

4 Consciousness

The life of the mind

So far, we have talked mainly about beliefs concerning things outside ourselves: the green of grass, the smell of roses, the feel of glasses in our hands. But there is much that we believe about what is internal to us. I believe that I am *thinking* about self-knowledge, that I am *imaging* cool blue waters, and that I *believe* I am a conscientious citizen. In holding these beliefs, I attribute rather different sorts of properties to myself: thinking, imaging, and believing. What sorts of properties—or at least phenomena—are they, and how do our beliefs about them give us justification and knowledge? For instance, are some of these self-directed beliefs the products of a kind of *inner perception*? This seems a natural view, and we have already seen how an understanding of perception can clarify memory. If there is some truth in the inner perception view of self-knowledge, exploring the analogy between outer perception and self-consciousness might help to explain how beliefs about our inner lives are justified or constitute knowledge.

Our most important kind of self-knowledge is not about our bodies, but about our minds—for instance about what we believe, want, feel, and take ourselves to remember. It will help to start by describing the kinds of mental properties illustrated by thinking, imaging, and believing. Since they are all broadly mental, this is a task in the philosophy of mind. But epistemology cannot proceed without considerable reflection on mental phenomena, and here it overlaps the philosophy of mind. Thinking, inferring, and believing, for example, are central in both branches of philosophy; and to understand self-knowledge, we need a good sense of what kinds of properties characterize us. We might begin with two kinds that, for our purposes, yield a basic division.

Two basic kinds of mental properties

Thinking is a kind of process and involves a sequence of events, events naturally said to be in the mind. Thinking in human beings has a beginning, a middle, and an end; it is constituted by mental events, such as considering a proposition; and these events are always ordered in time, often in subject matter, and sometimes in logic.

Simply **having an image**, in the minimal way one does when there is a static, changeless picture in the mind's eye, is being in a certain (mental) *state*. Unlike something that changes, the existence of such a state does not absolutely require the occurrence of any events during the time it exists. *Imaging* can be a process of calling up a succession of images or, as when one of them is held changeless in the imagination, static. I could image something for a time without any change whatever in my imaging, and without the occurrence of any mental event that might be part of the imaging.

Believing could also be called a mental state, but this terminology can be misleading in suggesting that having a belief is a state of mind, where that implies a global mental condition like worry or excitement. Unlike images and aroused emotions such as jubilation, beliefs do not tend to crowd one another out.

Beliefs differ from images in at least two further ways. **First**, beliefs need not be in consciousness and indeed we can be conscious of only a quite limited number at once. We all have many which, unlike my belief that I am now writing, we cannot call to mind without some effort. **Second**, believing need not in any sense be “pictorial.” Consider a belief present in consciousness, in the way my belief that the rain has stopped is. This belief is present because I have called it to my attention. I might have held it without attending to it or even to the fact it records.

Even a belief present in consciousness in a prominent way *and* about something as readily picturable as the Statue of Liberty need not involve anything pictorial in the way my imaging must. Suppose I believe that the Statue of Liberty has a majestic beauty standing high in the Bay of New York. Without picturing anything, I can entertain this proposition, and in that way have this belief in my consciousness. By contrast, imaging cool blue waters requires picturing a blue surface. To be sure, when we *call up* this belief about the statue, we tend to picture that structure. But I could later get the proposition in mind, as when I am listing some majestically beautiful landmarks deserving preservation, without picturing anything. I could even retain the belief if I had forgotten what the statue looked like and simply remembered my aesthetic judgment of it.

It will help in sorting things out if we observe a distinction that has already come up but needs more development. Let us call **mental properties like beliefs** *dispositional* and **mental properties like thinking** (processes properties, we might say) *occurrent*. The latter are constituted by mental events and are occurrences: they take place in the way events do and may be said to happen or to go on. The former are not occurrences and may not be said to happen, take place, or go on.ⁱ

The basic contrast is this. To have a **dispositional property**, or (perhaps not quite equivalently) to be in a dispositional state, is to be disposed—roughly, to tend—to do or undergo something under certain conditions, but not necessarily to be actually doing or undergoing or experiencing something or changing in any way. Thus, my believing that I am a conscientious citizen

is, in part, my being disposed to say that I am one, under conditions that *elicit* that sort of verbal manifestation of my belief, such as you're asking me whether I intend to vote. Yet I can have this belief without doing or undergoing anything connected with it, just as sugar can be soluble while it is still in a solid, unaltered lump. I can have the belief in dreamless sleep. By contrast, to have an *occurrent property* is to be doing, undergoing, or experiencing something, as sugar undergoes dissolution. Thus, if you are *thinking* about mental phenomena you are doing something, even if you are in an armchair; and if you are *imaging* a flowering crab apple tree, you are experiencing something, at least in the sense that you're imaging the tree is now present in your *consciousness*.

Having a static image, however, as opposed to *calling up* an image, is not a process as, say, silently talking to oneself is. Occurrent mental properties, then, must be subdivided. To mark the difference, we might call occurrent mental properties like thinking *experiential process properties* and occurrent mental properties like having a static image in mind *experiential state properties*.ⁱⁱ Clearly, both differ from dispositional mental properties; possessing those does not even require being conscious, much less having a kind of experience. All three kinds of mental properties turn out to be important for understanding the *epistemological role of introspection*.ⁱⁱⁱ

Introspection and inward vision

If we take a cue from the etymology of 'introspection', which derives from the Latin *introspicere*, meaning 'to look within', we might construe introspection as attending to one's own consciousness and, when one's mind is not blank, thereby achieving a kind of inner seeing. I might introspect my images, for instance, and conclude that my image of the spruce and nearby maple indicates that the spruce is taller than the maple. I might *have* to introspect my image of the maple to tell without looking back at it whether it has three secondary trunks. Introspection need not, however, be labored or even constitute an act. It may be simply a matter of becoming conscious of something in my mind. This can be as natural as something's coming into one's physical field of vision, rather than like making the effort of observation in order to see.

It is not only in consciously introspecting that one can vividly image. In Shakespeare's *King Lear* there is a scene in which Edgar wants to convince Gloucester, who has lost his sight, that he is at the top of a cliff. Edgar's description is so vivid that the deception succeeds:

How fearful and dizzy 'tis to cast one's eye so low!

The crows and choughs that wing the midway air Show scarce so gross as beetles. Halfway down Hangs one that gathers samphire, dreadful trade!

Methinks he seems no bigger than his head.

The fishermen that walk upon the beach
 Appear like mice, and yond tall anchoring bark . . .
 Almost too small for sight.

(Act IV, scene vi)

What Gloucester sees in his **mind's eye** is so vivid that he believes he is at the edge of a precipice. His visual consciousness is filled with images from Edgar's portrait. Here introspection is simply a matter of vivid consciousness of the imagery that is before the mind.

If introspective consciousness does produce inner seeing and other sensuous imagery (such as, commonly, sound), we can try to understand it by drawing on what we know about **perception**. For instance, we can explore introspectional counterparts of some theories of perception and sensory experience. But one limitation of that procedure is apparent the moment we reflect on the dispositional mental properties, for instance believing, wanting, or having a fear of cancer. We do not see such properties in any sensory way, nor even as we may be thought to see (in our mind's eye) an image of cool blue waters.

The analogy to vision might, however, still hold for introspection regarding **occurrent mental properties**. If it does, it presumably applies only to the mental state properties, such as imaging. For surely thinking is not seen. It need not even be heard in the mind's ear. I may hear my silent recitation of Shelley's 'Ozymandias', but thinking *need* not occur in inner speech, certainly not speech of that narrative, punctuated sort.

Perhaps it is only pictorial mental properties that we see through inner vision; and perhaps it is only other "**phenomenal**" properties, such as inner recitations, tactual imagings (say, of the coldness of a glass), and the like, that seem accessible to **inner analogues of perception**: to hearing in the mind's ear, touching in the tactual imagination, and so on. It is doubtful, then, that we can go very far by conceiving introspection as simply producing inward seeing. Still, it is worth exploring how the analogy to seeing holds up for the one important case of pictorial properties.

Some theories of introspective consciousness

Suppose that introspecting such things as images of cool blue waters does produce a kind of inner seeing. Are we to understand this seeing on **realist lines**, so that there must be some real object seen by the introspective eye?

Realism about the objects of introspection

One might think that the **sense-datum** view simply cannot be extended in this way to **introspection**. This is at least a natural assumption about self-understanding. For on the introspectional counterpart of the sense-datum view, seeing (in one's mind's eye) an image of cool blue waters would require

something like another image, one that represents the first one in the way sense-data represent a physical object seen by virtue of the perceiver's acquaintance with them. Call it a second-order image, as it is an image of an image.

What would second-order images be like? If I try to image my image of cool blue waters, I get that very image all over again, or I image something else, or I get something that is not an image at all, such as a thought of my original image. But this point does not show that there could not be second-order images. Perhaps there could be some that are less vivid than the originals they picture, just as my imaginational image of blue waters is less vivid than the sensory image I have in seeing those waters. In any case, sensedatum theorists could avoid positing second-order images. It would be more plausible to hold that to image simply is to be acquainted with sense-data of a kind different and typically less vivid than outer perception yields.

An adverbial view of introspected objects

A defender of an adverbial account of sensory experience would not countenance images as sense-datum-like entities with properties in their own right. Take first perceptual imaging that is later "copied" in retrospective imagination. Adverbialists will likely hold that there is really just *one* basic kind of imaging process, and that it occurs more vividly in perception than in imagination. Thus, imaging blue waters is simply imaginationally, rather than perceptually, sensing in the way one does upon seeing blue waters: sensing "blue-waterly," as we might adverbially express it. Since the adverbial view conceives imaging as a way of experiencing rather than as a relation to an object, there is no image as an object to be copied.

On the adverbial view, then, there is no need to posit second-order images to represent first-order mental images to us, and the less vivid imagings which might seem to represent mental images are best construed as less vivid occurrences of the original imaging process. This point does not show that there *cannot* be second-order images or similar interior objects of inner vision, but the adverbial view reduces the inclination to think that there are any by suggesting a plausible alternative account of the facts needing explanation. Chief among these facts is that in recalling an image (especially a sensory image), one may have a less vivid image which apparently stands to the former as an imaginational image of a scene stands to the sensory image of that scene from which the imaginational image seems copied. The adverbial account of sensory (and other) experiences might explain this by interpreting the recalled image, say of blue waters, as recollectively sensing blue-waterly, where this is like visually sensing blue-waterly, but less vivid.

Given these and other points, it seems doubtful whether any realist theory of the introspection of images—one that takes them to be objects existing in their own right and having their own properties—can justify a strong analogy between that kind of introspection and ordinary viewing. For it is

by no means clear that there is any *object* introspected to serve as the counterpart of an object of ordinary vision. For an adverbial approach to experience, although realism about the (physical) objects of perception is a highly plausible view, realism about the objects of introspection is not. The idea is roughly that mental properties, such as imaging, can adequately represent physical objects in our mental life; inner objects should not be postulated for this task.

The anti-realist [why, anyway? And, how do we know, it, other than 'inner understanding' or introspection?] element in this adverbial view should not be exaggerated. To deny that mental images are objects having their own properties and in that sense are not real does not in the least imply that imaging is not real. Imaging processes are surely real properties of persons, even though they are not relations between persons and objects of immediate, inner perception. Does this imply, say, that introspection has no object in the sense of something it is of (or about), such as imaging blue waters? Certainly not. But this "object" is a kind of *content*, not an entity to which the mind is related. On the adverbial view of introspection, this kind of object is determined by what the introspection is about and is not a thing with its own properties, such as colors and shapes, sounds and movements, depths and textures.^{iv}

The analogy between introspection and ordinary perception

The adverbial view in question may seem unable to do justice to the apparently causal character of introspection. There is surely some causal explanation of my being acquainted with, say, imaging blue waters rather than imaging the Statue of Liberty when I monitor a daydream of a rural summer holiday. Perhaps it is mainly in what causes the relevant imaging that such introspective consciousness differs from seeing. How might this difference be explained?

Suppose the adverbial account is true. Introspection may still be like simple perception in two ways. First, introspective viewing may imply some kind of causal relation between what is introspected in it, say an imaging, and the introspective consciousness of that state or process. Second, such viewing may imply a causal relation between the object of introspective knowledge—for instance one's imaging blue waters—and the beliefs constituting this knowledge.

In explaining the analogy between introspection and perception, I want to concentrate mainly on introspective beliefs as compared with perceptual beliefs; we can then understand how introspection, and indeed consciousness in general, can ground justification and knowledge. A major question is how we can determine whether what the theory says is true: whether, in introspecting, as when we concentrate on our own imaging, the beliefs we thereby form about what we are concentrating on are produced by that very thing or by some aspect of it such as its imagined blue color. Only to the extent that they are should we expect introspection to ground justification and

knowledge in the broadly causal way that perception does. Many considerations are relevant here, but let me cite just two sorts.

First of all, it is surely *because* I am imaging cool blue waters that, when, with closed eyes, I introspectively consider what I am conscious of, I believe I am imaging them (and am conscious of my imaging them). It is natural (and reasonable) to take this ‘because’ to express a causal relation, even if I *could* be mistaken in thinking there is one. If the causal basis of my belief is not some inner object seen (as on the sense-datum theory), it is presumably the state or process of imaging. This is, in any event, how an adverbial theory of sensory experience would view the causal relations here. Similarly, if I introspectively believe that I am thinking about introspection, I believe this because I *am* thinking about it: it is the thinking itself that causes me to believe that it is occurring. In both cases, the introspective beliefs are produced by inner processes, and indeed in a way that makes it plausible to regard the beliefs as true. Some inner processes are *like* seeing an object in still other ways, but these processes can all be understood without positing inner objects analogous to perceptible ones such as trees and seen by the introspective eye.

A **second** point is this. Suppose my believing that I am imaging cool blue waters is not caused by my imaging them. The belief is then *not* introspective at all. It is *about* what is introspectable, but it is not *grounded* in introspection, any more than a belief merely about a perceptible, such as the rich red in a painting in a faraway museum, is a perceptual belief. Here, then, is another important similarity between introspection and perception.

Introspective beliefs, beliefs about introspectables, and fallibility

It may seem that the case described—believing one is imaging something when in fact one is not—is impossible. But suppose I have been asked to image cool blue waters, yet I hate the water and anyway have a lot on my mind. Still, if I want to be cooperative, then even though my mind is mainly on my problems, I may call up an image. However, as I am not concentrating on calling up the image, the image that I actually get might be only of a blue surface, not of blue waters. I might now inattentively assume (and thereby come to believe) that I have called up the requested image of cool blue waters. This belief is produced by a combination of my calling up the wrong image, which I do not attentively introspect at all, and by non-imaginational factors such as my desire to cooperate. I might even retain the belief for at least some moments after I cease to image at all. In that case, it is not only not true; it is not even introspective.

This example suggests that even a true belief about one’s conscious states or processes would not be introspective without being causally connected with them. It would be about these introspectable elements but not grounded in “seeing” them in the way required for being an introspective belief. Other examples support the same point. Imagine that my task is to think about

introspection for a solid hour. I monitor myself and, on the basis of introspection, conclude from time to time that I am thinking about introspection. As I reflect on my topic, I continue to believe that I am thinking about introspection. Now when I truly believe this simply because I have repeatedly confirmed it and am confident of steady concentration, and *not* because I am still monitoring myself introspectively, my belief, though perfectly true, is not introspective.

The best explanation of this point seems, again, to be that my belief is not caused (in the right way, at least) by the thinking that should be its ground. It is a **retained belief** [*grounded belief*] about my ongoing mental activity; it is not **grounded** in that activity as a focus of ~~my~~ introspective attention. My belief that I am thinking about introspection is a **propositional belief** that I am now doing so, but it is not an **objectual belief**, regarding my present thinking, to the effect that it is about introspection. It is not grounded in my present thinking at all, any more than my belief that a painting I remember portrays a rich red coat is grounded in seeing that red.

An overall conclusion we may draw here is that although there may be no objects such as sense-data or imaginational copies of them which we introspect, the process by which self-consciousness leads to introspective beliefs, and thereby to knowledge and justified beliefs about one's own mind, is nevertheless causal. Like perception of the outside world and (though in a different way) recalling events of the past, it produces something like a **sensory impression** and, at least commonly, beliefs about what seems to be revealed to one by that impression. The causes of introspective beliefs, however, are apparently processes and events in the mind. They are not, or at least need not be, objects that reside therein.

Consciousness and privileged access

Suppose that introspective consciousness is causally grounded in the way we have seen. We should then raise some of the same epistemological questions about it that we raised about perception. For instance, is **introspection** subject to counterparts of **illusion** and **hallucination**? And if it is, how might it still be a source of justification and knowledge? Let us start with the question of how anything like illusion or hallucination might occur in consciousness.

Infallibility, omniscience, and privileged access

One might think that, regarding the **inner domain**, which is the subject of **introspective beliefs**, one cannot **make mistakes**. If so, one might conclude that illusion and hallucination regarding this domain are impossible. Indeed, David Hume maintained that since "**the contents of the mind**" are known by "consciousness" (by which he meant something at least much like introspection), "they must appear in every respect what they are, and be what they appear."^v

Hume's statement suggests two far-reaching claims about self-knowledge. One claim—that the contents of our minds must be what they appear to us to be—expresses the idea that introspective consciousness can give us beliefs that cannot be mistaken. The other claim—that these contents must appear to be what they are—expresses the idea that consciousness makes us so richly aware of the (introspectable) contents of the mind that it guarantees us full knowledge of them. These ideas need refinement before we can reasonably appraise them.

The first claim suggests a thesis of infallibility (impossibility of error): an introspectual belief—roughly one to the effect that one is (now) in an occurrent mental state (such as imaging) or that one is undergoing a mental process (such as thinking) or that one is experiencing something (such as pain)—cannot be mistaken. The infallibility thesis rests largely on the idea that we are in such a strong position regarding occurrent mental phenomena that we cannot err in thinking they are going on inside us.

The second claim suggests a thesis of omniscience (all-knowingness) with respect to the present occurrent contents of consciousness: if one is in an occurrent mental state, undergoing a mental process, or experiencing something, one cannot fail to know that one is. The omniscience thesis rests largely on the idea that occurrent mental phenomena are so prominent in consciousness that one cannot [*unclear*] help knowing that they are present.

Together, these two theses constitute the strong doctrine of privileged access. The infallibility thesis says that our access to what is (mentally) occurring in us is so good that our beliefs about its present make-up are infallible; there is no risk of error. The omniscience thesis says that our access to it is so good that we cannot fail to know what (mentally) occurs in us; there is no risk of ignorance. It is because no one else is in such a good position to know about our mental life, and because we ourselves are not in such a good position to know about the *external* world, that it is natural to speak here of *privileged* access. The strong doctrine of privileged access is associated not only with Hume but even more with René Descartes, who is widely taken to maintain it in his famous *Meditations on First Philosophy* (1641), especially in Meditation Two.

Suppose for the sake of argument that the infallibility and omniscience theses are true. One might think that this would preclude inward counterparts of illusion and hallucination. It would not. Having illusions and hallucinations does *not* imply having false beliefs or ignorance of anything. Looking from a sharp angle from corner to corner, you can see a book as having the shape of a (non-rectangular) parallelogram without believing that it has that shape; and I can hallucinate a spruce tree like one that has burned to the ground without believing it is before me. We may know the facts. For inner perception as for ordinary sense perception, phenomenal experience is one thing and belief another.

Suppose, on the other hand, that there *are* no inner objects, such as blue, watery images, to appear to us to have properties they do not possess, such as

wavy surfaces. Then illusions of the kind we have in perception, in which an object appears to have properties it actually lacks, could not occur, as there is no object to appear to us. Nor could a hallucination of, say, an image of blue waters be *of* such an object and true or false *to* it, though of course in a sense it “misrepresents” reality. Suppose, however, that there are inner objects that we see when we image. How would hallucinating an image of, for instance, a loved one differ from just *having* that image? A **sense-datum theorist** might say that the hallucinatory image would be less vivid or stable than a real one. But it would still be an image of the same thing and could also be just like a normal image in vividness and other respects. It would be wrong to say, then, that a hallucinatory image is necessarily a less vivid or less stable version of a normal image, and the difficulty of explaining the difference between hallucinatory and real images is an additional reason to avoid (as the adverbial view does) positing mental images as objects.^{vi}

Difficulties for the thesis of privileged access

It might be, however, that quite apart from **illusion** or **hallucination**, we can have **false beliefs**, or suffer some degree of ignorance, about our mental life. This is clear for *some* mental phenomena, such as dispositions like believing, wanting, and fearing. We can mistakenly believe that we do not have a certain ignoble desire (say to make a fool of a pretentious boss), particularly if it is important to our self-image that we see ourselves as having no **hostile desires**. For the same reasons, we can fail to know that we *do* have the desire. One can also discover a **fear** which, previously, one quite honestly disavowed because it was at odds with one’s sense of oneself as courageous.^{vii}

Dispositions, however, should not be conceived as *occurring*, and in any case it is the occurrent mental phenomena to which philosophers have tended to think we have the kind of privileged access expressed in the theses of infallibility and omniscience. Can we be mistaken, or at least ignorant, about our occurrent mental states or processes?

Consider first the **possibility of mistake**. Could I believe I am thinking about the nature of introspection when I am only daydreaming about the images and feelings I might introspect? It seems so, provided I do not attend closely to what is occurring within me. Granted, this would not be a wholesale mistake; it would be like thinking I am watching someone observing a game, when I have become preoccupied with the game itself and have ceased to attend to its observer.

Suppose, however, that the infallibility thesis is restricted to beliefs based on *attentive* introspection, where this implies “looking” closely at the relevant aspect of one’s consciousness. Call this the **restricted infallibility view**; it says only that *attentive introspectional beliefs are true*. If I carefully consider the proposition that I am thinking about introspection, and I believe it on the basis of attentive introspection (i.e. on the basis of my carefully focusing on the relevant aspect of my consciousness), could this belief be mistaken?

It may seem that error here is impossible. But suppose I desperately want to believe that I am thinking about introspection. Could this not lead me to take my daydreaming to be such thinking and even cause me to form an attentive introspective belief that I am doing such thinking? It seems so. Similarly, I could believe, on the basis of attentive but imperfect introspection, that I am imaging an octagon and then, concentrating harder and counting sides, discover that the figure has only seven.

If, for some occurrent mental states (such as thinking), it is possible to be mistaken in believing that one is now in them, then the omniscience thesis of privileged access should be abandoned along with the infallibility view. This holds even if the omniscience thesis, too, is restricted, as it should be, to cases of carefully attending to one's consciousness. The easiest way to see why fallibility cuts against omniscience is to note how omniscience would tend to guarantee *infallibility* and so would be cast in doubt if the latter is. Let me explain.

Given the extensive self-knowledge implied by omniscience, if, instead of thinking about the nature of introspection, I am only daydreaming, then I must know that I am daydreaming. But I will presumably not be so foolish as *also* to believe that I am thinking about introspection—something plainly different from daydreaming. Since I would know as well that I am occupied with, say, a series of images that portray me as swimming in cool blue waters, it is even less likely that I will believe I am thinking about introspection. It appears, then, that if I know every truth about—am omniscient about—my consciousness, then I presumably cannot believe any falsehood about it and so am infallible about it as well.^{viii}

It is at best extremely unlikely (and perhaps impossible) that these two things—knowing every truth about one's consciousness and nonetheless believing some falsehood about it—occur together, leaving one omniscient regarding one's own consciousness, yet inconsistent and fallible about it. One would know every truth about it yet would also somehow believe falsehoods incompatible with those truths. This being at best improbable, if I am fallible I am at least very likely not omniscient. Now recall our daydreaming example. It casts doubt even on the restricted thesis of omniscience. In that example, although I am in fact daydreaming, I would presumably not know that I am. If I do know that I am daydreaming, I would believe this, and then it is very doubtful that I would also believe I am thinking about introspection.

These points suggest that, contrary to the thesis of omniscience, I can fail to know certain things about my consciousness even when I am attending to it; but they do not imply that the omniscience side of the privileged access view is wildly mistaken, in that I might be ignorant of *every* truth about my daydreaming. Far from it. As I (objectually) believe the daydreaming to be thinking about introspection, I presumably at least know my daydreaming to involve words or colors or shapes. I have some knowledge *of* it, but I would still not know the proposition *that* I am daydreaming and thus would not be omniscient regarding the mental processes occurring in me.

The possibility of scientific grounds for rejecting privileged access

Perhaps there could someday be a source of significant evidence against even the restricted doctrines of privileged access. For it could turn out that every occurrent mental phenomenon is uniquely correlated with some distinct set of brain processes. Then someone could devise a “**cerebroscope**” for viewing the brain and could read off the contents of consciousness from the cerebroscopic data (a possibility with disturbing implications that require ethical scrutiny). What would guarantee that our **introspective beliefs** must match what the machine says about our mental lives? And what would a mismatch show?

Imagine that we could discover cerebroscopically a unique neural pattern for, say, believing on the basis of attentive introspection that one is imaging cool blue waters, at the same time as we discover the pattern for imaging a field of blue-green grass. It would be natural here to suppose the **subject is mistaking** the grassy imaging for a watery one. Might we not regard the sophisticated equipment as more likely to be right than the subject?

There is a problem with this reasoning. How could one establish the unique correlations except by relying on the accuracy of people’s introspective beliefs? Might it not be necessary to start by asking people what they are, say, imaging, to assume that they are correct, and only *then* record the associated brain state? And if learning the correlations would depend on the accuracy of introspective reports, how could the correlations show such reports to be mistaken?

A possible reply is this. First, let us suppose that learning the correlations would depend on the accuracy of introspective reports. Still, neuroscientists would not have had to rely on the accuracy of precisely the introspective belief being shown to be mistaken, and perhaps not even on the accuracy of highly similar beliefs. In any event, once they construct their instrument, they might no longer need to consult introspection to use it. They might throw away the very ladder they have climbed up on.

Imagine, however, that they do have to rely on just the **sorts of belief** we are examining, together with **evidence** regarding these beliefs’ **reliability**—evidence we already have **independently** of the cerebroscope. Would this imply that the cerebroscope could not provide powerful evidence against introspective beliefs?

Consider an analogy. We might use a **mercury thermometer** to construct a **gas thermometer**. We might calibrate a container of gas with a **piston** that rises and falls as the gas is heated and cooled. The new temperature readings might correlate perfectly with mercury readings in many instances: in measuring water temperature, wood temperature, and in other cases. The gas thermometer might then be used for the same jobs as the mercury thermometer *and* might gauge temperatures that the mercury thermometer cannot measure, say because they are above its boiling point. Could we not use a gas

thermometer to correct a mercury thermometer in some cases, or perhaps to correct all mercury thermometers within limits? We could. This seems so even if we originally thought the mercury thermometer infallible in measuring temperature, perhaps because we mistakenly considered its readings partly definitive of what temperature *is*. **We would rebuild the ladder we have climbed up on.**

Similar points might hold for beliefs about what is now occurring in one. If the analogy does extend this far—if the gas thermometer is to the mercury thermometer rather as the cerebroscope is to sincere testimony about current mental states—then even the restricted omniscience view fares no better than the restricted infallibility view. For even when I am attentive to what is in my consciousness, a cerebroscope could indicate that I do not believe (hence do not know) that a certain thing is occurring, such as a frightening image which I believe I have put out of mind.

Introspective consciousness as a source of justification and knowledge

It is important not to overextend our criticism of various claims of privileged access. After all, even the restricted infallibility and omniscience views are very strong claims of privileged access. They can be given up along with the strong theses of privileged access quite consistently with holding that our access to what is occurring in us is very privileged indeed. Let us explore this.

The range of introspective knowledge and justification

Nothing I have said undermines a **qualified epistemic principle**. This **self-knowledge principle** says: our attentively formed introspective beliefs about what is now occurring in us are *normally* true and constitute knowledge. The difficulty of finding grounds for thinking they even *could* be false provides some reason to consider them at least very likely to be correct. Similarly, when we are attentive to what is occurring in us, then if something (knowable) is occurring in us, such as a certain melody in the mind's ear, *normally* we know that it is occurring, or at least we are in a position to know this simply by attentively forming a belief about what is going through our mind. At least this qualified epistemic principle holds for the domain of our conscious life.

Granted, our “access” to our dispositional properties is not as good as our access to what is occurring in us. We need not be conscious of the former properties, whereas the very existence of one's imaging (or of an image if there are such objects) *consists in* its place in consciousness. Beliefs and other mental dispositions need not even enter consciousness, or ever be a subject of our thoughts or concerns. Some of them may indeed be *repressed*, so that we normally cannot easily become aware of them.^{ix}

Nevertheless—and here is a **justification principle applicable to the**

dispositional mental domain—our beliefs to the effect that we are now in a dispositional mental state, for instance wanting, fearing, intending, or believing, are normally justified. We might also say that such beliefs, though defeasibly justified, are always prima facie justified, so that they are justified overall unless some defeating factor, such as an abnormal psychological interference, occurs. Moreover, normally, when we have a want (or fear, intention, belief, or similar disposition), we are in a position to know (and justifiedly believe) this. We can, then, usually know this if we need to. We very commonly do not know it, however; for such things may not enter consciousness at all, and there is often no reason to take any notice of them or form any beliefs about them. This kind of ignorance is innocuous.

There are many issues and details I have not mentioned; but we can now generalize about introspection (roughly, self-consciousness, i.e., consciousness turned inward) in relation to belief, justification, and knowledge, and summarize our main epistemological conclusions. Plainly, many beliefs arise from introspection, and the points that have emerged suggest a general epistemic principle concerning self-knowledge which, though far weaker than the infallibility thesis, is far-reaching: normally, introspective beliefs grounded in attentive self-consciousness are true and constitute knowledge. (This principle is slightly different from the self-knowledge one stated above.)

A second epistemic principle—an attentional epistemic principle concerning self-knowledge—though far weaker than the omniscience thesis, is that normally, if we attentively focus introspectively on something going on in us, we know that it is going on, under at least some description. I may not know that I am humming the slow movement of Beethoven's Pathétique Sonata, but I do know I am humming a melodic piano piece.

The corresponding justification principles suggested by our discussion seem at least equally plausible. First—to cite an introspective justification principle—normally, introspective beliefs grounded in attentive introspection are justified; and, second, normally, if I attentively focus on something going on in me, I am justified in believing that it is going on in me. To be sure, some such beliefs are better justified than others, and even some that are not attentive are justified. All of them are plausibly regarded as prima facie justified.

There are many possible principles regarding our justification and knowledge about ourselves, and there are many possible qualifications of the one just stated. But those principles are sufficient to suggest the power of introspection as a source of justification and knowledge. The examples I used to argue that introspection is fallible do not show that the apparently false introspective beliefs were unjustified or that true ones are not knowledge. A false belief, particularly if it is of a kind usually justified, can still be justified; and a true belief of a kind that can sometimes be false may itself constitute knowledge.^x

The defeasibility of introspective justification

These points about the high degree of privileged access we apparently do have may create **a danger of overestimating the strength of introspective justification**. From our examples, it might be thought that attentive introspection, even if not absolutely infallible, generates a kind of justification that at least cannot be defeated. Even if I am somehow mistaken about whether I am imaging blue waters, if I believe this on the basis of introspection, it would seem that I am *in the right*, **even if objectively I am not right**.

How could I fail to be justified in believing that I am imaging cool blue waters, if my belief is grounded in attentive introspection? If the question seems rhetorical, this may be because one thinks there simply is nothing else I should have done besides attending and hence no possible defeaters of my justification by appeal to another kind of ground for belief. Let us explore this.

Granting that I could not fail to be justified *unless* I could have good reason to believe I may be mistaken, still, perhaps I could fail to be justified if I had good reason for believing I am mistaken, such as evidence from repeated cerebroscopic results indicating that I have been mistaken in many quite similar cases. It is far from obvious that I could not have sufficient evidence of this sort. It seems wisest, then, to conclude that although introspective justification tends to be very strong, it remains prima facie rather than absolute and can be defeated by counter-evidence.

In any case, plainly beliefs grounded in attentive introspection, such as my belief that I am now imaging blue waters, are normally justified to a very high degree. Moreover—and here we have still another justification principle—*normally, my simply being engaged in attentive introspection also yields* **situational (propositional) justification for beliefs about what I am attending to**, even when it does not in fact yield any such beliefs. If I somehow “notice” my imaging blue waters yet do not form the belief that I am doing so, I am nonetheless (prima facie) justified *in* believing, and have justification for believing, that I am, just as, if I see a bird fly past and take no notice of it, I am still justified in believing it is flying past me. The analogy to perception seems sound here, and that is one reason why introspection is considered a kind of inner observation and (unless it somehow yields no content) a kind of inner perception.

Consciousness as a basic source

If we now ask whether consciousness, including especially introspective consciousness, is a *basic* source of belief, justification, and knowledge, the answer should be evident. It is. In this, as in many other respects, it is like perception. But it may well be that the **degree of justification** which consciousness (including introspection) generates is greater than the degree generated by perceptual experience, other things being equal.

The special strength of justification on the part of beliefs about elements in consciousness has led some philosophers to think that these beliefs are a kind of foundation for knowledge and for the justification of all other beliefs. Descartes is often thought to have so regarded introspectively grounded beliefs or knowledge. Whether knowledge and justification need a kind of foundation and whether, if they do, these beliefs are the best candidates to serve as a foundation—better than, say, perceptual and memory beliefs—are the major questions pursued in Chapter 9.

There is a further epistemologically significant difference between perception and consciousness, especially as manifested in introspection, as sources of knowledge (and justification). We can by and large introspect *at will*—roughly, just by (sufficiently) wanting to—though we may also do it quite spontaneously. Moreover, there is no limit to how many things we can come to know by introspecting, if only because we can, without limit, call up images and construct thoughts. But we cannot perceive at will; and what we can know through perception is limited by what there is outside us to perceive and by external conditions of observation, just as what we can know through remembering or recalling is limited by what has actually happened (or what propositions are true) and by the conditions of belief or image retention crucial for remembering or recalling.^{xi}

Introspective consciousness, then, is unlike perception and memory in enabling us to acquire a considerable amount of knowledge whether external circumstances cooperate or not. Whatever one can “observe” in one’s own mind is a possible subject of study, and it appears that many of the beliefs we attentively form concerning our mental lives tend to constitute genuine knowledge. Very roughly, introspective consciousness is a substantially active faculty; perception and memory are largely reactive faculties.

Granted, some content—like sensations of pain—comes into consciousness uninvited. Still, we can very freely call to mind both propositional and imagistic content. Some of it may come from memory, which shows that introspective consciousness may draw on that as well on resources created by imagination or intellect. By contrast, sensory content, such as perceptual images, enters our mind only when our senses are taken, by our own observational efforts or by contingencies of experience, to it. In the inner world, by sharp contrast with the external world, there is far more at our beck and call. This is perhaps another reason why introspectively grounded beliefs have sometimes seemed to be such good material to serve as foundations for knowledge and justification. In addition to the high degree of justification self-consciousness commonly confers on beliefs, it is an active source of both justification and knowledge.

There is a trade-off, however. Through perception, we acquire (primarily) justified beliefs and knowledge about the external world; without these, we would likely not survive. Through introspection, we acquire (primarily) justified beliefs and knowledge only about the internal world; with only this, our knowledge and justification would be sadly limited, even if we could survive.

This is not to underplay the importance of the internal world: without good access to it we would have little if any self-knowledge and, for that reason, probably at best shallow **knowledge of others**.

Self-knowledge is also important as a back-up when questions arise about one's justification or knowledge regarding external objects. Confronted with a strange object, one may carefully consider the stability, coherence, and variations of one's **perceptual experiences** of it in order to rule out **hallucination**. Told that one merely imagined a **car's passing**, one may try to recall the event and then scrutinize both the vividness of one's imagery and one's confidence that the belief comes from memory rather than merely imagination. Without the kind of self-knowledge possible here, we would have less knowledge about the external world. Both perceptual and introspective knowledge are vital, and both, as we shall soon see, can be extended, by good reasoning from the raw materials they supply, far beyond their beginnings in our experience.

Notes

ⁱ To say that beliefs are dispositional properties does not imply that the concept of believing is dispositionally *analyzable*, hence equivalent to a set of conditional propositions; and I do not think the concept is dispositionally analyzable.

ⁱⁱ To be sure, images can be possessed memorially, as is my image of the Statue of Liberty when I do not have it in mind; and 'imaging' can designate a process, as when I call up the series of images corresponding to looking at the statue from the Brooklyn Heights Promenade and glancing northward to Lower Manhattan, thence to the Brooklyn Bridge, and up the East River beyond the bridge.

ⁱⁱⁱ Both kinds of properties are experiential, in that they represent features of experience. Both, then, might be considered phenomenal, but sometimes the term 'phenomenal property' is restricted to the sensory kind that characterizes either the five senses or "inner sense," by which we feel sensations of pain and pleasure.

^{iv} Such contentual objects are often called *intentional objects*, largely on the ground that, like lofty deeds we intend to perform but do not do, they need not exist.

^v See David Hume, *A Treatise of Human Nature* (first published in 1739–40) (Part IV, section II), ed. by L.A. Selby-Bigge (Oxford: Oxford University Press, 1888).

^{vi} One might still distinguish genuine from hallucinatory images by insisting that to be a genuine image *of* (say) a loved one is to be an image caused by the corresponding sense, say seeing that person. But this view has an odd consequence. Through hearing a detailed description I could have an accurate image of Maj that is in a sense *of* her, as it perfectly "pictures" her, even if I have never seen her; but this would be mistakenly classified as a hallucinatory image, by the causal conception just stated. There are certainly different kinds of images and various ways in which they can mislead, but the analogy between perception and introspective consciousness does

not extend in any simple way to the possibility of inner illusions and hallucinations. I cannot pursue the matter further here, but for a detailed non-technical discussion of mental imagery see Alastair Hannay, *Mental Images: A Defence* (London: George Allen & Unwin, 1971) and my critical examination of this book in ‘The Ontological Status of Mental Images’, *Inquiry* 21 (1978), 348–61.

^{vii} Some of these cases seem to occur in *self-deception*, a phenomenon that raises profound questions for both epistemology and the philosophy of mind. For a comprehensive collection of papers on it (including one offering my own account), see Brian P. McLaughlin and Amélie O. Rorty (eds.), *Perspectives on Self-Deception* (Berkeley and Los Angeles, CA: University of California Press, 1988).

^{viii} The thesis of omniscience might be restricted to *introspectable* truths, as opposed to such truths as that there are 101 berries visible on the blackberry bush I am imaging, which I could know only on the basis of memory (and arithmetic) as well as introspection. The infallibility thesis might also be plausibly restricted in a similar way. This point bears on the connection between the two theses but should not affect the argumentation in the text.

^{ix} The thesis of omniscience might be restricted to *introspectable* truths, as opposed to such truths as that there are 101 berries visible on the blackberry bush I am imaging, which I could know only on the basis of memory (and arithmetic) as well as introspection. The infallibility thesis might also be plausibly restricted in a similar way. This point bears on the connection between the two theses but should not affect the argumentation in the text.

^x For reasons to be considered in Chapter 13, skeptics tend to deny this.

^{xi} There is less disanalogy in the negative cases: we cannot always cease at will to concentrate introspectively on our mental life, as illustrated by preoccupying pains; and we cannot in general cease perceiving at will; we must, for example, do so by closing our eyes or turning off a radio. This blocks the path of observation, just as an aspirin might block the path of pain.