *isomorphic skeptical hypothesis*.[9] Imagine a computer that simulates the world item by item, feature by feature. For example, suppose you have the experience of seeing a cat eating. The computer has a file for a cat and an entry in the file that says the cat is eating. This file, rather than a real cat, causes you to have a visual experience like that of a real cat eating. Next, suppose you have the experience of the cat stopping eating and then grooming itself. The ordinary explanation of why you have this experience is, of course, that you are seeing a real cat that has stopped eating and is now grooming itself. But, according to the isomorphic skeptical hypothesis, what's going on instead is that the computer's cat file has been updated. The entry for eating has been deleted and replaced by an entry for grooming, which causes the corresponding experience in you. Overall, the isomorphic skeptical hypothesis denies that your experience is caused by everyday objects with familiar properties. Your experience is caused instead by computer files with electronic properties, such that the files exactly mimic everyday objects with

9. *Isomorphic* means "similar in structure."

familiar properties. The isomorphic skeptical hypothesis is supposed to match the real world hypothesis explanation for explanation. If it does, then the real world hypothesis is no better than the isomorphic skeptical hypothesis from an explanatory point of view, and explanatory considerations give us no basis for accepting the real world hypothesis rather than the skeptical alternative.

However, it is doubtful that the isomorphic skeptical hypothesis is truly the equal of the real world hypothesis in explanatory terms. Here is the basic thought, setting aside some important points: The real world hypothesis ascribes various familiar properties to ordinary objects. For example, if you have an experience of seeing a round peg, you ascribe the property of being round to an object, namely a peg. According to the real world hypothesis, the round peg behaves like something that is round (it looks round to you, it fits into a round hole, and so forth). Moreover, it appears that very little needs to be said to explain why a round thing like the peg behaves as it does. The peg behaves like a round thing *because it is round*. By contrast, on the side of the isomorphic skeptical hypothesis, something that *isn't* round is supposed to behave as though it were round. The isomorphic skeptical hypothesis can't explain such behavior the way the real world hypothesis does; the computer file doesn't behave like a round thing because it is round. Some *further, more complicated* explanation needs to be given as to why the computer file behaves systematically as though it were a round thing.[10] In this way, the explanatory apparatus of the isomorphic skeptical hypothesis turns out to be more complicated than that of the real world hypothesis after all.

The claim here is that the difference between the real world hypothesis and the isomorphic skeptical hypothesis is comparable to the difference between the "round earth hypothesis" and the "flat earth hypothesis." On the round earth hypothesis, the earth's really being round explains why it behaves as though it's round (e.g., why you're able to get back to the same spot by going in what is apparently one direction). On the flat earth hypothesis, the earth behaves in various respects as though it were round (e.g., you're able to get back to the same spot by going in what is apparently one direction), but the flat earth hypothesis doesn't explain this behavior by ascribing roundness to the earth. Instead, that hypothesis needs to be loaded up with some additional rigamarole to account for why the earth behaves as though round when it's not (e.g., why compasses and other navigational devices systematically mislead us, etc.). Ultimately, just as you are entitled to reject the flat earth hypothesis because of its explanatory deficiencies, so, too, you are entitled to reject the isomorphic skeptical hypothesis in light of its similar explanatory shortcomings.

Let's sum up. The deceiver argument is philosophically a stroke of brilliance, which seems to make skepticism about the external world inevitable. However, there are in principle two ways out of trouble. We may deny that the argument's conclusion follows

10. The same point applies to other properties besides roundness, of course. There is more to say about exactly why the isomorphic skeptical hypothesis turns out to be more complicated than the real world hypothesis. See Jonathan Vogel, "Cartesian Skepticism and Inference to the Best Explanation," *Journal of Philosophy* 87 (1990): 658–66, and Jonathan Vogel, "The Refutation of Skepticism," in *Contemporary Debates in Epistemology*, ed. Matthias Steup and Ernest Sosa (Blackwell, 2005). [Vogel's note.]

from premise 2 or we may deny the truth of premise 2. The first maneuver doesn't work. The conclusion does follow from premise 2, by way of the underdetermination principle. However, the second path lies open. The real world hypothesis provides better explanations than skeptical hypotheses do. The explanatory superiority of the real world hypothesis gives you reason to believe it and to reject its skeptical competitors. Thus, premise 2 of the deceiver argument is wrong, and the argument as a whole fails. Explanationism gives us the answer to the problem of skepticism about the external world.

TEST YOUR UNDERSTANDING

1. Does Vogel think that the problem with the deceiver argument is that the underdetermination principle needs to be assumed in order to derive the conclusion?
2. What is the example of the gas gauge supposed to show?
   a. That gauges can sometimes give incorrect readings.
   b. That seeming to see a tree does not make you justified in believing that there is a tree before you.
   c. That premise 2 is true.
3. Vogel discusses two failings an explanatory hypothesis may have. What are they?
4. Which of the two failings is supposed to afflict the isomorphic skeptical hypothesis?

NOTES AND QUESTIONS

1. There is arguably a disanalogy between Vogel's example of Dr. G and ordinary cases of knowledge by perception, say, knowing that there is a book on the table by vision. Suppose Dr. G knows that Roger has an allergy because this hypothesis best explains Roger's symptoms. Then this involves fairly sophisticated reasoning by Dr. G from her evidence that Roger is sneezing, that Roger has moist eyes, and so forth. If knowing by vision that there is a book on the table is similar, presumably it also involves sophisticated reasoning. Is this plausible? How do you think Vogel would respond to this worry?
2. For another example of this kind of response to the skeptic, see Laurence BonJour, "A Version of Internalist Foundationalism," in *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues*, ed. L. BonJour and E. Sosa (Blackwell, 2003), 3–96.

   A general discussion of "inference to the best explanation" is in Gilbert Harman, "The Inference to the Best Explanation," in Chapter 4 of this anthology. See also the "Notes and Questions" to William Paley, "The Argument from Design," in Chapter 1 of this anthology.