
Introspective Awareness of Sensations*

Christopher S. Hill

ABSTRACT. My goal is to formulate a theory of introspection that can be integrated with a strongly reductionist account of sensations that I have defended elsewhere. In pursuit of this goal, I offer a skeletal explanation of the *metaphysical nature* of introspection and I attempt to resolve several of the main questions about the *epistemological status* of introspective beliefs.

The first three sections of this paper are concerned with several topics that fall in the area of overlap between empirical psychology and philosophy of mind. I distinguish between two forms of introspective awareness of sensations, one passive and the other active, and I give an account of their most salient properties. I then defend this account at some length. In the course of my defense, I criticize theories that attempt to explain introspection by representing it as an analogue of sense perception.

In the fourth and fifth sections I develop a skeletal theory of the epistemological status of introspective beliefs about sensations. It is often claimed in philosophy that all such beliefs are infallible, and also that the scope of introspection is so comprehensive that we can be said to be omniscient with respect to the realm of sensations. I argue that these claims are much too strong, but I also maintain that they contain an element of truth. I then turn to consider the question of whether there is logical room for scepticism concerning the deliverances of introspection. Is it possible to develop an argument concerning introspection that is similar to one of the familiar sceptical arguments concerning sense perception? I try to show that the answer is negative.

I

Frequently, when we say that a subject is introspectively aware of a sensation of type ϕ , we are making a claim which has the following truth conditions: (i) S has a sensation of type ϕ ; (ii) S believes that he or she has a

sensation of type ϕ ; and (iii) the belief cited in (ii) is both caused and confirmed by the sensation cited in (i). When a claim of this sort is true, S has what I call *Basic Awareness* of one of his or her sensations.

Basic Awareness is a form of consciousness. Thus, according to the definition, to be in a state of Basic Awareness is to stand in a certain epistemic relation to a sensation. It is to have knowledge of the existence of a sensation on the basis of one's current experience. But there is a sense of "conscious" in which it expresses this very epistemic relation. In general, when someone has knowledge of the existence of something in virtue of his or her current experience, it is appropriate to say that he or she is *conscious of* that thing. At the present moment, for example, it is appropriate to say that I am conscious of the pen in my hand.

There is a second sense of "conscious" that is germane to Basic Awareness. Used with this second sense, "conscious" does not count as a relational term at all but is rather an adjective that we use to pick out the members of a certain category of internal states. Thus, we frequently speak of *conscious experiences*. Since the objects of Basic Awareness are sensations, and sensations are conscious experiences, "conscious" can be said to express a monadic property that is possessed by the objects of Basic Awareness. By the same token, the noun "consciousness" can be said to stand for the totality of all of one's conscious experiences.

Some philosophers have described introspection as a "reflexive" phenomenon — as a state in which consciousness is engaged in self-contemplation.¹ In view of the present distinction between two uses of "conscious," it is easy to appreciate the temptation to view consciousness in this way. After all, the distinction implies that it is literally correct to say that in Basic Awareness one is conscious of consciousness, and that Basic Awareness involves consciousness of consciousness. These statements mean only that in Basic Awareness one stands in a certain epistemic relation — namely, the relation *being conscious of* — to components of the

totality of conscious experiences. It should be mentioned, however, that it is also somewhat misleading to make such statements; for they can suggest that Basic Awareness involves something which stands in a certain relation to *itself*, and a look at the foregoing definition shows that reflexivity is foreign to its nature. It is literally true to say that Basic Awareness involves consciousness of consciousness, but only because the two occurrences of "consciousness" have different meanings. It is literally false to say that Basic Awareness is reflexive.

II

Although it is not unusual for contemporary philosophers to write as if Basic Awareness were the *only* form of introspective awareness of sensations, it is necessary to recognize a second form. There is no clause in the definition of Basic Awareness which indicates that it calls for an action of any kind: Basic Awareness is an essentially *passive* phenomenon. But our ordinary discourse testifies to the existence of a form of introspective awareness that is essentially active in nature. Thus, we are disposed to describe ourselves and others as "attending to an itch," "concentrating on a toothache," "focussing on a burning sensation," and "scrutinizing a visual image." Attending counts intuitively as an action, and so do concentrating, focussing, and scrutinizing.²

This second form of introspective awareness, which I call *Active Introspection*, comes in three main varieties.

First, one can engage in Active Introspection by deciding to attend more closely to a sensation of which one is already aware. For example, after being marginally aware for a day or two of a spot in one's visual field, one might become interested in it and decide to subject it to scrutiny. As one examines it, one will become aware of aspects and components of the spot that had not previously been in evidence.

Second, instead of deciding to attend more closely to a sensation, one can decide to maintain the level of attention that one is currently bestowing upon it. This variety of Active Introspection is closely related to the first variety: they differ only in that the first variety never requires more than a single act of will, while the second one sometimes requires a continuing effort. (A continuing effort is required when other phenomena are competing with a given sensation for one's attention.) Because of their similarity, I will hereafter refrain from

singling the second variety out for special mention. What I say about the first variety can be taken as applying, *mutatis mutandis*, to the second variety as well.

Third, there is a variety of Active Introspection that occurs when one has a description of a type of sensation in mind, and one becomes interested in determining whether it is currently possible to bring a sensation which answers to that description within the scope of one's attention. In short, this third variety occurs when one endeavors to direct one's attention on a *new* sensation, where a sensation is new if one is not currently attending to it. Thus, for example, having lost touch with the aftertaste of one's most recent cup of coffee as a result of being temporarily occupied with other matters, one might suddenly recall the aftertaste and undertake to attend to it anew. Or one might have a description of a certain region in phenomenal space in mind, and one might become interested in determining what sensations are currently available there. Moved by this interest, one might focus one's attention on the region in question. As Pillsbury has pointed out, efforts of this sort usually meet with results: "If you will attend fixedly for a few moments to any point on the external skin, you will find coming into consciousness a number of itching, or prickling sensations which you had not previously noticed, and would in all probability not have observed were it not for the increased attention to that part of the body."³

There are two views about the nature of Active Introspection that seem to be worthy of consideration. One view analyzes it in terms of an analogy with vision, and the other explains it by comparing it to the action of adjusting the volume of a radio. I call these two views the *Inner Eye Hypothesis* and the *Volume Control Hypothesis*.⁴

Both hypotheses acknowledge and endorse the foregoing claims about the varieties of Active Introspection. Thus, both assert that Active Introspection can take the form of attending to a sensation of which one is already aware, and that it can also take the form of undertaking to attend to a new sensation. After this point, however, the hypotheses diverge. They offer analyses of attention that are quite different.

The Inner Eye Hypothesis maintains that there is a level of representation which mediates between sensations and beliefs about sensations. The representations at this level are said to be states of an internal scanning device, and this scanning device is said to stand in much

the same relationship to sensations as the physical eye does to extramental objects and events. One attends to an extramental entity by arranging for one's physical eye to be in the right position to pick up information about the entity. It is much the same, according to the Inner Eye Hypothesis, in the case of introspection: one attends to a sensation by adjusting one's internal scanning device in such a way that it becomes attuned to information about the sensation. Further, extramental entities can exist without standing in any informational relations to the physical eye, and their internal qualities are never affected by their coming to stand in such relations. The Inner Eye Hypothesis claims that the same things are true, *mutatis mutandis*, of sensations and one's internal scanning device. It asserts that sensations can exist without being scanned, and also that the internal qualities of sensations do not change when one scans them.

According to the Inner Eye Hypothesis, one attends more closely to a sensation of which one is already aware by arranging for one's scanning device to produce representations of the qualities of the sensation that are more vivid and more detailed, and one undertakes to attend to a new sensation by "turning" or changing the "orientation" of one's device in such a way as to enable it to pick up information about sensations that meet a certain description. In short, the Hypothesis analyzes the changes associated with attention to sensations in terms of changes in the relations between sensations and one's scanning device and in terms of changes in the internal state of the device.

In order to formulate the Volume Control Hypothesis succinctly, it is helpful to have the concept of a *phenomenal field* in hand. Let us say that the phenomenal field of individual *i* at time *t* is the totality consisting of all of the conscious experiences that *i* has at *t*.

The Volume Control Hypothesis denies that there is a level of representation which mediates between sensations and beliefs about them, and as part of this denial, it rejects the claim that attending to sensations involves an internal counterpart of the physical eye. It claims that to attend to an already familiar sensation is simply to increase the prominence of the sensation within one's phenomenal field, and that to attend to a new sensation within one's phenomenal field is simply to modify one's phenomenal field by summoning a sensation into existence. But what is it to bring a sensation into greater prominence within a phenomenal

field? And what is it to summon a sensation into existence? The Volume Control Hypothesis answers the first question by asserting that an increase in the prominence of a sensation is a change in certain of its intrinsic qualities, not a change in relations it bears to something else. It also likens the action of bringing a sensation into greater prominence to the action of changing the sounds emanating from a radio by adjusting the volume control. As for the second question, it attempts to explain what it is to summon a sensation into existence by asserting that each sensation derives ultimately from a packet of information in an unconscious portion of one's mind — a packet which has the potential to become a sensation with a particular set of phenomenal characteristics, but which must be subjected to further processing before it can achieve its potential. The Hypothesis claims that summoning a sensation of a certain kind into existence is simply a matter of actualizing the potential of a packet of the right sort. Thus, according to the Hypothesis, attending to a new sensation is like converting radio waves into sound by turning on a radio.

Here is an alternative version of what the Volume Control Hypothesis tells us. An episode of Active Introspection begins, it says, in one of two ways: either with a decision to bring an already familiar sensation into greater prominence, or with a decision to undertake to summon a sensation answering to a certain description into one's phenomenal field. In the first case, the decision is followed by a phase during which one adjusts one or more controls. These controls cause the packet of unconscious information corresponding to one of one's current sensations to be subjected to further processing, and this additional processing brings the sensation into greater prominence. (The processes that lie between the decision and the increase in prominence are unconscious. In this respect, they are like the searching that lies between one's conscious decision to look for a certain item in memory and the subsequent arrival of a memory image in consciousness.) In the second case, one attempts to determine whether there is an unconscious packet of information which is of the right kind to produce a sensation answering to the description one has in mind. (Here of course the nature of one's search depends on the nature of one's description. A search for a packet corresponding to "an itch" would be more complex and would take more time than a search for a packet corresponding to "an itch in my left shoulder.") In the event that a packet of the right

kind is available, one adjusts the controls that activate the appropriate information-processing mechanisms, and the packet is converted into a sensation. As in the first case, one eventually winds up with some new beliefs about one's phenomenal field.

I am now in a position to state the two views about the nature of introspection that I am particularly concerned to recommend. They can be summarized as follows: first, there are exactly two main forms of introspective awareness of sensations, Basic Awareness and Active Introspection; and second, the picture of Active Introspection that is provided by the Volume Control Hypothesis is fundamentally sound. My justification for thinking that there are at least two forms of introspective awareness of sensations has already been given (consider the definition of Basic Awareness in relation to my examples of Active Introspection), and my justification for thinking that there are at most two forms is that there are no data which make it necessary to postulate a third form. As for the Volume Control Hypothesis, my justification for accepting it is given in the next section.

In its present guise, the Volume Control Hypothesis is less a theory of Active Introspection than a rather sketchy metaphor. To proceed beyond the metaphorical level, it would be necessary to address several questions that are extremely difficult. First, what is the nature of the phenomenological changes that are associated with attending to an already familiar sensation? That is to say, what are the changes that are constitutive of the activity I have called bringing a sensation into greater prominence? Second, what is the nature of the unconscious processes that underlie these phenomenological changes and explain them? Presumably, when one undertakes to bring a sensation into greater prominence, one does so by committing a larger share of one's processing capacity to the task of making phenomenological sense of the unconscious packet of information that is most intimately associated with the sensation. But how exactly does one's conscious decision to attend to the sensation bring about this effect, and what is the nature of the information processing devices that one calls upon? Third, what is the nature of the unconscious processes that underlie and explain the changes in one's phenomenal field that occur when one undertakes to summon up a new sensation which answers to a certain description? What are the unconscious processes which lead from one's decision to try to summon up a new sensation of a certain kind to the eruption of a sensation of that kind into one's phenomenal field?

All of these questions are forbiddingly complex. Take, for example, the first one. I am inclined to say that an increase in prominence involves changes along two distinguishable dimensions — intensity and degree of internal complexity. But what is intensity? It seems to come to different things in different cases. Attention can increase the phenomenal volume of an auditory sensation, the vividness of a visual sensation, the severity of a pain, the importunity of an itch, and the strength of a feeling of pressure. I am strongly inclined to view these changes as similar, and it seems natural to use "intensity" as a label for the respect of comparison in terms of which their similarity is to be understood. However, in view of the obvious differences between such dimensions as phenomenal volume and severity of pain, the claim that they are similar is badly in need of clarification and defense. Further, what is internal complexity? Clearly it is a function of the number and character of the constituents of a sensation. But to develop this answer it would be necessary to know, for example, what it is about two auditory sensations that makes one a part or a constituent of the other. We would also need to understand the part-whole relation in the case of visual sensations, gustatory sensations, and so on. Unfortunately, it is far from clear how exactly to make sense of the notion of a part-whole relation in phenomenological terms, and it turns out to be no easier to find the common denominator of the family of phenomenological part-whole relations than it is to find the common denominator of the family of types of phenomenological intensity.

I am not at present concerned to give a fully developed theory of Active Introspection, but only to provide a sketch of its most salient features. Thus, the foregoing questions fall outside the scope of this discussion. It should be mentioned, however, that it is possible to find a certain amount of relevant information elsewhere in the literature.^{5,6}

III

Although I think that the Inner Eye Hypothesis is wrong, I must concede that there are several considerations which seem *prima facie* to support it.⁷ We should turn now to the task of assessing these considerations. Until we have reached a final decision about their merits, it will remain possible that the Inner Eye Hypothesis provides a more or less accurate picture of Active Introspection.

First, it can seem that the way we conceptualize the actions we perform when we engage in Active Introspection is more in accord with the Inner Eye Hypothesis than with the Volume Control Hypothesis. When I bring an already familiar sensation into greater prominence I never consciously think of myself as changing its intrinsic qualities, and when I undertake to summon up a new sensation I never consciously think of myself as undertaking to bring a sensation into existence. Rather I think of myself as “taking a closer look” at a sensation, or as “looking for” a sensation of a particular kind.

Second, it seems possible to distinguish between internal phenomena that count as states of awareness and internal phenomena that count as objects of awareness. In particular, it seems possible to distinguish between the states of awareness that are directed on sensations and the sensations on which such states are directed. Thus, for example, it seems that the former have a different location in phenomenal space than the latter. (In my own case, states of awareness seem to be located somewhere in the center of my head, while the sensations that count as objects of awareness are distributed throughout the phenomenal firmament.) This sense of difference and separation fosters the idea that sensations can exist without being objects of awareness, and it thereby contributes to the plausibility of the Inner Eye Hypothesis. (Recall that the Hypothesis explicitly claims that it is possible for a sensation to exist without being scanned by an internal scanner.)

Third, the Hypothesis gets a certain amount of support from cases in which we find ourselves acting as if we were under the influence of sensations, but in which we apparently have no awareness of sensations of the appropriate kinds. Thus, for example, one may find oneself scratching a leg and come to realize that one is doing so for a reason — the leg is itching. Again, as Armstrong points out, “[a]fter driving for long periods of time, particularly at night, it is possible to ‘come to’ and realize that for some time past one has been driving without being aware of what one has been doing.”⁸ The Inner Eye Hypothesis can easily accommodate cases of this sort, but it is not immediately apparent that they are fully compatible with the Volume Control Hypothesis.

All of the considerations that we need to take into account are now before us.⁹ Each of them seems *prima facie* to show that the Inner Eye Hypothesis is superior to the Volume Control Hypothesis. I will argue, however, that the superiority of the former is more illusory than real.

In describing the first consideration, I mentioned that we normally conceptualize the process of attending to sensations in terms of a perceptual model. I now wish to claim that nothing of much interest follows from this point. In particular, it by no means follows that it is in any deep sense correct to conceptualize attending to sensations in terms of a perceptual model, or even that we believe it to be deeply correct to do so. The fact is that we find it natural to describe a number of different phenomena in terms of perceptual models. We do not find it natural to use such models because we sense that there are deep underlying similarities between perception and the phenomena we are describing. There are similarities, of course, but they tend to be rather superficial. We notice them only because sense perception is never very far from our thoughts.

Think of a laboratory technician who is trying to determine the composition of a sample by chemical analysis. The technician may find it perfectly natural to say that he or she is “taking a closer look” at the sample. In saying this, however, the technician does not mean to assert that he or she is doing something that is fundamentally akin to what we do when we subject an object to closer visual scrutiny. When one subjects an object to closer visual scrutiny, one simply changes the relations between the object and one’s eyes. One does not change the intrinsic qualities of the object. But a technician who is analyzing a sample may well be changing many of its intrinsic qualities.

The second consideration is the perception that states of awareness that are directed on sensations are distinct from the sensations on which they are directed. As I indicated earlier, I think that this perception has to be taken seriously. At the same time, however, I believe that it is a bad mistake to see the perception as evidence for the Inner Eye Hypothesis. This is because it has no content beyond the content I have already attributed to it: it tells us that certain states of awareness are distinct from sensations, but it doesn’t tell us anything about the *nature* of those states. In order to count as evidence for the Hypothesis, the perception would have to imply that the states in question occupy a level of representation that lies between the level of sensations and the level of beliefs. But it carries no such implication. It is *consistent* with the view that the states occupy an intermediate level, but it is also consistent with the view that the states are identical with beliefs.

The perception we are now considering suggests that it is possible for sensations to exist without being objects of awareness. It should be mentioned, perhaps, that the

Volume Control Hypothesis can accommodate this idea. This can be seen by taking note of two facts. First, there is a perfectly good and perfectly familiar sense of “aware” in which we can say that to be aware of something is to know on the basis of current experience that the thing exists. Second, the Volume Control Hypothesis does not imply that a sensation must inevitably be accompanied by a belief that is directed on the sensation. It is entirely compatible with the idea that sensations can exist without being the topics of states of belief, and by the same token, it is entirely compatible with the idea that they can exist without being the topics of states of knowledge of any kind. (Because of this second fact, there is no difficulty in combining the Volume Control Hypothesis with the epistemological position I develop in the next section.)

This brings us to the third consideration, that is, to the fact that there are cases in which we find ourselves acting as if we were under the influence of sensations but in which we apparently have no sensations of the appropriate kind. It can seem that cases of this sort count heavily against the Volume Control Hypothesis. In view of the previous paragraph, however, we can see that this impression must be wrong. The Volume Control Hypothesis is fully compatible with the existence of sensations of which we are not aware.

We have found no reason to prefer the Inner Eye Hypothesis to the Volume Control Hypothesis. It turns out, however, that there is a strong reason for preferring the latter to the former. As I will now try to show, the Inner Eye Hypothesis gives a false picture of Active Introspection.

Attending to a sensation normally involves one or more *qualitative* changes. Thus, consider a case in which someone decides to focus on a sensation that has heretofore been at the margin of consciousness: if the sensation is an itch, attending to it will probably make it more importunate; if it is a pain, attending to it will probably make it more severe (at least for a while); if it is an auditory sensation, attending to it will probably increase its phenomenal volume; if it is a visual sensation, attending to it will probably increase its vividness; and so on. Or consider what it is like to attend to a new sensation: when one attends to a new itch one begins to experience itchiness; when one attends to a new pain one begins to experience pain; when one attends to a new auditory sensation one begins to have a new auditory experience; and so on. Beginning to experience itchiness is a qualitative change, and the same

is true of beginning to experience pain and beginning to have a new auditory experience.

Now the Inner Eye Hypothesis describes Active Introspection as a process that involves the following stages: first, one adjusts the orientation of one’s internal scanning device; second, the device picks up new information about a sensation somewhere in one’s phenomenal field; third, the device enters an internal state which counts as a representation of the sensation; and fourth, this representation causes one to form one or more beliefs about the sensation. (In a case in which one is attending more closely to an already familiar sensation the representation is a new and improved version of a prior representation, and in a case in which one is attending to a new sensation the representation differs in character and content from all of its immediate predecessors.) According to the Inner Eye Hypothesis, it is possible to give a complete account of Active Introspection without postulating any events beyond the events associated with these four stages. In particular, it is possible to give a complete account without claiming that Active Introspection brings about changes in the internal qualities of sensations.

As this account shows, the Inner Eye Hypothesis implies that attending to a sensation is ultimately a matter of forming a representation of the sensation. How then can the Hypothesis do justice to the fact that attending to a sensation normally involves a qualitative change? *It can do so only by claiming that the process of forming a representation of a sensation is itself a qualitative change.* It follows that the Inner Eye Hypothesis is committed to *two* levels of qualitative states. One level consists of sensations, and the other consists of the states of one’s internal scanning device that count as representations of sensations.

We are now in a position to see that the Inner Eye Hypothesis is badly flawed. First, it runs afoul of the principle of simplicity that counsels us to eschew duplications of kinds of entities unless they are forced upon us by data or by systematic considerations. In the case at hand, this principle advises us to refrain from postulating two levels of qualitative states unless doing so carries with it a substantial gain in explanatory power. Since the Inner Eye Hypothesis has no more explanatory power than the Volume Control Hypothesis, we have an obligation to set the Inner Eye Hypothesis aside.¹⁰ Second, the Hypothesis misdescribes the “location” of the qualitative changes that occur when one engages in Active Introspection. The qualitative changes

that occur when one attends to a sensation are changes in the qualitative nature of the sensation to which one is attending. (Thus, for example, as we noticed a bit earlier, attending to an auditory sensation normally brings about an increase in phenomenal volume. It is clear that a change of this sort is a change in the qualitative nature of the sensation that is the object of one's attention.) It follows immediately that the Inner Eye Hypothesis is wrong, for it implies that the qualitative changes associated with Active Introspection do not involve sensations but rather the states of one's inner scanning device by which sensations are represented.¹¹

IV

Since Descartes, philosophical discussions of introspective beliefs about sensations have tended to focus on two principles about the epistemological status of such beliefs. These principles may be expressed as follows:

It is logically necessary that if p is a proposition to the effect that S currently has a sensation with a certain phenomenal quality, and S believes that p , then it is true that p .

It is logically necessary that if p is a proposition to the effect that S currently has a sensation with a certain phenomenal quality, and it is true that p , then S believes that p .

The first principle, which I call the *Infallibility Thesis*, claims that our beliefs about the phenomenal qualities of our sensations are necessarily free from error. The other principle asserts that our knowledge of the phenomenal qualities of our sensations is necessarily complete. Following Armstrong, I call this second principle the *Self-Intimation Thesis*.¹²

In order to appreciate the appeal of these principles it is helpful to contemplate one of the main differences between perceptual beliefs about extramental phenomena and introspective beliefs about sensations. In dealing with extramental phenomena we frequently have occasion to contrast appearance with reality. The relationship between extramental phenomena and our beliefs about them involves qualitatively individuated representations (i.e., sensations) that count as appearances, and these appearances can be misleading in a number of ways. However it is at once a teaching of common sense and a consequence of the theory of

introspection presented in the previous sections that there is no appearance/reality distinction in the case of sensations. There is no set of appearances that mediate between sensations and our beliefs about them. In the case of sensations the appearance *is* the reality. This fact can seem to exclude erroneous judgements about sensations and also ignorance of their phenomenal qualities. How could we form false beliefs about sensations if there are no misleading appearances to betray us? And how could we fail to be aware of our sensations if there is no distinction between the sensations that exist and the sensations that put in an appearance?

I acknowledge the force of this argument. In my judgement, it shows that the Infallibility Thesis and the Self-Intimation Thesis both contain important elements of truth. Nevertheless, I think it is necessary to reject these principles. I will give grounds for rejecting them in the next few paragraphs, and I will then attempt to state new principles which lack their flaws but which also capture the parts of their content that are true.

In order to assess the Infallibility Thesis adequately, it is necessary to distinguish between two kinds of errors that we are prone to make. First, as we have just noticed, errors can occur when beliefs are based on appearances that fail to do justice to the entities to which the beliefs refer. When we are misled in this way by imperfect information, we make errors that may be called *errors of ignorance*. Second, errors can arise when we have adequate information about the entities with which we are concerned but we fail to take this information fully into account in forming beliefs about the entities. Errors of this sort, which may be called *errors of judgement*, are usually due either to some form of inattention or to the influence of expectation upon judgement. They tend to occur when we are hasty in forming beliefs, when we are suffering from information overload, when we are preoccupied, and when we are being lazy. They also tend to occur in situations in which anticipation causes us to lower the thresholds of application that are associated with some of our concepts.

Although we are perforce innocent of committing errors of ignorance in forming beliefs about our own sensations, we do run the risk of misconstruing our sensations by committing errors of judgement. Take the phenomenal quality *being a visual sensation that is hendecagonal in form*. One might misclassify an instance of this quality as a dodecagonal sensation if one were forced to reach a conclusion about its shape in a

hurry, and one might also do so if one were concerned with other matters at the time. Further, one might misclassify an instance if one had been given reason in advance to expect that one's next visual sensation would include a figure with twelve sides. This example involves phenomenal qualities that are fairly *complex*, that is, qualities that cannot be apprehended without doing more information processing than the bare minimum. But it is also possible to commit errors of judgement in forming beliefs about *simple* qualities of sensations, such as the quality *being a pain*. My favorite example is a case that was presented by Rogers Albritton a number of years ago in a seminar. This case involves a college student who is being initiated into a fraternity. He is shown a razor, and is then blindfolded and told that the razor will be drawn across his throat. When he feels a sensation he cries out: he believes for a split second that he is in pain. However, after contemplating the sensation for a moment, he comes to feel that it is actually an experience of some other kind. It is, he decides, a sensation of cold. And this belief is confirmed when, a bit later, the blindfold is removed and he is shown that his throat is in contact with an icicle rather than a razor.

In addition to more or less anecdotal arguments of this sort, which provide grounds for thinking that it is both logically and nomologically *possible* for there to be erroneous beliefs about sensations, there are grounds for thinking that we often form such beliefs (or at least, that we are disposed to form them) in the actual world. Thus, there is evidence — indirect, but nonetheless strong — that we are prone to misclassify sensations that are markedly similar. It is well established that subjects in experiments frequently confuse similar colors, similar tones, similar tastes, and so on.¹³ These findings do not count directly in favor of the claim that we are prone to confuse similar *sensations*, for the subjects in question are typically asked to concern themselves with external stimuli. However, their judgments about stimuli may be taken as evidence concerning their dispositions to make judgements about sensations. They would hardly be capable of confusing two stimuli if they were disposed to make fully accurate judgements concerning the sensations to which the stimuli correspond.

It appears, then, that the Infallibility Thesis is just plain wrong. What about the Self-Intimation Thesis? There are a number of reasons for thinking that it is no more worthy of our assent than the Infallibility Thesis. I will cite three of them.

First, it is intuitively correct to say that beliefs about the phenomenal qualities of sensations are *caused* by the sensations to which they refer. When I believe that I am in pain, my belief is linked to a pain by a series of events whose members are related as cause and effect. But we know that causal links are contingent in the sense that they satisfy the following principle: if a state or event *x* causes a state or event *y*, then it is logically possible for *x* to exist without being accompanied by *y*. Hence, it cannot possibly be true that sensations are accompanied by beliefs as a matter of logical necessity. Second, errors of judgement count no less heavily against the Self-Intimation Thesis than against the Infallibility Thesis. When one makes an error of judgement one forms a false belief. But more: one also *fails* to form a true belief. Failures of this sort are counterexamples to the Self-Intimation Thesis. Third, in addition to cases that involve errors of judgement, it seems that there are numerous cases in which sensations are not classified at all. Thus, as the reader can readily confirm by considering his or her own experience, inattention often keeps us from forming beliefs about a number of the details of our visual sensations. Normally, when we consider our beliefs about a recent visual experience, we find that we have firm beliefs about the components to which we have attended but only the vaguest idea as to the identities of the other components.

In addition to undermining the Self-Intimation Thesis, these arguments challenge a family of theories of the nature of consciousness that has enjoyed considerable popularity. Philosophers influenced by Descartes have sometimes urged that it is possible to analyze the characteristic *being a conscious experience* in terms of an inherently doxastic characteristic like *being an internal state that is known to exist with certainty* or *being an internal state whose distinctive nature is known with certainty* (that is, in terms of a characteristic which implies that its instances are the topics of states of belief). However, if the foregoing arguments are sound, it is entirely possible for there to be conscious experiences which have no inherently doxastic characteristics. Indeed, the second and third arguments indicate that there are *actual* experiences which lack such characteristics. It follows that it is not even possible to use one of these characteristics as the basis for a delineation of consciousness that is synthetic and empirical.

It might be thought that it is possible to save this approach to the problem of explaining consciousness by

appealing to characteristics with a modal dimension. Thus, instead of appealing to the characteristic *being an internal state that is known to exist with certainty*, one might try to base an analysis on the modal characteristic *being an internal state that CAN BE known to exist with certainty*, and instead of appealing to *being an internal state whose distinctive nature is known with certainty*, one might try to base an analysis on *being an internal state whose distinctive nature CAN BE known with certainty*. After all, we have not yet seen an argument to the effect that there are sensations that *cannot* be apprehended by their possessors.

In fact, however, there is reason to think that this amendment is inadequate. It seems that human beings are sometimes unable to apprehend their sensations because they lack the conceptual resources to do so. Consider, for example, a young baby who encounters a visual sensation that has a number of characteristics with which the baby is not yet fully familiar. Not many of us would want to maintain that the components of this sensation can be known to exist with certainty; for it seems unlikely that young babies can be said to possess the concept of existence. Nor are there many who would want to maintain that the baby is able to form beliefs which do full justice to the distinctive natures of all of the components. If one were to make this claim, one would be in the position of having to defend either the view that all of our concepts of phenomenal qualities are *a priori*, or the view that a baby can devise or refine a number of new concepts at exactly the same time as he or she is deploying the concepts in forming beliefs. Neither of these views has much intuitive appeal.

The Infallibility Thesis and the Self-Intimation Thesis are much too strong, but it would be a mistake to abandon them altogether. The appearance/reality argument that we considered at the beginning of this section shows that they are not completely lacking in merit. I suggest that they should be replaced with two principles that I call respectively the *Direct Awareness Thesis* and the *Manifest Nature Thesis*:

It is logically necessary that if x believes that y has a certain phenomenal quality, where y is one of x 's current sensations, and x 's belief is based on y , then x has not been misled by appearances.

It is logically necessary that if (i) p is a proposition to the effect that y has a certain phenomenal

quality, where y is one of x 's current sensations, (ii) it is true that p , (iii) x believes either p or the denial of p , (iv) this belief is based on y , and (v) x has not committed an error of judgement, then x believes that p .

These principles may require some minor qualifications. However, as far as I have been able to determine, they are refreshingly free from major flaws.

V

Humanity's claim to be able to obtain knowledge by sense perception has been a target for scepticism since the dawn of philosophy. In contrast, there have been very few sceptical challenges to our claim to be able to obtain knowledge by introspection. What accounts for this difference? Is it *impossible* to extend the arguments that have been formulated by sceptics so as to obtain new arguments that apply to introspective beliefs? If so, then why? If one were to try to construct a sceptical argument concerning introspective beliefs, at exactly what points would one encounter problems?

We can bring one of the main problems into focus by considering a line of thought that is frequently used to justify scepticism about sense perception. Let PB be the set of propositions that represent the perceptual beliefs of a certain normal subject, S . Further, let SH be the hypothesis that is obtained by conjoining the following four propositions: (i) S is a brain in a vat; (ii) S is connected to a computer that monitors all of S 's thoughts; (iii) all of the sense experiences S has had up to now have been caused by events inside the computer; and (iv) in the future the computer will provide experiences like the ones that S has had in the past (i.e., experiences which confirm the members of PB and which lead S to adopt new beliefs that are consistent with the members of PB). With these abbreviations in hand, we can summarize a standard sceptical argument as follows:

- (1) PB is logically incompatible with SH .
- (2) S is not justified in rejecting SH ; for S is unable to rule it out by adducing empirical evidence, and S is unable to bring forward any *a priori* arguments which count against it.

- (3) If x is justified in believing the members of a set A of propositions, and A is incompatible with the proposition p , then x is justified in rejecting p .
- (4) Hence, by (1), (2), and (3), S is not justified in believing the members of PB .
- (5) If x knows that a proposition is true, then x is justified in believing the proposition.
- (6) Hence, by (4) and (5), S cannot be said to know that the members of PB are true.

Although it seems *prima facie* that (2) is highly questionable, sceptics have managed to argue convincingly that it should be accepted. Thus, they have pointed out that all empirical evidence is ultimately sensory in nature, and that there are no sense experiences which show that the members of PB are to be preferred to SH . (SH gives rise to exactly the same retrodictions and predictions about the realm of S 's sense experience as PB .) Moreover, when their opponents have tried to rule out hypotheses like SH by *a priori* arguments, sceptics have always managed to come up with replies that are at least moderately convincing.

Now it is clear that there is no hope of using a line of thought that is fundamentally similar to (1)–(6) to justify scepticism about introspective beliefs. SH exploits the gap between the appearances on which our perceptual beliefs are based and the underlying reality to which the beliefs refer; SH would be incoherent if we could not hold appearances fixed while imagining changes in the underlying reality. It follows, of course, that there can be no counterpart of SH that applies to introspective beliefs. Thus, as the Direct Awareness Thesis informs us, there is no gap between the entities on which our introspective beliefs are based and the entities to which they refer.

Perhaps, however, it is possible to construct a different sort of sceptical argument. When one peruses the contemporary literature on the mind-body problem, one sooner or later encounters the concept of *ersatz pain*.¹⁴ Ersatz pain is described as an internal state which lacks the qualitative character of real pain but which is equivalent to real pain in point of causal relations to other internal states, to input, and to outputs. Like pain itself, it is said, ersatz pain occurs when there is bodily damage or the body is exposed to extremes of temperature or pressure. Moreover, ersatz pain can cause the same constellation of beliefs and desires as real pain, and it is therefore able to influence

behavior in the same way. Now as Sydney Shoemaker has pointed out, there is a certain amount of tension between the view that one can be said to know that his or her actual pains are real pains and the view that it is logically possible for a human being to be in an internal state that satisfies the definition of ersatz pain.¹⁵ Thus, suppose that it is logically possible for a human being to be in a state of this kind. Among other things, it follows that it is logically possible for a state other than real pain to cause S to remember the real pains that he has experienced in the past, and also that it is logically possible for a state other than real pain to cause S to think that his current state is qualitatively similar to the past experiences to which the memories in question refer. In view of these consequences, it is extremely tempting to conclude that it is logically possible for S to be in a state which he is incapable of distinguishing from real pain. But if this is true, then how can S rule out the proposition that his actual state is a state of ersatz pain? And if he is unable to rule the proposition out, then how can he claim to know that his actual state is a state of real pain?¹⁶

In the hope of bringing this line of thought into sharper focus, let us imagine that a sceptic has put forward the hypothesis (hereafter called '*EPH*') that unbeknownst to S an evil demon has recently replaced S 's disposition to have real pains with a disposition to have ersatz pains. (If you are a materialist, imagine that this transformation consists in the substitution of a bank of artificial neurons for a bank of real neurons.) According to the sceptic, S cannot claim to be justified in rejecting this hypothesis on the basis of introspection. Is the sceptic right?

What would have to be the case in order for S to be justified by introspection in rejecting *EPH*? Well, it is natural to say that S is justified in rejecting *EPH* on introspective grounds if and only if S has introspective evidence which supports the proposition that S can still experience real pain. But what is it to have introspective evidence of this sort? The answer is obvious: one can be said to have introspective evidence for a proposition about one's sensations just in case one actually has sensations that confirm the proposition. It follows that S has evidence of the right sort if and only if he is currently experiencing real pain or he has experienced real pain in the recent past. Thus, it seems that S need not construct an *argument* in order to be justified by introspection in rejecting *EPH*. He need only be in a certain state. That is, he can overturn the sceptic's

hypothesis by *being in pain*. (Here we find some possibly unexpected support for Aeschylus's contention that knowledge comes through suffering!)

It is evident that this line of thought calls the sceptic's position into question. The sceptic wants to claim that *S* is not justified by introspection in rejecting *EPH*. According to our current perspective, a defense of this claim should take the form of a proof that (despite *S*'s beliefs to the contrary) *S* is not currently experiencing real pain and has not experienced real pain in the recent past. That is to say, in order to defend the claim the sceptic must try to establish an empirical proposition about *S*'s state of mind. But to undertake this task is to abandon scepticism. The philosophical sceptic has no interest in establishing substantive propositions about the empirical world, but is rather concerned with normative questions about justification. It is the goal of the philosophical sceptic to show that we lack justification for certain of our substantive beliefs without taking a position on whether those beliefs are true or false.

It may be helpful to compare scepticism about introspective beliefs with one of the traditional forms of scepticism. Let us take another look at the situation involving *PB* and *SH*. The members of *PB* can be said to be *confirmed* by *S*'s sense experiences, but it is arguable that the former are not *selectively* confirmed by the latter. It can seem that *S*'s sense experiences have just as strong a tendency to confirm *SH* as to confirm the members of *PB*. In other words, it can seem that they are *neutral* between *SH* and the members of *PB*. Now this apparent neutrality gives the sceptic a right to maintain that *S* is not fully justified in believing the members of *PB* on the basis of sensory evidence. The sceptic can claim to have this right without denying that *S* has the sensory states that common sense attributes to him, and without trying to show that the members of *PB* are false. On the other hand, the situation involving *EPH* is quite different. It cannot be argued with any plausibility that *S*'s sensory evidence is neutral between his introspective beliefs and *EPH*. If *S* is currently experiencing pain, then he has evidence which confirms his belief that he is in pain. By the same token, he has evidence which shows (decisively) that *EPH* is false. Hence, in order to claim that *S* is not fully justified in believing that he is in pain, the sceptic must argue that this belief is not confirmed by sensory evidence. That is to say, he must deny that *S* is currently experiencing pain. But to deny this proposition is to abandon scepticism.

At this point, our sceptic might protest that we have reached a conclusion that is unfavorable to him only because we stacked the deck at the outset by accepting an excessively liberal answer to the question of what should count as a refutation of *EPH*. He might urge that it is not enough that *S* possess sensory evidence which strongly and selectively confirms the proposition that he is in pain. In addition, the sceptic might say, *S* must be capable of distinguishing between situations in which he has such evidence and situations in which he does not.

In effect, then, the sceptic is proposing that there are two conditions that one must satisfy in order to be completely justified in believing that *p*. First, there is a condition that might be called the *Evidence Condition*. According to this condition, in order to be completely justified in believing that *p*, one must possess evidence which strongly and selectively confirms the proposition that *p*. Second, there is the following *Discernibility Condition*: in order to be completely justified in believing that *p* on the basis of evidence of type ϕ , one must be able to distinguish between situations in which one has evidence of type ϕ and situations in which one lacks such evidence.

Prima facie, at least, it seems that if the sceptic can defend the Discernibility Condition, he will be in a position to conclude that *S* is not completely justified in believing that he is in pain. For the following proposition is *prima facie* correct:

- (7) *S* is incapable of distinguishing between situations in which he is in pain and situations in which he is in ersatz pain.

And (7) implies that *S* is not always capable of distinguishing between situations in which he has introspective evidence that he is in pain and situations in which he lacks such evidence. To be sure, it may be that the intuitions which cause us to favor (7) are misleading. However, as is shown by the prominence that (7) enjoys in the literature, these intuitions are strong and widely shared. Someone who wishes to reject (7) must shoulder the burden of proof.

Is it reasonable to accept the Discernibility Condition? Well, if it is to serve the sceptic's purposes, the Condition must be interpreted in such a way that it entails (8).

- (8) One is not completely justified in believing that *p* unless it is true that, for every situation σ in which one lacks evidence that *p*, if σ were actual, one

would be able to recognize that one was in a situation in which one lacked evidence that p .

Nothing weaker will do. This is because the sceptic's purposes will not be served unless S fails to satisfy the Discernibility Condition. S fails to satisfy the requirement that (8) imposes, but he *would* satisfy a less general requirement (provided that the lack of generality was not due to restrictions that were entirely *ad hoc*). After all, S has no problem in recognizing that he is not in pain in normal situations. It is only in *outré* situations like the one described by *EPH* that his ability to recognize that he is not in pain breaks down.

It is, I think, obvious that (8) is unacceptable. Take, for example, someone who is now taking a bath. Apart from sceptical arguments like (1)–(6), it seems entirely appropriate to assume that this individual is completely justified in believing that he is taking a bath. However, there are possible situations in which he lacks evidence that he is taking a bath but in which he is unable to recognize that he lacks it because Alzheimer's disease has erased the concepts that such recognition requires. Equally, there are possible situations in which he lacks evidence of the kind in question but is incapable of recognizing that he lacks it because he is unconscious. If (8) was correct, then the existence of possible situations of these sorts would prevent us from saying that our bather has evidence in the actual situation which completely justifies his belief. But this is absurd.

Is it really true that nothing weaker than (8) will do? Can't the sceptic reformulate (8) by restricting the variable ' σ ' to exclude situations in which one is either unconscious or not in possession of one's faculties? Isn't there a way of doing this that is not *ad hoc*? In particular, what about (9)?

- (9) One is not completely justified in believing that p unless it is true that, for every situation σ in which one lacks evidence that p and in which one's cognitive faculties are in good working order, if σ were actual, one would be able to recognize that one was in a situation in which one lacked evidence that p .

(9) is immune to counterexamples involving unconscious subjects and also to counterexamples involving subjects who suffer from Alzheimer's disease. Moreover, there is no justification for describing the italicized condition as *ad hoc*. But still, there is a problem. It appears that (9) allows us to say that S is completely

justified in believing that he is in pain. This is because the new restriction on ' σ ' excludes situations in which *EPH* is true. Situations in which *EPH* is true could not possibly count as situations in which S 's cognitive faculties are in good working order, for they are situations in which S has a stable disposition to misclassify certain of his current internal states and a stable disposition to misremember the intrinsic natures of certain of his past internal states. Thus, even though it is true that S would be unable to recognize *EPH*-situations if they were actual, it is impossible to get a result that would be useful to the sceptic by combining this fact with (9).

It is important to distinguish between the Discernibility Condition and another principle that it superficially resembles. According to this other principle (hereafter called the *Exclusion Principle*), in order to be completely justified in believing that p on the basis of evidence of type ϕ , one must be in a position to exclude (i.e., to rule out) all hypotheses which imply that one is *not* in possession of evidence of type ϕ . The principle may also be formulated as follows: in order to be completely justified in believing that p on the basis of evidence of type ϕ , one must be able to *determine whether or not* one has evidence of type ϕ . As this second formulation shows, it is easy to confuse the Exclusion Principle with the Discernibility Condition. However, careful scrutiny reveals that the two principles are quite different. The Exclusion Principle is a claim about one's *actual* epistemic state: it says that one must be able to use one's actual cognitive faculties and the information that is actually available to one to rule out the members of a certain set of hypotheses. On the other hand, when it is interpreted in such a way that it can be seen to imply (8), the Discernibility Condition claims that one must be capable of certain epistemic achievements in *other* possible situations — including possible situations in which one's cognitive faculties are quite different than they are in the actual situation.

Although there are philosophers who would maintain that it is too strong, to my mind the Exclusion Principle is quite plausible. Because of its similarity to the Discernibility Condition, those of us who find the Exclusion Principle plausible will often be tempted to embrace the sceptic's position. However, we can now see that it would be a mistake to allow this similarity to influence us. The similarity is only skin deep.

I will conclude by pointing out that S does in fact satisfy the Exclusion Principle in the case we have been

considering. In the case at hand, the relevant proposition is the proposition that *S* has a pain, and the relevant evidence is one of *S*'s current pains. Thus, in the case at hand, the Exclusion Principle comes to this: in order to be completely justified in believing that he has a pain on the basis of evidence which consists in his having a pain, *S* must be in a position to exclude all hypotheses which imply that *S* does *not* have a pain. Now we have already seen that *S* satisfies the Evidence Condition: he has evidence which strongly and selectively confirms the proposition that he has a pain. Because this evidence *selectively* confirms the proposition in question, *S* is able to rule out all conflicting propositions. But this means that he is able to rule out all hypotheses which imply that he does not have a pain. So he satisfies the Exclusion Principle in virtue of satisfying the Evidence Condition. Q.E.D.

It is not always true that someone who satisfies the latter requirement will also satisfy the former. Quite the contrary. This holds in the present case only because the proposition that *S* believes is identical with the proposition which describes his evidence for the belief. (That is to say, the proposition which is the value of '*p*' in the Exclusion Principle in the present case is identical with the proposition which is the value of 'one has evidence of type ϕ '.)

Notes

* I owe a large debt to Sydney Shoemaker. I have been helped considerably both by conversations with him and by the lectures he gave in the summer of 1985 in his N.E.H. Summer Seminar on Self-Consciousness and Self-Reference. I have also received valuable advice from Richard Lee, David Roach, David A. Schroeder, Lynne Spellman, and (especially) Anthony L. Brueckner, Willem de Vries, and David H. Westendorf.

¹ For example, Sartre seems to have held a view of this kind. Hazel E. Barnes describes his position as follows: "The very nature of consciousness is such, he says, that for it, to be and to know itself are one and the same Consciousness of an object is consciousness of being consciousness of an object. Thus by nature all consciousness is self-consciousness . . ." See p. xi of Barnes's 'Translator's preface' in Jean-Paul Sartre, *Being and Nothingness* (Philosophical Library, New York, 1956).

We find a closely related view in Aron Gurwitsch's *Marginal Consciousness* (Ohio University Press, Athens, Ohio, 1985). See especially pp. 3–5.

² It should be mentioned that there are contemporary philosophers who are keenly aware of the existence of an active form of introspection and of the philosophical problems it poses. D. M. Armstrong is a case in point: "Normally, introspective consciousness is of a pretty relaxed sort. The inner mental eye takes in the mental scene, but

without making any big deal about it It is, however, possible to undertake introspective scrutiny, to bend one's energies to try to discover the exact nature of what is going on in one's mind." See Armstrong's contribution to D.M. Armstrong and Norman Malcolm, *Consciousness and Causality* (Blackwell, Oxford, 1984). The quoted passage occurs on p. 120.

³ See W. B. Pillsbury, *Attention* (Macmillan, New York, 1908), 6.

⁴ The idea of drawing an analogy between attending to sensations and adjusting the volume of a radio is due to D. C. Dennett. However, my way of exploiting this analogy in building a model of attention appears to be quite different from the way that Dennett prefers. Dennett's position appears to be a kind of blend of the two hypotheses that I discuss. See 'Why you can't make a computer that feels pain', in Dennett's *Brainstorms* (Bradford Books, Cambridge, 1978), 190–229. See esp. p. 202.

⁵ Not surprisingly, there are some interesting discussions of the changes in the phenomenal field that accompany attention in the writings of such introspectionist psychologists as James, Wundt, Titchener, and Pillsbury. (Some of these discussions provide fairly strong support for my contention that the prominence of sensations can be analyzed in terms of intensity and degree of internal complexity.) See, for example, the following works: William James, *The Principles of Psychology*, Vol I (Dover Publications, New York, 1950), 402–458; Wilhelm Wundt, *Outlines of Psychology*, 3rd English edition (Alfred Kroner, Leipzig: 1907), 228–251; Edward B. Titchener, *A Text-Book of Psychology*, revised edition (Macmillan, New York, 1909), 53–54 and 265–83; and W. B. Pillsbury, *op. cit.* 2–10. (For a survey of the literature about one of the key concepts in 19th century discussions of attention, see I. M. Bentley, 'The psychological meaning of clearness', *Mind* XIII (1904), 242–253.)

⁶ Given that contemporary psychologists are often reluctant to become deeply involved in issues concerning qualitatively individuated states, it is natural to assume that it will be some time before modern psychology provides any information about the information-processing mechanisms that underlie Active Introspection. However, this assumption is not quite true: cognitive psychologists are concerned to explain what is involved in attending to *extramental* objects and events, and some of their findings have an indirect bearing upon questions about attending to sensations. According to cognitive psychology, attending to an extramental entity should be seen as a matter of committing a sense receptor and one or more information-processing mechanisms to the task of obtaining a detailed and trustworthy representation of the entity. (See, for example, the following papers: Donald E. Broadbent, 'Task combination and selective intake of information', *Acta Psychologica* 50 (1982), 253–290; Daniel Kahneman and Anne Triesman, 'Changing views of attention and automaticity', in Raja Parasuraman and D. R. Davies (eds.), *Varieties of Attention* (Academic Press, New York, 1984), 29–61; William A. Johnston and Veronica J. Dark, 'Selective attention', *Annual Review of Psychology* 37 (1986), 43–75.) Since the task of obtaining a detailed and trustworthy representation of an extramental entity consists partly in obtaining a prominent sensation that represents the entity, and since changes in prominence are partially constitutive of attending to sensations, it is appropriate to say that theories about the information-processing mechanisms that underlie attending to extramental objects and events are implicitly germane to certain aspects of Active Introspection.

⁷ In view of the popularity which the Inner Eye Hypothesis and closely related ideas have enjoyed in the history of philosophy, it is hardly surprising that there are considerations which suggest *prima facie* that the Hypothesis is true. (As is well known, Leibniz argued for the existence of a quasi-perceptual relation between minds and experiences that he called *apperception*, and Locke and Kant explicitly likened introspection to sense perception.)

⁸ See D. M. Armstrong, *The Nature of Mind* (Cornell University Press, Ithaca, 1981), p. 59.

⁹ A more comprehensive treatment would no doubt mention two additional considerations. First, there is a line of thought that Leibniz sometimes used in defending his doctrine of *petites perceptions*. In order to hear the noise of a multitude of waves, he maintained, we must have sense experiences corresponding to the individual constituents of the multitude. However, we cannot be said to be aware of these experiences considered in isolation from their fellows, for as individuals they are too "small" to be detected. Second, someone might attempt to justify the Inner Eye Hypothesis by urging that it is better equipped than the Volume Control Hypothesis to explain the phenomenon known as *blind sight*. (Patients with a certain neurological disorder claim that they are unable to see objects that are held in front of them, but they are nonetheless able to point to the objects in question. For a brief account of this interesting phenomenon, see Randolph Blake and Robert Sekuler, *Perception* (Alfred A. Knopf, New York, 1985), pp. 115–116.)

I do not discuss Leibniz's argument because it depends on a tacit assumption that seems to me to be highly questionable (namely, the assumption that an experience which represents an extramental entity must have components which correspond to all of the components of the entity). As for blind sight, I do not discuss it because I do not see that it points to *sensations* of whose existence we are not aware. The data can be fully explained by supposing that pointing behavior can be guided by unconscious packets of information like the ones postulated by the Volume Control Hypothesis.

¹⁰ In an earlier paper I sketched an argument against Occam's Razor which, if it were sound, could be used to cast doubt on appeals to simplicity like the one in my objection to the Inner Eye Hypothesis. (See the fifth footnote of 'In defense of type materialism', *Synthese* 59 (1984), 295–320.) However, I now question the relevance of the example on which my earlier argument is based.

¹¹ There is also a third objection to the Inner Eye Hypothesis, namely, that it conflicts with a stable intuition about what it is that *confirms* our introspective beliefs. It preserves the intuition that our beliefs about sensations are confirmed by qualitative states. However, it implies that the qualitative states in question are the states of internal scanning devices which count as representations of sensations. This implication is at variance with the intuition that beliefs about sensations are confirmed by the sensations to which they refer.

¹² See D. M. Armstrong, *A Materialist Theory of the Mind* (Routledge and Kegan Paul, London, 1968), p. 101. (Armstrong tells us that he borrowed the term from Ryle.)

Despite the fact that Armstrong's theory about the nature of introspection appears to be quite different from the theory I develop in Sections I–III (he favors a version of the Inner Eye Hypothesis), his views about the epistemological status of introspective beliefs are quite similar to the views that are expressed in Section IV. (I have

found it extremely helpful to consult his writings on introspection in developing my own position.)

¹³ Psychophysicists have long found it necessary to allow for a variety of errors in establishing absolute thresholds and difference thresholds. Here is a typical observation: "One sticky problem, though, for the concept of the threshold is that of judgement errors. All of the psychophysical methods that have been discussed [in this text] have some procedure for balancing out 'errors' in judgement that may interfere with the observer's ability to report accurately his sensory experiences." See Ronald H. Forgas and Lawrence E. Melamed, *Perception*, Second Edition (McGraw-Hill, New York, 1976), p. 38.

¹⁴ See, for example, Ned Block, 'Are absent qualia impossible?', *The Philosophical Review* LXXXIX (1980), 257–274. In Note 4 Block attributes the term "ersatz pain" to Larry Davis.

¹⁵ See Sydney Shoemaker, 'Functionalism and qualia', *Philosophical Studies* 27 (1975), 291–315. See also Shoemaker's 'Absent qualia are impossible', *The Philosophical Review* XC (1981), 581–599.

¹⁶ Shoemaker has used an argument that is closely related to this one in defending the view that pain is a functional state, that is, an internal state that is definable in terms of its causal relations to other internal states, to inputs, and to outputs. (See the papers cited in Note 18.) His argument runs roughly as follows:

- (1) If pain is not a functional state, then it is logically possible for ersatz pain to exist (i.e., it is logically possible for a human being to be in an internal state that satisfies the definition of ersatz pain).
- (2) If it is logically possible for an individual to be in an internal state that satisfies the definition of ersatz pain, then it is possible for the individual to be in a state other than a state of real pain which he or she is unable to distinguish from a state of real pain.
- (3) If it is possible for an individual to be in a state other than a state of real pain which he or she is unable to distinguish from a state of real pain, then he or she does not know that he or she is experiencing real pain.
- (4) When a human being is in a state of real pain, he or she knows that he or she is in a state of real pain.
- (5) Hence, by (1)–(4), pain is a functional state.

It seems to me that premise (3) is highly questionable. (As the reader will observe, my reply to the sceptic's argument is based on an objection to a closely related premise.)

Department of Philosophy
127 Kimpel Hall
University of Arkansas
Fayetteville, AR 72701, U.S.A.