Chapter 5

OBSERVATIONAL KNOWLEDGE

According to classical empiricist doctrine, well-founded beliefs about matters of fact and existence can receive their support only from observation, memory, or inferences whose premises were ultimately supplied by observation or memory. On the face of it, observation is more basic than memory, because anyone who remembers that P must formerly have known that P, and when P is a contingent matter of fact, this prior knowledge could ultimately have arisen only from some kind of kind of observation. Memory and the sort of inference appropriate to matters of fact— Hume called it "experimental inference"—have always raised problems for empiricists, but observation is problematic in its own way. This chapter will deal with problems connected with the nature and scope of observational knowledge; the chapter to follow will be focused on memory, experimental inference, and the resolution of a skeptical problem that is raised in the present chapter.

A Problem about Observation

Although empiricists have always insisted that observation is our ultimate source of evidence about matters of fact, observation is actually far more problematic than it appears. In fact, it does not accord with some well-considered remarks by Hume, the most important figure in the history of empiricism. When Hume discussed philosophical skepticism in the last section of his *Enquiry*, he emphasized that the observational process typically results in beliefs or opinions, which may or may not amount to knowledge. His occasion for emphasizing this was his recollection of certain "trite topics" that skeptics "in all ages" dwell upon—specifically:

> the imperfection and fallaciousness of our organs on numberless occasions; the crooked appearance of an oar in water; the various aspects of objects according to their different distances; the double images which arise from pressing one eye... and many other appearances of a like nature.

Trite as these topics may be, they do prove, Hume admits, that "the senses alone are not implicitly to be depended on" and that we must "correct their evidence" by reason and by considerations derived from "the nature of the medium, the distance of the object, and the disposition of the organ...."[1] The evidence supplied by the senses in the observational process amounts to a thought or opinion (something propositional) because it can be corrected or corroborated by various considerations.

In these astute remarks Hume unwittingly raised a serious problem for empiricism, one that he made no effort to solve himself. The problem concerns the considerations that should be used in correcting or corroborating the thoughts generated by sensory experiences. What is the basis for these considerations? How can they be rationally supported? We should expect Hume to reply that they are supported "by experience," but it is not clear how, on his view, experience could accomplish such a thing. If the thoughts excited by a sensory experience must

[1] Hume (1777), p. 151.

invariably be corrected or corroborated by empirical considerations that owe their epistemic authority to experience, they too must ultimately have been excited by experience and appropriately corroborated. But what could have corroborated them? We seem to be faced with an infinite regress of empirical considerations or by some kind of corroborative circularity among them. The alternative of an intrinsically acceptable empirical consideration seems to be out of the question for a good empiricist like Hume.

Locke attempted to cope with the problem arising here by speaking of "the grounds of probability" bearing upon a belief or opinion. His conception of probability was primitive by modern standards, but it arose from his earnest attempt to improve upon the certainty requirement for rational opinion, which is not really appropriate, he believed, for empirical matters. Rational certainty is achieved, he thought, by demonstration or some kind of immediate insight, neither of which is applicable to uncertain matters of fact. Unlike rational certainty, which is an all or none affair, probability comes in degrees, Locke said, and it applies to beliefs or opinions created by arguments whose effect is naturally similar to that of demonstrations. In what way similar? The answer is "Similar in producing conviction." The conviction produced by empirical arguments is weaker than that produced by proofs, but it is equally concerned with truth and falsity.

The empirical arguments Locke was referring to concerned causes and effects, and the evidence appropriate to them was broadly observational. Hume would later call such arguments "experimental," but Locke did not have any general name for them. When we use them to assess human testimony about observed matters of fact, our evidence should include a variety of factors, Locke said.[2] Among them are the number of people who claim to have observed the phenomenon, their integrity as observers, and their skill in making such observations. If the testimony is taken from a book, we should consider the design or purpose of the author. As for testimony generally, we should ascertain "the constancy of the parts and the circumstances of the relation"--by which Locke probably means such things as the consistency of the testimony, its coherence, its "constancy" as to voice, competence, information, and so forth. In addition, we should consider the presence of circumstances that may have a bearing on all the preceding, such as fear, sorrow, love or hate, and the presence, nature, and circumstances of "contrary testimonies."

These grounds properly include the empirical considerations that Hume identified, because they are certainly pertinent to the truth of a claim like "I saw an oar bend when it entered the water." Locke's contention was that, to be rational, we should examine all the grounds of probability bearing on a proposition and then, upon a "due balancing of the whole, reject or receive it with a more or less firm assent proportional to the preponderancy of the greater grounds of probability on one side or the other."[3] In addition to wondering how a ground of probability can be also wonder about the principles to be followed in balancing probabilities. Neither Hume nor Locke identified these principles, and it is hard to see how they might have proceeded to do so. However that might be, we have the surprising result that, according to fairly explicit testimony that can be found in the writings of both Locke and Hume, observational evidence should be assessed by general considerations of an empirical kind. This is surprising, because classical empiricism is generally

---

[2] Hume discusses the way his thinks human testimony should be critically assessed in the second part of his chapter on miracles; his views are generally similar to Locke's but they are directed to a particular topic, miracles, and they are set forth less systematically. See his *Enquiry* of Human Understanding, Section X, Part II.

[3] Locke (1984), vol. 2, Bk. IV, ch. 15, sect. 5, pp. 366.

associated with the doctrine that general beliefs about the world must be founded on a generalization from experience. This testimony seems to turn things upside down.

Evaluating Observational Beliefs

If we think about the considerations Hume and Locke mentioned in relation to the assessment of human testimony or observational evidence, we can identify four basic sorts of things that should be weighed when we are evaluating a belief to the effect that some person perceives something. Suppose, for example, that our friend Tom forms the belief that he sees a gray cat in a nearby room. If we are seriously concerned to ascertain whether he does indeed see such a thing,[4] we should begin by reflecting on the perceptual process he employed--in this case, seeing. Tom may or may not be good at this process. His vision may or may not be acute; he may or may not be able to distinguish grays from tans; and it is even possible that he is blind and responding, in the present case, to a hypnotic suggestion. Another matter is the nature of what he ostensibly observes. Are there such things as gray cats? Are cats easy to see? Are they large, small, opaque, or transparent? Are there other things that look like them--things that a person like Tom might naturally (in these circumstances) take to be a cat? Then there is the character of the circumstances (the conditions) in which Tom's belief is formed. Was the light good enough for a person, a person like Tom, to see a cat, a gray one, at the distance in question? And finally there is the sort of observer (the sort of person) Tom is. Is he intelligent, sane, critical--or gullible and demented? Is he obsessed with cats? And so on and so forth.

Some of these questions might be answered by observing Tom's behavior and recalling how he has behaved on various occasions in the past, but these answers will raise further questions of the same kind about our own observations. Even if we could answer these further questions without circularity, we would be faced with general questions about observers, observable objects, observational processes, and conditions of observation whose answers are relevant to the evaluation of *anyone's* observational beliefs or reports, ours or the wisest and most perceptive among us. But to support these answers by reference to anyone's observations would be reasoning in a circle.

Empiricists rejecting Hume's "mitigated" skepticism often felt justified in avoiding these questions because they were convinced that empirical knowledge rested on a special kind of observation that rendered them unproblematic. This kind of observation might better be called "direct apprehension," since the objects it is concerned with are supposed to be subjective sensory objects. In chapter one I mentioned that Bertrand Russell considered our basic empirical knowledge to be obtained by "acquaintance," a process he understood to involve a direct confrontation with the objects of our immediate experience. He called these objects "sense data,"[5] but they could equally be described as *sensa* or what we sense when we have sensory experiences. Obvious examples of sense data are itches, feelings of pain, after-images, and the supposed sensory objects involved in the experiences we have when we look at purple clouds, fall foliage, or indeed when we have any perceptual experience. Since these supposed objects were thought to be directly and wholly presented to a subject, they were considered to be what the subject perceived them to be. No empirical assumptions relating to their nature, the nature of the process by which they are apprehended, the nature of the circumstances

[4] The question is whether he sees such a thing, not whether such a thing exists. The answer to the latter might be used as evidence for the answer to the former, but the questions are nevertheless different and should be answered in different ways.

[5] Russell (1953), p. 198. For a very perceptive up-to-date discussion of sense data, see Huemer (2007).

under which they happen to be apprehended, and the nature of the apprehending subject (the person) were considered pertinent to what the subject knows in apprehending them.

I have spoken of sense data as supposed objects because Russell's successors soon came to doubt that such things actually existed.[6] The doubters generally conceded that we have all sorts of sensory experiences, but they argued that sensory experiences were states of sensing that did not include the sensed objects that Russell called sense data. The arguments they offered for this surprising negative view—surprising, because it seems obvious that we do apprehend something subjective when we have after-images or double vision—were based on general considerations of a theoretical sort. Perhaps the decisive negative argument, the one most effective in convincing philosophers to reject sense datum theories, was to the effect that accepting such theories is tantamount to asserting that a sensuous curtain stands between perceiving minds and the world they normally believe they are perceiving.[7] This sensuous curtain shields the external world from our perceptual activity and renders it fundamentally unknowable. It becomes an incomprehensible Kantian "thing-in-itself."

Philosophers rejecting sense data on these grounds often thought they did not have to cope with the questions I asked about perception because they were convinced that sensory experiences bereft of sense-data could yield knowledge without presupposing empirical information about perceivers, perceived objects, the process of perceiving, and the effect of background conditions on what subjects might suppose they are perceiving. Roderick Chisholm stubbornly insisted on this, and he convinced many philosophers that he was right. In perceptual experience we are "appeared to" in certain ways, he said, and being appeared to is a "self-presenting state," one that necessitates the certainty, for the subject, that he or she is in that state.[8] Since being certain about the state one is in involves certainty about what that state is like, one's certainty about the character of one's current sensory experience depends only on that experience, not even partly on the empirical assumptions I have described.

Chisholm's position on this matter is far from convincing. Although some sensory experiences do, at least in some circumstances, seem to present themselves to our consciousness,[9] there is no good reason to suppose that we *cannot* make errors about them. Our access to them may be privileged, but it is hardly infallible. There is good evidence for this. Scientific studies have shown that people make all sorts of errors about the character of their mental states and sensory experiences; and they are often entirely unaware of what they are thinking, sensing, or feeling, particularly when their attention is focused on something else. The errors they make about their sensory experiences do not depend on a particular conception of them—whether they are understood as involving sensory objects (sense data) or not. People are simply not infallible in identifying, describing, or otherwise ascertaining the specific character of their sensory experiences or their conscious states generally.[10]

---

[6] See Barnes (1944-45), Quinton (1955), and, again, Huemer (2007).

[7] Ayer (1956), p. 117.

[8] See Chisholm (1976), p. 26.

[9] Normal human beings have sensory experiences whenever they perceive anything, but they are rarely aware of their sensory experiences, having their attention focused on things and persons in the world around them. For discussion, see e.g. Johnson (2006), BonJour (2007), Crane (2006).

[10] I discuss some pertinent studies in Aune (1967), pp. 31-38; other studies concerned with reports on mental processes generally are discussed in Nisbett and Wilson (1977), a study every philosopher should be familiar with. See also Williamson (2000), ch. 4.

Does Knowledge Need a Foundation?

The idea that we could make errors about the character of our sensory experiences—as opposed to the physical realities whose existence those experiences normally indicate—might seem perplexing or even alarming to philosophers who assume that if we have any empirical knowledge at all, it must rest on a foundation of something directly knowable. Classical empiricists commonly made this assumption. The basis for it is a regress argument that goes back to Aristotle.[11] According to this argument, if an empirical fact is known by means of some inference, the premises used in the inference must be known to be true. If those premises cannot be known to be true non-inferentially—by some kind of direct inspection—they will have to be known by reference to some more basic premises, which will have to be known to be true as well. Since one cannot possibly know something P on the basis of knowing some Q that is knowable, ultimately, on the basis of P itself, either the regress stops with *some* non-inferentially known fact or facts or P is not really known at all. Since nothing is a better candidate for being directly knowable for a subject S than the character of S's own experience, knowledge of such experience must be knowable in that way.

This argument is irresistible if we suppose that we have empirical knowledge and also believe that the certainty requirement is applicable to it. According to that requirement, anything we actually know is either immediately certain (certain without reference to anything else) or a provable consequence of other things that are immediately certain. But as I argued in chapter one, this requirement is not applicable to routine examples of empirical knowledge. According to the standards we normally use in everyday life, we know many things that are not immediately certain or provable consequences of other immediate certainties. I know that I live in the State of Massachusetts, but my knowledge of this, well-supported as it is, [12] is not an immediate certainty or something I have inferred from immediate certainties. Thus, when existing standards are assumed, the regress argument has no force and there is no plausible basis for inferring that our empirical knowledge, all of it, rests on some directly knowable foundation of certain truth.

If we allow that empirical knowledge nevertheless requires some starting place, one that may be neither indubitable by the subject nor the result of some actual inference, we still do not have to agree that it consists wholly or even partly of facts about the knower's sensory experience. What would suffice is some report or assertion (verbal or mental) that is reliably correlated with the sort of occurrence or state that makes it true. The man's assertion that the bird whose cry he hears on a remote lake in Maine is a loon has this kind of reliability, and so does the woman's assertion that the façade she sees through the window of a train (when no barn-façades are in the area) is that of a barn. In both cases the observers would normally be taken to know what they say they are hearing or seeing. Their evidence in these cases is, of course, defeasible and could therefore be overridden if countervailing evidence should become available—evidence about the prevalence of phony barn façades and evidence about fake loon calls broadcast on northern lakes by scientific maniacs with powerful amplifiers. Yet in the absence of such evidence

[11] Aristotle, *Posterior Analytics, 72b18.*

[12] I have, of course, an enormous amount of evidence for it—so much so that I can hardly survey it all or identify the most important elements belonging to it..

there would normally be no question that the observers have the relevant knowledge.[13] This evidence falls short of what is required by the certainty requirement for knowledge, but when that requirement is applied there is almost no room for empirical knowledge at all.[14]

It is important to realize that the existence of stopping places or non-inferred items of ostensible knowledge does not imply that persons having it need not possess knowledge of other, related things. If the man in Maine did not know what a loon is, did not know they cried in a striking way as they flew about the lakes they inhabit, he would not know that he was hearing a loon; and the woman would not know that the façade she sees is that of a barn if she did not know what a barn is and what it is typically used for. But this background knowledge need not provide premises from which the reports about the loon and the barn were inferred. To have imperfect knowledge—the kind not requiring rational certainty—a subject typically has to satisfy a cluster of epistemic conditions whose general character I discussed in chapter one, but these conditions do not require that some inference be made. Some of the required background knowledge may be more general, moreover, than the knowledge provided by observational reports. The structure of empirical knowledge is therefore quite different from the one suggested by Aristotle's regress argument.

The fact that ascriptions of knowledge are commonly assessed and even made on the basis of defeasible presumptions about the causes and effects of familiar phenomena suggests a solution to Hume's problem about how the evidence of our senses could possibly be corrected by reason and empirical facts about perceivers and perception. For any empiricist, reason itself—or pure reason, as Kant described it—has a very limited role in assessing observational claims. If these claims are to be assessed by background information, that information must have arisen empirically from observations of a less critical, or less cautious, kind. This is no doubt the way we got the information in the first place. We began innocently and naively, but we soon became more critical. Conflicting testimony required us to change our minds in many instances, and our naïve presumptions gave way to corrected and improved ones.[15]

Here is a simplified picture of how the correction and improvement is accomplished. There is an initial presumption, tacit rather than explicit, that able-bodied people who are reasonably young can equally discern what is present to their senses. This presumption is not baseless; it is supported by the agreement such people commonly reach about what visible, audible, fragrant, or foul. But this agreement is imperfect; the parties sometimes disagree about what is discernible under these or those conditions. Fortunately, regularities occur among the occasional disagreements, and we eventually conclude that, just as some people are stronger or can run faster than others, some people are better than others at seeing, hearing, smelling, or tasting. We also conclude or, better, learn that some perceptual conditions facilitate while others hinder the identification of colors, sounds, tastes, and smells. As we reflect on conclusions of this kind, we draw distinctions and adopt explanatory hypotheses that correct our original

---

[13] Huemer (2001) says that observers' "seemings" that they are perceiving a tree or house are "presumed true, until proven false' (p. 100). BonJour (2004) rejects the direct realism of Huemer and other philosophers, claiming that they "offer little or nothing by way of a positive account of how perceptual beliefs are [actually] justified according to their view." I agree that more can and should be said about this matter; I make some remarks about it later in the chapter and near the end of chapter six.

[14] See Fogelin (1994), p. 140.

[15] The idea that what might be considered initial opinions give way to improved ones in the process of rational inquiry is nicely modeled by Gupta's discussion of interdependent definitions in Gupta (2006), chapter 3.

presumptions. We now consider some observers more reliable that others; we now regard some perceptual conditions as optimal for certain sensible qualities; and so on.

The generating conditions for the higher-order knowledge needed for the assessment of ground-level observations is thus human inconsistency and judgmental conflict. Whether we are concerned with a single person's experiences or with intersubjectively available external objects, inconsistent reports are not only possible but actual. We resolve the inconsistencies by drawing distinctions and adopting explanatory principles, which we may eventually have to revise again. Sometimes our predications fail, and other explanations come to mind. We nevertheless become accustomed to resolving conflicts by rejecting some observational reports in favor of others. We disallow some because of the conditions in which they are made; we reject others on account of an ostensible defect in the perceiver's sense organ; and we disallow many because of the perceiver's distraction, lack of concentration, or carelessness and inattention. As our knowledge grows in respect to observers, ways of observing, observable objects, and conditions pertinent to observational success, we naturally take account of an ever-wider range of phenomena when we assess observational reports. Background theory thus becomes increasingly important. Observation loses its autonomy and becomes subject to higher-order principles.

The new principles we introduce in coping with observational conflicts do not concern only outer things; they also concern the nature of our experience when we make observations. The red, green, or gray things we perceive are located in space some distance from us; but sometimes we perceive something we want to describe with these color-words when nothing so describable available externally. These anomalies prompt us to think of our perceptual experiences as occurrences that somehow have qualities themselves. Wilfrid Sellars famously used a myth, the myth of Jones, to explain the origin of human discourse about sense-impressions,[16] but the concept of a sense-impression is not something every philosopher, let along every intelligent adult, will acknowledge having. So-called disjunctivists about the objects of experience say that we either sense external things or suffer hallucinations, but we do not sense something inner or subjective (objects or processes) in both cases: when we perceive something we do not sense something in addition to what we perceive.[17] The metaphysics of experience is now a contested subject,[18] and I lack the space to pursue it here. But I have no doubt that we do have perceptual experiences when we sense external objects and that these experiences possess qualitative features of their own—features that we can normally describe only by words strictly applicable to external things. One morning in the distant past I had the sort of experience spiritualists describe as seeing an apparition. I have always thought of it as a hallucination, the only hallucination I can remember ever having. My only way of describing it is this: "It was an experience of ostensibly seeing a drab motionless woman suspended in front of my bedroom window."

There is no primitive stage of our intellectual development when we did not think of ourselves as embodied creatures perceiving one another amid the objects of a common world. Our talk of our "selves" makes clearest sense only in connection with the thinking animals we actually are. But we normally look away from our selves when we perceive things, and not perceiving our arms, legs, or noses, we can

---

[16] Sellars (1959).

[17] See Crane (2006), section 3.4, and also Johnson (2006), pp. 286-89. A sense-impression as Sellars describes it may not be an object of inner sense, but disjunctivists would almost certainly reject such a thing anyway, as Johnson (2006) appears to do on p. 288.

[18] See the essays in Gendler and Hawthorne (2006).

be tempted to think of ourselves as something very different from an embodied creature. We might even come to believe, as Roderick Chisholm did, that we are spirits who move about the world and perceive it only "by means of" the body we inhabit and such things as the glasses that are perched on the body's nose.[19] But our ultimate evidence for our beliefs about ourselves and our world is the experience we have, not our inner states. That experience, which is fundamentally focused on the "objects" of our senses, needs to be understood theoretically, like anything else. Its character and place in the scheme of things needs to be investigated. As we pursue this investigation, our conception of our experience and our selves develops and changes. Current debates in the philosophy of mind show that this development is far from over. What we know about our sensory experience is therefore modified by higher-level inferences; it does not provide an immutable foundation on which the rest of our knowledge rests.

Alternatives to Foundationalism

The structure that I have just described does not accord with the usual alternative to the view, call it "foundationalism," that knowledge rests on a foundation of uninferred certainties or, as some say,[20] likely truths. The usual alternative is coherentism. According to this view, only beliefs can add credibility to beliefs,[21] and they can do so only in the context of a larger system of beliefs, one in which each member gains an indirect justification from the size and coherence of the system as a whole. The notion of coherence involved here is explained in different ways by different coherentists, but the alternative explanations are generally similar. BonJour, who was once a coherentist,[22] said that a coherent system must be consistent, both logically and probabilistically, and that its consistency is enhanced by inferential connections between its constituents and diminished both by unexplained anomalies among them and by relatively autonomous subsystems including them. [23]

The interplay between the justification a belief receives from another belief and the justification both beliefs indirectly receive from the system to which they belong obviously requires very careful treatment. How large must a system be to make a given belief strongly justified, all things considered? How large and how coherent must it be if a given belief amounts to knowledge? Fogelin (1994) once asked rhetorically if any human system of beliefs has ever satisfied BonJour's standards for coherence, and he answered in a way suggesting he thought the answer is clearly no. If this answer is right, as I think it is, presumably no belief has ever been justified and no one, by BonJour's coherentist standards, has ever known anything. This would have been an exceedingly unwelcome outcome for BonJour when he was a coherentist, because he developed those standards as a means of avoiding skepticism[24].

Although BonJour was convinced that foundationalism had unacceptable consequences, his coherentism involved the "internalist" conception of epistemic justification that was central to foundationalism. According to this conception, if

---

[19] Chisholm (1991), p. 171.
[20] BonJour 1999), p. 230.
[21] See Davidson (1986), p. 311.
[22] He abandoned coherentism in BonJour (1999), where he defended a form of foundationalism. An unrepentant coherentist, whose views are far more complicated than BonJour's, is Lehrer (1997). My reason for rejecting any form of coherentism appears at the end of this section. But also see BonJour (1999).
[23] BonJour (1985)
[24] BonJour (1985), p. 80.

belief A is not self-justified, it is justified by an inference from some belief B.[25] I have not explicitly criticized this conception of justification; I have simply not used it. In speaking of knowledge I have instead followed Lewis and spoken of evidence.[26] I could just as well have spoken of justification, but if I had done so, I would not have understood it as most internalists do.[27] The man who identified the loon's call made no inference; his belief that he was hearing a loon could be described as justified in the circumstances but not justified by itself: the available sounds are vitally important. In rejecting the typical[28] foundationalist's assumption that a belief is either self-justified or justified by means of another belief, I am therefore also rejecting the coherentist's assumption that a belief can be directly justified only by another belief.

In spite of this dissent from the internalist assumptions of the typical foundationalist and the coherentist, the picture of empirical knowledge that I favor might be described as a picture, really, of organized belief. A philosopher committed to what I called the certainty conception of knowledge might in fact insist on this description. That is all right by me. We can of course—within limits—use the word "knowledge" as we want, and if we want to restrict knowledge to instances of rational certainty, we can certainly do so. But as I argued in chapter one, existing usage does not impose this certainty restriction on the word "knowledge." There is therefore nothing incorrect in speaking of knowledge as I am doing here. I shall later, as I implied at the end of chapter one, discuss reasons for occasionally seeking greater certainty than what we ordinarily regard as satisfactory, and I shall treat these reasons sympathetically. But inherently uncertain "knowledge" is what we ordinarily have, seek, and argue about. The structure of that knowledge is what I have been discussing.

It is the fallibility of the best judgments we usually call knowledge that destroys any supposed invariant foundation for empirical knowledge. As we gain information about the nature of perceivers, perceived objects, perceptual processes, and background conditions that affect perception, our assessment of particular observation claims becomes so theory-dependent that we cannot realistically isolate an independent "observation language." In fact, in continuing to learn details about the microstructure of our world, we can easily reach the point of using paradigmatically theoretical language in making routine observations. This language may be highly exotic in university laboratories, but even wags in high-school lunchrooms have long been wont to speak of drinking H_2O or seasoning a dish with a little more NaCl.

These last observations might seem to support coherentism, but they do not actually do so. Although observational knowledge increasingly involves higher-level principles, the totality of what we know, or think we know, typically contains a lot of disorder.[29] This disorder is increased if we think of the knowledge of a single person rather than (as often in philosophy) "our" knowledge, the knowledge of some idealized community. Even in the best universities, physicists may be badly informed about the latest developments in psychology or molecular biology; and

---

[25] Not all philosophers who are internalists about justification would accept this. Feldman speaks of a feeling of wamth as a reason for believing something, but feelings are not premises from which conclusions can be inferred. See Feldman (2005), p. 273.

[26] See chapter one, pp. 13f. The notion of justification did tacitly come into my account of knowing for certain. See my definition on p. 34 and the paragraph immediately following it.

[27] See footnote 25 above.

[28] BonJour is now a foundationalist who holds that a basic belief can be justified by a perceptual experience, which is not and does not include a belief. See BonJour (1999), p. 230

[29] Gupta (2006) describes our view of the world as "a collage of conflicting pictures" that "contains both empirical anomalies and conceptual paradoxes" (p. 200).

mathematicians or philosophers may know next to nothing about diplomatic history or agronomy. A plain person, one who is not a professional knower, may have a real hodge-podge of knowledge; the aggregate will depend heavily on special interests, such as photography or the propagation of Hosta lilies, and on how much reading the person has done, and in what subjects. When we ask the impersonal question "What is known about the structure of space or the interface between chemistry and physics?" we may learn that what the best and brightest collectively know about these matters is very well organized, but there are gaps in even collective wisdom, and knowledge in some areas—for instance, the breeding of tigers in captivity—may have very little to do with any scientific discipline. Far from being an organized system growing from some single source, knowledge in the sense of what is known empirically has no general, specifiable structure. It is something of an aggregate of aggregates with a jungle of twisted and gappy connections.

These last observations apply to Quine's "holism" as well as to coherentism. In "Two Dogmas…" Quine famously said:

> The totality of our so-called knowledge or beliefs, from the most casual matters of geography and history to the profoundest laws of atomic physics or even of pure mathematics and logic, is a man-made fabric which impinges on experience only along the edges…. A conflict with experience at the periphery occasions readjustments in the interior…. Any statement can be held true come what may, if we make enough adjustments elsewhere in the system.[30]

Although Quine was disputing the supposed autonomy of analytically true statements in this passage, his words have been taken, rightly or wrongly, to support the idea that our beliefs form an interconnected web, elements of which are indirectly supported by the degree to which the whole structure comports with nonverbal experience. As I shall show in the next chapter, however, empirical confirmation is really not this holistic. Individual statements are always confirmed together with some others, but the totalities thus confirmed are not as extensive as Quine's words suggest. Just think of how you might confirm Tom's belief that it is freezing outdoors. Seeing that there is snow on the ground and that the plate of water left out for the birds is now a plate of ice would normally be sufficient; it would not require any consideration of evolutionary theory or the laws of supply and demand. Sometimes whole theories may be relevant to the confirmation of some empirical statement, as it often is in subjects such as astrophysics,[31] but this is far from usual. The whole of twisted and gappy connections I described above is actually compatible with everything we know about empirical confirmation.[32]

Knowledge and World: Some Problems

What I have been saying about knowledge here applies to imperfect knowledge, or knowledge ordinarily understood. As I have observed, this knowledge is based squarely on defeasible presumptions. Although these presumptions, which are known to be generally reliable, are commonly accepted and rarely questioned in everyday life, they provide well-known targets for philosophical criticism. The evidential basis for the criticism is sometimes very reasonable, and it deserves to be met. I will discuss some of it in the rest of this chapter.

---

[30] Quine (1953), p. 42.
[31] See Gribben (1998), p. 184.
[32] See chapter 6 below.

A basic theme in much of this criticism is the idea, which I have just been criticizing, that our beliefs about the world around us must ultimately be inferable, if they are rationally defensible, from the character of our sensory experience. Apart from the claims I mentioned earlier—about what is directly and primarily knowable, empirically—two additional reasons are often given for this idea. The first pertains to the transmission of information. Although it may appear that we are in direct contact with the physical things we perceive, there is actually a significant logical gap between the information we receive and the spatially separate objects transmitting it. To have actual knowledge of those objects, there must therefore be a flow of in formation from them to us: we must absorb that information and consciously take account of it.[33] Doing this, whether we consciously realize it or not, requires understanding and rational principles. We must in fact draw conclusions about what is external from information that is produced within us.[34] At the very least, we must infer causes from received effects.

The second reason frequently given is that the external objects we might think are presented to us in experience are actually very different from anything that is presented there. Descartes was the first philosopher to emphasize this point, insisting that external objects could be exhaustively described in relation to their geometrical and kinematic qualities.[35] We normally think of external objects as colored, noisy, or fragrant on account of the effects they ultimately have on our consciousness, but these effects, which Descartes regarded as ultimately determined by the geometry of external objects, light, and our sensory receptors, are ideational states productive of beliefs about those objects. Our current scientific beliefs about the external world are, of course, not the same as those of Descartes, but they agree with
his in not ascribing the sensuous qualities we discern in our perceptual experience to the objects themselves. Our sense organs and nervous system contribute to their character just as much as their more remote external causes do.[36]

Although both these beliefs have historically led to skeptical doubts about the very existence of a world external to our consciousness, it is obvious that one of them is based on the idea that we know the external world exists and also know a lot about its nature. A serious epistemological question that may yet be asked is "How is this knowledge possible?" or, less cryptically, "How can we possibly know what we think we know about this so called external world?" Hume in effect raised this question and came to the conclusion that it cannot be answered affirmatively: this supposed knowledge is not possible; we cannot really have it. His reasons for this conclusion are worth recalling, because they are still relevant to philosophical thinking on the subject of the external world. Some philosophers have recently defended principles that are tantamount to the ones Hume assumed.[37]

One of the reasons Hume gave involved a particular conception of legitimate non-deductive inference. What he called "experimental inference" is causal inference, or inference relying on a causal principle such as "Scratching dry, well-made matches on a rough surface in the presence of air causes them to light." Inference relying on such principles—for instance, an argument concluding that a match satisfying the conditions mentioned in the causal principle will, having been scratched, light—would now be considered deductive rather than experimental, but the cognitive process Hume described as giving rise to our belief in these principles

---

[33] See Dretske (1981), ch. 6.
[34] BonJour defends this nicely in his 1999 essay.
[35] I discuss Descartes' view of the external world in Aune (1991), chapter one, section 6.
[36] See the Introduction to the essays on the science of color in Byrne and Hilbert (1997).
[37] I show this in chapter six, when I discuss problems about inductive inference.

would be considered inductive today. As Hume explained it, this process involved a generalization from experience. On experiencing a "constant conjunction" between occurrences of a kind A and a kind B, we form the belief that B-occurrences are caused by A-occurrences, and the strength or firmness of our belief is determined, he said, by the number of these transitions that we experience. Hume did not actually describe the process of forming these beliefs as an inference because he could not specify an appropriate rule of inference. He thought we simply and naturally formed the appropriate belief when the constant conjunctions we experience become "sufficiently" numerous for us. The number of conjunctions required in the case of this or that person is purely an empirical matter.[38]

In spite of Hume's celebrated doubts about the rationality of this belief-forming process and the relative frequency of obtaining true beliefs by means of it, he considered it an acceptable process of belief-formation, the only such process available to us. But if we need to infer external causes for our subjective experiences by a process of this kind, we cannot possibly succeed. To employ the process we will have to experience a constant conjunction between our subjective experiences and their external causes, and to do this we will have to experience those causes directly—and this, Hume thought, is something we cannot do. If we could do it, we would have no need to infer their existence by any kind of reasoning.

The impotence of Hume's form of experimental inference to justify our supposed knowledge of the existence and nature of an external world does not necessarily raise a problem for contemporary empiricists, because many of them who think that the existence of external objects needs to be inferred by some kind of non-deductive reasoning accept forms of inference that do more than generalize from experience. One currently favored form, used by BonJour in justifying his belief in a world external to his consciousness, is Inference to the Best Explanation.[39] As it happens, there are serious problems with this form of inference; so one current way of providing such a justification remains questionable. (I shall discuss this matter further in the next chapter, where I discuss various forms of experimental inference.) But Hume had, as I said, other reasons for thinking that we cannot really know what we think we know about the external world, and one such reason has been given in a recent argument purporting to defend a conclusion directly contrary to Hume's. In Hume's thinking this reason supported a serious doubt about the meaningfulness of talk about a domain of objects that we cannot directly observe. The recent argument supports a similar doubt, but it is intended to undermine the kind of skepticism that Hume espoused.

It was a basic tenet of Hume's philosophy that meaningful words express genuine ideas and that genuine ideas arise from experience. Hume called the experiences from which genuine ideas arise "impressions" and claimed that any ostensible idea must, to be genuine, be derivable from one or more impressions. To be derivable from a single impression an idea must be a copy of that impression; to be derivable from a group of impressions, an idea must be complex and each idea ingredient in it must be a copy of some impression. The impressions copied by a person's genuine ideas must, of course, be impressions that person has actually had. Every genuine idea is the effect of one or more antecedent impressions. Since the supposed idea of a world external to one's consciousness could not have arisen from internal impressions, this supposed idea is bogus and cannot confer meaning on any word. If the term "external world" is meaningful, it cannot therefore have the meaning it seems to have; it cannot refer to anything that does not belong to a person's experience.

---

[38] I discuss Hume's epistemology in some detail in Aune (1991), ch. 3.
[39] Se BonJour (1999).

By current standards Hume's principle of meaning sounds very crude, but it was taken up and polished by nearly two hundred years of empiricist activity. The concept of experience underwent significant changes during this period, and the required connection between experience and meaning changed significantly as well. Some early twentieth-century empiricists held that empirically significant statements must be "reducible" to statements that can be verified by experience; logical positivists maintained that they must themselves be so verifiable. As I noted in the preface to this book, logical empiricists repudiated "reductionism"; they abandoned verificationism for confirmationism, the thesis that meaningful sentences must, at least in the context of some theory, be subject to empirical confirmation: they must support predictions that could in principle be verified and, if verified, would increase their probability. The last attempt by a logical empiricist to work out a satisfactory confirmation criterion of empirical meaningfulness was recorded in Rudolf Carnap's "Methodological Character of Theoretical Concepts."[40] In 1958 David Kaplan discovered a problem with the criterion Carnap offered in this paper, and Carnap subsequently abandoned the project.[41] Evidently he did not think an acceptable criterion for empirical meaningfulness could be found.

Although logical empiricists accepted Carnap's verdict on attempts to specify a criterion of empirical meaningfulness, the one-time empiricist Hilary Putnam has recently defended a new principle of meaning closely related to Hume's principle. Putnam's principle is a version of what is known as "Semantic Externalism," and it has a very positive bearing on Hume's problem about the external world. Putnam introduced his principle in an effort to show that Hume's problem cannot meaningfully arise. It is arguable, however, that Putnam's principle has a general, unsatisfactory consequence that was characteristic of Hume's principle: we cannot meaningfully say what we want to say, or think we are saying, about domains to which we lack experiential access. Putnam supported his principle by a now-famous thought experiment about brains in a vat, a thought experiment that adds vitality to skeptical doubts about a world external to our consciousness. It raises a problem even for philosophers who wish to maintain that we directly perceive an external world.

Semantic Externalism

Putnam introduced his semantic externalism by commenting on a famous paper by Alan Turing. In 1950 Turing proposed what he called an "imitation game" as a means of determining whether an appropriately programmed computing machine could reasonably be considered conscious. His idea was that if a scientific investigator, having examined a sufficient number of typewritten responses to questions designed to determine whether the respondent is an intelligent human being or a computer programmed to mimic the responses of a human being, could not distinguish the human respondent from the computer, then the investigator would be entitled to conclude that the computer is a thinking thing.[42] Putnam, after describing Turing's imitation game in some detail, asked whether a similar test could determine whether the words produced by a machine would actually refer to what a person using those words would be referring to. Putnam's answer was no. However natural and well composed the machine's responses to the investigator's questions

---

[40] Carnap (1956).

[41] To my knowledge, Kaplan never published his criticism, but he described Carnap's reaction to it in Kaplan (1971).

[42] Turing (1950). Actually, Turing used the question whether the investigator can distinguish the computer responses from the human responses as a replacement for the question "Can machines think?" He thought the latter question was "too meaningless to deserve discussion."

may be, if the machine has no sense organs or other hookup with the objects it is ostensibly writing about and no motor organs for interacting with those objects, it will not, he said, be *referring* to anything at all. If it is merely playing an imitation game, the words it produces will in fact be comparable to the sounds produced by a record player: they will not be genuinely referring uses of language.

Without attempting to specify the minimum conditions necessary for genuine reference, Putnam proceeded to apply his general conclusion about reference to a hypothesis that a skeptic might cite in support of a negative assessment of ordinary perception. The hypothesis concerns a number of brains immersed in a vat of nutrients and connected to a super computer in such a way that they have the sensory experiences of seeing, hearing, smelling, and physically interacting with a spatially extended external world of objects and persons. According to the hypothesis, the brains will believe they inhabit the spatio-temporal world we believe we are experiencing, but they will be wrong. The world of their experience will be a delusional world; their reality will be utterly different from what they think it is. The same could be true of us, the skeptic says. For all we know, we ourselves could be similar brains in vats. If we cannot eliminate this possibility, we cannot reasonably contend that we are what we think we are and that we know what we think we know.

Putnam used his thesis of semantic externalism to attack this skeptical scenario. Since the brains in the vat of nutrients are not and, he assumed, never have been in causal contact with a world of objects external to their consciousness, their thought-words cannot actually refer to such objects as vats, trees, and other persons. If their thoughts refer at all, they refer to the entities that stimulate those thoughts: their subjective experiences or elements of the computer programs that produce those experiences. This fact about reference undermines the skeptic's contention, because it implies that a sentence or thought "I am a BIV (a brain in a vat)" could not possibly be true. If a subject thinking this thought could, by means of it, think about actual brains in vats—that is, mentally refer to them—the thought would be false: the subject's thoughts and experiences would be connected to an external world containing vats and other things. If, on the other hand, the subject were not connected to real external objects, the sentence or thought "I am a BIV" would not refers to vats and could not thereby say something true about them. In one way or another, therefore, an utterance or thought "I am a BIV" could never truly affirm that the subject is a BIV. In consequence, it could not support a genuinely skeptical hypothesis.

Putnam expressed the conclusion of the argument in stronger terms than I have used here; he said we can *know* that we are not brains in a vat. This stronger conclusion does not appear to follow from the premises of his argument, however. Consider the assertion "I am a BIV." According to the argument, if I am a BIV my words "I am a BIV" do not have the reference they appear to have; they refer to subjective experiences or features of a computer program. If I am not a BIV but a rational animal, then my words do refer to a BIV and are false. Since I am either a BIV or a non-BIV, it follows that my words "I am a BIV" either have an exotic meaning or they assert something false of a real person. But neither disjunct of this consequence is shown to be true by Putnam's argument, and the disjunction as a whole does not entail that I have the categorical knowledge Putnam says I have.

If I had some direct knowledge of what my words (or thoughts) refer to, I could eliminate one of the disjuncts in question, but if Putnam's thesis of Semantical Externalism is true, direct knowledge of this kind is out of the question since it depends on causal factors external to my consciousness. I cannot therefore argue:

> My words "I am a BIV" refer to an embodied me and a real external vat. Therefore my words or my thought do not have the exotic meaning mentioned in the last paragraph. Therefore, my words "I am a BIV" assert something false about a real person. Therefore a real person and a real vat exist. I know this last fact. Therefore, I know that I am not a BIV.

My epistemic predicament appears to be described by the disjunctive conclusion that either my words "I am a BIV" have an exotic meaning or they assert something false of a real person. I do not know what my words actually mean and I therefore cannot eliminate either disjunct and so ascertain my true status as a thinking being.

In an encyclopedia article Anthony Brueckner suggested that the conclusion Putnam wanted can be obtained by a variant argument:[43]

a. If I am a BIV, then it is not the case that if my word "tree" refers, it refers to trees.
b. If my word "tree" refers, it refers to trees.
c. Therefore, I am not a BIV.

The first premise here is supposed to be a consequence of SE, Putnam's semantic externalism. The second premise is supposed to state a semantic fact that speakers can know a priori about their language, whatever it is and wherever they are speaking it. Thus, by virtue of knowing what "refers" means and knowing the meaning of quotation marks, speakers can supposedly know that disquotation is applicable to any successfully referring expression of their language. Since these two premises entail that the conclusion C is true, any speaker or thinker to whom "I" applies can supposedly know that he or she is not a BIV.

The argument is not satisfactory because a speaker (or thinker) to whom the disjunction I mentioned applies would not know what either premise is referring to. Suppose the referent of "I" is a certain BIV. The premise will then be true, but given SE the speaker could not understand what it is supposed to say—could not think the corresponding thought—because the speaker cannot comprehend a reference to trees. Similarly, if such a thinker entertained premise B, it would be thinking, "If my word 'tree' refers, it refers to trees*," the asterisk implying that the subject is thinking about what Putnam calls "trees-in-the-image," not trees in the intended sense. Non-BIVs could, of course, express the thoughts appropriate to the premises and conclusion, but if SE is true, they would not know what thoughts they would be expressing and so would not know that C is true.

The idea that we do not have a direct, privileged access to what our words or, more generally, our ideas refer to is contrary to standard empiricist doctrine, but Putnam accepted it, saying "meanings just aren't in the head."[44] If Putnam is right about this—if the meaning, the referential character, of a word or idea is in a significant way determined by input/output causal relations holding between that word or idea and objects in the world—then the empiricist idea that analytic truths do not (as Hume put it) depend on anything that is anywhere existent in the world must apparently be false. Putnam's Semantic Externalism is therefore an extremely important thesis. Not only does it, at least as Putnam believes, have serious consequences for what we can know empirically, but it appears to undermine the empiricist's conception of analytic truth.

---

[43] Brueckner (2004).
[44] Putnam, p. 19.

Criticism of Semantic Externalism

Is Putnam's Semantic Externalism a tenable doctrine? Is it well supported by the considerations Putnam offered in its defense? This last question is obviously weaker than the first, for considerations other than the ones Putnam used may support it more strongly than his did.[45] But it is easier to answer this weaker question, and answering it may make it unnecessary to answer the stronger one.

The first thing to say about Putnam's defense of his semantic externalism is that he provides no clear account of the connection that he thinks is necessary for genuine reference. In one passage he appears to say that a genuinely referring predicate must be associated with "language entry rules" or "language exit rules":

> There are "language entry rules'' which take us from experiences of apples to such utterances as "I see an apple," and "language exit rules" which take us from decisions expressed in linguistic form ("I am going to buy some apples") to actions other than speaking. Lacking either language entry rules or language exist rules, there is no reason to regard the conversation of the machine…as more than syntactic play (p. 11).

But this claim is far too strong. Not every meaningful predicate is what a philosopher of science would call an observation term. We can surely talk about electrons, photons, and a host of other things without possessing language entry rules or language exit rules that feature the relevant predicates.

It is conceivable that in speaking of language entry and language exit rules Putnam meant to assert something far weaker—namely, that reference is possible only in a language containing basic predicates that are associated with such rules. This weaker thesis recalls the old doctrine of complex ideas, which I mentioned earlier in connection with Hume. According to this doctrine, some ideas are simple while others are complex. Complex ideas are built up from simple ones; we construct some of them and others arise from our interactions with complex objects. The ideas of a mermaid and a centaur are human creations; the idea of a dog or giraffe was no doubt originally generated in some human beings by the experience of perceiving such an animal. Other human beings got the idea from parents and friends who explained what these animals are like.

Do the words "mermaid" and "centaur" refer to anything? Given the sense in which Putnam uses the word "refer" in *Reason, Truth, and History,* we would have to say no. Putnam stipulates that he uses "refer" to stand for a relation that holds between a word, symbol, or idea and something that actually exists (p. 1, note). These words are perfectly understandable, however; we know what features a thing would have to possess to be a mermaid or centaur. Since predicates are general terms that purport to refer to many different things, the sort of reference they have is what empiricists used to call "multiple denotation."[46] As far as the word "mermaid" is concerned, this kind of reference is clarified by the formula:

M1 ∀x("mermaid" refers to x iff x is a mermaid).

A more revealing statement about the reference of "mermaid" is the following:

M2 ∀x("mermaid" refers to x iff x is like a woman from head to waist and a fish from waist to tail.)

---

[45] I discussed Tyler Burge's version of the doctrine in the penultimate section of chapter four.
[46] See Martin (1958), ch. 4.

M2 specifies a reference condition for "mermaid," and if a predicate of some language or conceptual system is associated with such a condition, we can say that it has a referential use even though it may lack an actual referent.

A weaker, more plausible claim that an empiricist might want to make about meaningful predicates is that they have a referential use only when they are associated with a reference condition that is specifiable by means of predicates that are themselves directly or indirectly attached to existing objects. The attachment to existing objects that these predicates have could be explained further by mentioning language entry rules, which a subject conforms to in making observations. The weaker claim I have been describing is not precise, but it is no more indefinite than Putnam's remarks about a predicate's causal connection to its referents. It does, however, raise a problem that is pertinent to the limits of acceptable reference.

The problem concerns the specificity of the relevant reference condition. If the condition is supposed to furnish necessary and sufficient conditions for a successful reference, it amounts to a definition of referential meaning in observational terms. This is evident from the fact that a language-entry transition conforming to a language entry rule is a propositional response to an experiential stimulus, an example of which would be thinking "That's red" when one has an appropriate red-sensing experience. But a definition of referential meaning in observational terms is tantamount to a positivist's conception of referential meaning. If the general condition that must be satisfied for acceptable reference involves a looser connection with language entry rules—one loose enough to permit reference to unobservable entities—it may then be perfectly acceptable, but it will not support Putnam's semantic externalism, for referring terms will not themselves have to be attached to anything that can prompt a language-entry response. Only a weak condition impresses me as realistic. We can meaningfully refer to leptons and quarks, which are in no way observable; but if we can do this, BIVs should be able to refer to brains, vats, and distant objects.

If we recall the basic structure of Putnam's argument for his semantic externalism, we can see that the case he made was exceedingly weak. He began by describing Turing's imitation game, which was intended to provide a test for answering the question, "Could a computing machine that successfully performed a certain imitative task be reasonably considered conscious?" He then asked if an analogous imitative task could show that a machine actually referred to something. He answered no, saying:

> What we have is a device for producing sentences in response to sentences. But none of these sentences is at all connected to the real world (p. 10).

He did not pose the general question "How must a language or discourse be connected to the world if words occurring in it are to refer to things in the world?' but he did say that unlike sentences that the machine might produce,

> Our talk of apples and fields is intimately connected with our *nonverbal* transactions with apples and fields. There are "language entry rules" which take us from experiences of apples to such utterances as "I see an apple" and "language exit rules" which take us from decisions expressed in linguistic form ("I am going to buy some apples") to actions other than speaking. Lacking either language entry rules or language exit rules, there is no reason to regard the conversation of the machine…as more than linguistic play (p. 11).

Putnam's last sentence here (on a plausible reading) is pretty clearly true, but it does not imply that every word that refers to something is associated with language

entry or language exit rules. No doubt some rules of this kind are needed if the words of a language or discourse are actually *applied* to objects in the world, but Putnam does nothing to show that all referring words require such rules. Until he shows this, his case for semantic externalism is basically unsupported.

The idea that many referring terms are not associated with language entry rules is actually required for important claims Putnam makes about substances such as water. He makes these claims in slightly different ways in different essays. According to one statement, the referent of the word "water"[47] is identified by means of paradigmatic samples whose chemical composition is ascertained by experts in chemical analysis. (He says the natures of other natural kinds are ascertained or identified by experts with other specialties: botanists are the experts to whom one appeals for information about the nature of plants, for instance.) Chemists tell us that water is H_2O; a substance on another planet that is superficially similar to our water but is not H_2O would not be water. According to another statement, paradigm samples of water are identifiable as such because of their appearance and the functional role of similarly appearing stuff in our world. Water is a transparent liquid that quenches thirst and makes plants grow; it falls from the sky as rain, fills lakes and ponds, and so on. Experts assure us that the substance in our world having these features and playing this role is H_2O. Since water *is* this substance—since it *is* H_2O—nothing could be water that is not H_2O. It is a necessary truth, one known empirically, that water is this chemical substance.[48]

If we ask what H_2O is, we will not want to be told "It is water." We will want information about the chemical formula. We will want to hear some story about hydrogen, oxygen, and the way these elements are related in H_2O molecules. But when experts start talking about hydrogen and oxygen atoms, they will be talking about entities that we cannot observe or interact with individually. There will be no language entry rules and language exit rules associated with the terms "hydrogen "atom and "oxygen atom". Putnam was fully aware of this fact. So he could not

---

[47] The word "water" is here functioning as a name, not a predicate; it names what I would call a natural substance rather than, as Putnam says, a natural "kind." As I see it, kinds are abstractions, not concrete realities. It is worth observing here that Putnam's semantic externalism is better suited to names than to predicates.

[48] Although Putnam's semantic externalism is better suited to names than to predicates, his claims about water in this paragraph are obviously highly idealized and add little support to his externalist conclusions. No one supposes that a homogeneous substance actually fills all our lakes, ponds, and streams or that the liquids in those different geographical sites are chemically identical. Although we have very good reason to believe that any water we drink, swim in, or sail on consists largely of H_2O, our normal means of identifying a sample of water does not depend on this belief or on any other chemical lore. A chemist can tell us what proportion of a given liquid is H_2O or what other compounds it contains, but the decision to apply the label "water" to the liquid in the Cuyahoga river (which once caught fire), the Campus Pond at my university (which is often black and murky owing to the presence of thousands of migrating aquatic birds), the Dead Sea (which is heavily saline), or a diluted gallon of what was once Chardonnay wine, will not depend on such a person's decision. In fact, if our acid rain began to contain substantial amounts of the chemicals making up the XYZ liquid that fills the rivers and ponds of Putnam's Twin Earth without any significant effects on its ability to quench the thirst of animals or contribute to the growth of familiar plants, ordinary people would call it "water" without hesitation and continue to do so if, owing to some extraordinary natural change, it became pure XYZ. These and comparable other facts make it evident, I believe, that a meaningful reference to water does not depend, conceptually or semantically, on any set proportion of actual H_2O in the liquid a normal person is thinking of. A person with a smattering of chemistry might, of course, conceive of water as H_2O, but this conception would be anomalous in practice, for no water most persons have ever drunk is close to being pure H_2O. Good drinking water is heavily dependent on its mineral content. For similar ideas on the relation between water and H_2O, see LaPorte (2003), ch. 4.

himself endorse the idea that general terms can refer to objects only if they come within the scope of some language entry or language exit rule.

To make a plausible case for the view that BIVs cannot refer to brains, vats, or any other object of what we normally consider the world, one will have to resort to considerations that Putnam did not identify. Is there anything about the assumed experience of such beings that would preclude their thinking about themselves, their tank, and the world outside their tank? If so, what is it?

I frankly cannot identify such a thing. Putnam is the only post-positivist philosopher I am aware of who has officially denied that beings as intelligent and as susceptible to empirical stimulation as we are cannot refer to what we can refer to, and the reasons for his denial are clearly unsuccessful. Lest the reader suppose that the patent difference between bodiless BIVs and our mobile selves must render their references fundamentally different from ours, I should emphasize that the similarities between their kind and our kind are in some ways just as great as—and possibly even greater than—the differences. In fact, their intelligence and mental agility is supposed to be the same as ours, and their sensory input and conscious output—their sensory experiences and their awareness of what they are doing—are supposed to be "qualitatively" identical to ours. The differences between their experiences and thoughts and ours are limited to the way both are connected to external things and, consequently, to their supposed referential features. BIVs therefore have the ostensible experience of communicating with others and receiving responses from them; they have the experience of being members of a social community; and they ostensibly learn from others and provide instruction in return. But if we, from empirically identical experiential inputs and outputs, can develop a language that permits reference to trees and meadows, there is no apparent reason (other than the untenable ones Putnam provides) for thinking that the BIVs could not do exactly the same.[49]

Later in *Language, Truth, and History,* Putnam criticizes the very notion of objective reference, arguing that it belongs to the perspective of "metaphysical realism," which he rejects.[50] He calls his own perspective "internalism" and says that for those accepting this perspective reference makes sense only "within a conceptual scheme." We "cut up the world into objects when we introduce one or another scheme of description," he says, and because "the objects signs are alike *internal* to the scheme of description," it is "trivial to say what any word refers to *within* the language the word refers to":

> What does "rabbit" refer to? Why, to rabbits, of course! What does "extraterrestrial" refer to? To extraterrestrials (if there are any)…. For me [Putnam says] there is little to say about what reference is within a conceptual system other than these tautologies. The idea that a causal connection is necessary is refuted by the fact that "extraterrestrial" certainly refers to extraterrestrials whether we have ever causally interacted with any terrestrials or not![51]

In saying this, Putnam dismisses, virtually without argument, the skeptical problem that he attempted to dispel by a serious argument in the first part of his book.

---

[49] A detailed explanation, based on neural inputs and brain physiology rather than external objects, of how human beings can form a "mental" representation of themselves and their environment can be found in Trehub (1991).

[50] Putnam's metaphysical realism is a peculiar doctrine that no actual philosopher, to my knowledge, ever espoused. I criticize it in Aune (1985), pp. 126-28. William Lycan caricatures it in a wonderfully amusing way in Lycan (1988), p. 190.

[51] Putnam (1981), p. 52.

The minimal argument Putnam gives for this later position concerns the perspective of the person framing the BIV hypothesis. Certainly no BIV would advance this hypothesis, he says; and if a non-BIV were to advance it, "the world would not be one in which *all* sentient beings were Brains in a Vat." So, he concludes, the BIV hypothesis "presupposes from the outset a God's Eye view of truth, or, more accurately, a No Eye view of truth—truth as independent of observers altogether."[52] And this is incompatible with his internalist perspective. A metaphysical realist might attempt to pose the problem, but his or her assumptions about reference and truth would render the attempt futile, since a BIV could not, on those assumptions, entertain the hypothesis at all.

This way of disposing of the BIV hypothesis is far too simple. The question of how we can know that certain sorts of unobservables exist is theoretically significant, and the perspective of the being who advances the BIV hypothesis is not sufficient to refute it. Any sane person believes that he or she inhabits an objective world of animals and things, but Putnam's story of brains in a vat is coherent and describes a conceptual possibility, which philosophers normally find interesting to think about. If, as we believe, we can think about objects we cannot actually observe—if we can think about electrons and photons no less than prime numbers and algebraic functions—the same should be true of BIVs: they should also be able to think about things that are not, for them, observable. Contemplating such a possibility does not require some philosophically objectionable "perspective."

A Skeptical Problem Restated

When we contemplate the possibility of BIVs thinking about objects they cannot actually observe, an old epistemic problem arises again. It arises from the similarities I emphasized between our thoughts and experiences and those of the BIVs in Putnam's story. Although we certainly believe that we experience shoes and ships external to us in space, we know that there must be a flow of information from those objects to us, and the last part of this flow is qualitatively the same as what a deluded BIV is supposed to experience when it thinks it is perceiving a shoe or a ship. Another similarity is present in what Putnam calls language-exit transitions. When we implement an intention to reach for an apple, we have the experience of reaching for an apple—and a qualitatively identical experience would occur in a BIV according to Putnam's story. The sequence of events beginning with real external objects and our perception of them as well as the sequence of events beginning with our motivating intentions and ending with our overt actions contain segments consisting of conscious experiences; and these intermediate segments may be presumed to be empirically the same in us and in the BIVs. How, then, can we *know* that we are actually physically different from the BIVs?

As I observed in chapter one, we normally brush aside the possibility that we might be anything like BIVs. The idea is too far-fetched to be taken seriously in everyday life. But philosophical reflection, at least epistemological reflection, is not a staple of everyday life. In everyday life we say we know all sorts of things whose truth we ascertain or surmise only by means of presumptions that sometimes fail. One such presumption, the defeasible presumption I mentioned in chapter one when I discussed the example of the phony barn perceived through the window of a train, is that we actually see what we seem to be seeing in the light of day. Lewis mentioned other such presumptions in giving his Rule of Reliability and his Permissive Rules of Method.[53] Everyday ascriptions of knowledge are based on these

---

[52] *Ibid*, p. 50.

[53] See chapter one, p. 12.

presumptions, but they are defeasible and always questioned by philosophers in search of certainty. They are interested in proof, and presented with Putnam's story, they will want to know if there is anyway of proving that we are not BIVs. Putnam purported to provide such a proof in advancing his semantic externalism argument. His argument failed. Is an alternative available?

For an empiricist, a proof is out of the question here. It is even out of the question for a philosopher like Hume, who, as I mentioned in chapter one, spoke of a kind of proof in principle applicable to empirical propositions: it is provided by "such arguments from experience as leave no room for doubt or opposition." A typical empiricist—and this includes Hume—would contend that even in this weak sense a proof is not possible for the proposition that what I called mediating experiences are in fact connected to external objects. Why not? Because the connection is causal and purely contingent. Such connections cannot be ascertained a priori. They can be "known" only by a posteriori inference, and this kind of inference yields probabilities rather than certainties.

As it happens, there are unresolved problems about the logical structure and rational acceptability of the forms of inference by which such connections can presumably be ascertained. These forms of inference are, in fact, needed to provide rational support for a significant variety of familiar beliefs—for instance, those about the remote past, the experiences of other people, and unobservable objects such as electrons. Beliefs about these things have always been problematic for empiricists. I shall discuss these problems, together with the subject of memory, one of the empiricists' three sources of empirical knowledge, in the chapter to follow.