

CHAPTER I

Forms of consciousness

It is customary to distinguish five forms of consciousness: *agent consciousness* (which is what we have in mind when we say of an agent that he is “losing consciousness” or “regaining consciousness”), *propositional consciousness* (which is expressed by the “conscious *that*” construction), *introspective consciousness* (which is what we have in mind when we say things like “His affection for me is fully conscious, but his hostility is not”), *relational consciousness* (which is expressed by the “conscious *of*” construction), and *phenomenal consciousness* (which is a property that mental states possess when they have a phenomenological dimension – that is, when they present us with such qualitative characteristics as pain and the taste of oranges).

I will have something to say about all of these forms of consciousness in the present work, though some of them will receive much more attention than others. To be more specific, I will have very little to say about agent consciousness and propositional consciousness, for insofar as the philosophical problems associated with these two forms of consciousness are problems of mind (as opposed to problems associated with agency and problems associated with knowledge), they are reducible to problems that arise in connection with other forms of consciousness. They are not in need of separate treatment. I will have more to say about all of the three remaining forms, but one of them, phenomenal consciousness, will be considered at much greater length than the others. The reason for this inequity is that phenomenal consciousness has a disproportionately large ability to produce philosophical puzzlement. Historically the task of explaining phenomenal consciousness has been thought to be the most challenging responsibility in philosophy of mind, and perhaps even the most challenging responsibility in all of philosophy. This view is also widely held by contemporary philosophers.

In addition to the five forms of consciousness that have just been identified, I will also discuss what I take to be a sixth form – a form that can reasonably be called *experiential consciousness*.

It seems that we have at least two notions of experience. One of these applies to mental states that have proprietary phenomenology, and therefore comes to much the same thing as phenomenal consciousness. The other notion of experience has a more general significance. It applies to states with a proprietary phenomenology, just as the first notion does, but it also applies to thoughts, judgments, suppositions, volitions, and all other mental states that count as *occurrent propositional attitudes*. (As is customary, I use the term “propositional attitude” to refer to mental states that can be described by verbs that take a sentential complement. Believing is a mental state of this sort, because beliefs can be described by combining “believes that” with a declarative sentence. I say that a propositional attitude is occurrent if it is an event. A thought is an event, and so is a volition, but most beliefs and desires are enduring states that remain with one even when one is asleep.)

Now when we reflect on the broader notion of experience, we find, I suggest, that it entails that the items to which it applies are conscious. To be an experience is to be conscious. This intuition is ratified by dictionaries. Thus, the third entry for “experience” in my *Webster’s* says that experiences are “the conscious events that make up an individual’s life,”¹ and the fourth entry in the *OED* says that to have an experience is to be “consciously the subject of a state or condition.”² I will take this testimony at face value – experiential consciousness really is a form of consciousness. And I will assume that experiential consciousness cannot be reduced to any of the other forms. It poses problems of its own and requires separate treatment.

Finally, it is sometimes maintained that there is a seventh form of consciousness, *access consciousness*, which is possessed by all and only those higher mental states that are poised for control of speech, reasoning, and intentional action.³ There will be some consideration of this proposal in the present chapter, but it will not figure prominently in later discussions. For reasons that will emerge, the standard characterizations of access consciousness are best seen as approximations to a correct account of experiential consciousness. They are suggestive, but problems emerge if we try to think of them as picking out a form of consciousness that is independent of the others.

¹ *Webster’s Ninth New Collegiate Dictionary* (Springfield, MA: Merriam-Webster Inc., 1989), p. 437.

² *The Oxford English Dictionary*, 2nd edition (Oxford: Oxford University Press, 1989), Volume 5, p. 563.

³ Ned Block, “On a Confusion about a Function of Consciousness,” *Behavioral and Brain Sciences*, 18 (1995), 227–247. Reprinted with some changes in Ned Block, Owen Flanagan, and Güven Güzeldere (eds.), *The Nature of Consciousness* (Cambridge, MA: MIT Press, 1997), 375–415.

I have been speaking thus far of different *forms* of consciousness, but it is in some ways preferable to speak of different *concepts* or *notions* of consciousness. I will often adopt the latter mode of expression. This will be done partly for stylistic reasons, but also with a view to highlighting the fact that our concepts of consciousness may partially misrepresent the phenomena to which they refer.

The ensuing sections of this chapter will expand on the present descriptions of the various forms of consciousness, characterize some of the problems to which they give rise, and sketch the main themes of later chapters.

1.1 AGENT CONSCIOUSNESS

Agent consciousness is a property that adult human beings possess throughout their waking lives, and also when they are dreaming. It can be possessed by agents other than adult humans, but it presupposes the ability to have experiences and to engage in various forms of reasonably high level cognitive activity. We are sure that a slug does not enjoy agent consciousness, and we are reluctant to attribute it to insects.

Consider an agent who is waking up, or who is coming out of a coma. We describe the agent as regaining consciousness. In what exactly does this transition consist? Surely what we have in mind, when we say that the agent is regaining consciousness, is that he is starting to think and feel again, to perceive the world, and to experience bodily sensations. If conscious states of these sorts were not occurring, then, I suggest, the concept of agent consciousness could not get a foothold. In short, it seems that enjoying conscious mental states is a *necessary* condition of agent consciousness. But it is also a *sufficient* condition. It would be absurd to deny consciousness of someone who is consciously thinking about a topic, or consciously perceiving an object or an event. When there is a stream of conscious events, or even a tiny rivulet, agent consciousness is necessarily present as well.

In view of these facts, we can appreciate that the notion of agent consciousness is not a basic or foundational concept. It can be explained in terms of the concepts we use to attribute consciousness to individual mental states. It appears that there are just three concepts that have this role – the concept of introspective consciousness, the concept of phenomenal consciousness, and the concept of experiential consciousness. Accordingly, it is possible to explain agent consciousness by saying that an agent is conscious just in case he is in one or more mental states that are introspectively conscious or phenomenally conscious or experientially conscious.

It is sometimes maintained that agent consciousness comes in degrees, and that this fact should be explicitly recognized in any explanation of its nature. Insofar as this is the case, however, it can be accommodated by appealing to the number, variety, and internal complexity of conscious states that an agent is enjoying. When someone says, for example, that Bill is gradually losing consciousness, what the speaker seems to mean is that there is a decrease in the number and/or variety of Bill's conscious states, and/or that his experiences are less complex. Equally, an increase in the level of one's consciousness is a swelling of one's stream of consciousness – that is, an increase in the number of events in the stream, or in the number of kinds of events in the stream, or in the complexity of the individual events.

1.2 PROPOSITIONAL CONSCIOUSNESS

Attributions of propositional consciousness are tantamount to attributions of propositional *knowledge*. Thus, generally speaking, when it is appropriate to say that someone is conscious that *p*, it also appropriate to say that he knows that *p*; and when it is appropriate to say that someone knows that *p*, it is also appropriate to say that he is conscious that *p*.

This equivalence thesis is occasionally challenged on the grounds that propositional consciousness is a *special kind* of propositional knowledge. More specifically, it is sometimes said that propositional consciousness is propositional knowledge that is *active* or *operative* – propositional knowledge that is currently on line and immediately available for use by a range of high level cognitive faculties. But this is not true. Thus, it can make perfectly good sense to say that a soldier is conscious that there will be a battle tomorrow even though the soldier is asleep. Further, consider the following exchange:

JACK: Is Bill conscious of the fact that Mary dislikes him?

JILL: As far as I can tell, Bill doesn't think about Mary much these days, if indeed he thinks about her at all; but yes, I'm sure he's still conscious that she dislikes him. It's hardly the kind of thing that he would forget.

Jill's reply to Jack strikes us as a bit awkward – we would find the exchange a bit more natural if Jill had said that Bill *knows* that Mary dislikes him. Even so, however, we have no trouble finding an interpretation of the exchange on which Jill's reply is literally true. And of course, if the reply is literally true, then it cannot be the case that “conscious that” is used only to attribute *operative* knowledge.

It appears, then, that this objection to the equivalence thesis is mistaken. But it is understandable, for there is a secondary sense of “conscious that” on which it does carry the suggestion that knowledge is currently operative. Thus, it would be quite natural to say the following: “Bill knows that he’s supposed to be home by midnight, but he isn’t presently *conscious* that this is so – at present he has no mind for anything but the music and his dancing.” Statements of this sort are often regarded as true. But this could not be the case unless “conscious that” *can* be used to attribute active knowledge. Accordingly, it must be the case that there are two senses of “conscious that,” a primary one on which it is equivalent to “knows that,” and a secondary one on which it has a more narrowly circumscribed use. (I say that the latter sense is secondary because it is generally necessary to stress “conscious” in order to bring it to the fore.)

With these observations in mind, we can easily see that the primary sense of “conscious that” presents no special challenges to someone who is trying to understand consciousness. Of course, it presents plenty of interesting problems to someone who is trying to understand *knowledge*; but a philosopher of consciousness is interested principally in the nature of the mind, not in the ability to acquire knowledge of the world. Hence, unless he is also an epistemologist, a philosopher of consciousness should set the primary sense of “conscious that” aside as irrelevant to his main concerns.

The same applies, with a qualification, to the secondary sense of “conscious that.” To understand the nature of *active* propositional knowledge, one should focus principally on propositional knowledge. But one should also consider what it is for a mental state to be active or operative. Because of this, the task of explaining the secondary sense is not the sole responsibility of epistemologists. Part of it belongs to philosophers of mind. Even so, however, I think that the philosopher of consciousness can set the secondary sense of “conscious that” aside. Insofar as it involves the notion of an active mental state, it presents some problems that are relevant to his main concerns. But as we will see, those problems arise independently in connection with experiential consciousness and access consciousness. These are no additional problems that are proprietary to the secondary sense of “conscious that.”

1.3 INTROSPECTIVE CONSCIOUSNESS

A mental state counts as introspectively conscious if it is *actually* an object of introspective awareness, but it seems that we are also willing to ascribe introspective consciousness to mental states if the relevant agents could

easily *become* aware of them, whether by simply redirecting their attention, or by asking a question like “What exactly is it that I am perceiving now?” Thus, we are quite prepared to say something like, “All of John’s feelings about his sister are conscious, though some of his feelings about his brother are repressed.” When we say something of this sort, we don’t mean to imply that all of the mental states to which we are attributing consciousness are objects of explicit introspective awareness. Rather, we mean that the states would easily become objects of awareness if the agent would turn his attention to them. In general, we distinguish between *actual introspective consciousness* and *potential introspective consciousness*, and we apply this distinction both to mental events and to continuing states like one’s attitudes toward one’s siblings.

It seems that the states that are characterized by actual introspective consciousness are highly circumscribed, at least in comparison to the states that have potential introspective consciousness. Thus, there is very little introspective evidence for the idea that we are always actively monitoring our mental processes, taking explicit note of their nature, and it seems very unlikely that the brain would find it advantageous to expend energy on constant monitoring of this sort. On the other hand, it seems likely that potential introspective consciousness is distributed quite widely. Indeed, in adult human beings, at any rate, it may well be true that all experiences, including both occurrent propositional attitudes and events with a qualitative dimension, could easily become objects of introspective awareness. Still, actual introspective consciousness is more fundamental than potential introspective consciousness, metaphysically speaking, for potential consciousness exists only when actual consciousness might occur. Accordingly, I will mainly focus on actual consciousness in the present work.

There are two varieties of actual introspective consciousness – actual consciousness of mental *occurrences*, and actual consciousness of *enduring* mental states that are stored and static. Actual consciousness of occurrences takes place, for example, when an agent takes note of a passing thought, and also when an agent judges that he is perceiving an object of a certain sort. On the other hand, there is actual consciousness of a stored state when an agent judges that he believes (and has believed) that Albany is the capital of New York. Now it might seem at first that awareness of continuing states is quite different than awareness of occurrent states; but closer consideration shows that there is reason to think that these two forms of awareness are closely related. Thus, it seems likely that it is necessary to activate or “refresh” a continuing state in order to appreciate

its existence introspectively. For example, in order to appreciate that one has the given belief about Albany, it seems necessary to activate the belief in some way – perhaps by asking the question, “What is the capital of New York?” Once this state has been activated, one can move from the resulting occurrent manifestation of it (e.g., the answer, “Albany!”) to an introspective judgment that acknowledges the continuing state. If this is right, then the present case is similar to cases in which one achieves actual, occurrent awareness of an occurrent state. In the former as in the latter, one moves from an occurrent mental state to an introspective judgment.

In addition to showing that actual introspective consciousness of enduring states is closely related to actual introspective consciousness of mental occurrences, this line of thought shows that the former in fact depends on the latter. Introspective consciousness of occurrences is more fundamental than introspective consciousness of enduring states. Because of this dependency, in the present work I will be concerned principally, though not exclusively, with consciousness of occurrences.

Introspective awareness of a mental state is independent of information about the external world, and it is also immediate, in the sense that it does not derive from chains of reasoning. If one arrives at a belief about one of one’s mental states via inferences from one’s behavior, or because a therapist has presented one with a theory of one’s states that convincingly explains certain of one’s dreams and emotional experiences, the belief does not qualify as introspective. It would probably be a mistake, however, to join various writers in maintaining that introspective beliefs must be entirely non-inferential. Generally speaking, achieving introspective awareness of a mental state is a matter of passing from the state itself to a judgment about the state. It is natural to think of transitions of this sort as inferences.

As we have noted, a state counts as introspectively conscious just in case it is an object of introspective awareness. It follows that questions about introspective consciousness are principally questions about the nature of introspection. One of the main questions about introspection is concerned with the nature of the vehicles of introspective awareness – the meta-cognitive states that represent or register our first order states. A number of contemporary philosophers endorse the view that these vehicles are perceptual in character, or at least quasi-perceptual, while others contend that they involve conceptualization and take the form of judgments. This controversy will receive considerable attention in a later chapter. Another very important question about introspection concerns the processes by which introspective awareness is produced. At first sight, at least, these processes seem to be highly variegated. Consider, for example,

a case in which you judge that you are currently perceiving a red square. It is, I think, quite natural to suppose that this introspective judgment is produced by a relatively straightforward inferential process – a process that is similar to the one that produces the non-introspective judgment that there is a red object in front of you. Now consider a case in which I ask you whether you think that the population of London is larger than that of New York. Suppose that an answer to this question is stored somewhere in your memory, and that, once it is activated, the memory leads you to self-ascribe the belief that London is the larger city. How exactly did you arrive at this judgment? Not, it seems, by a direct inference that has roughly the same form as perceptual inferences, but rather by searching various memory files by a Google-like procedure. More particularly, it seems that my question primes a mental search engine with the key words “London,” “New York,” and “population,” and then sets it in motion. These two examples pose the following question: Is introspection highly multiform, in the way that the examples suggest, or do the various introspective processes have a common nature that comes into view when one considers them more closely? As with the question about vehicles of introspection, we will be examining this question about introspective processes at some length.

I conclude this introductory discussion by noting that introspective consciousness is of fundamental scientific importance. This was originally emphasized by Freud. Many components of the Freudian model of the mind are now widely, and correctly, dismissed as pseudoscientific, but the notions of actual and potential introspective consciousness continue to play substantial roles in psychology. The mental states that are accessible to introspection are precisely the ones that can be reported in speech (in subjects with normal linguistic capacities, at any rate), and it is the ability of an agent to report his mental states that provides the principal evidence for scientific theories of consciousness. Moreover, there is reason to think that states that are introspectively conscious are also accessible to a large range of high level cognitive faculties, including those that are responsible for reasoning and for control of behavior. It seems that potential introspective consciousness marks a causal joint in the mind, a locus of causal relevance and authority that is of central importance in cognitive processing.

I.4 EXPERIENTIAL CONSCIOUSNESS

It is important to get a fix on this form of consciousness in the early going, for it turns out to be highly relevant to each of the remaining

forms. To understand them, it is necessary to understand their respective relationships to experiential consciousness.

As we saw, experiential mental states fall into one or the other of two large categories – occurrent propositional attitudes and states with a proprietary phenomenology. Thoughts, volitions, and passing wishes count as experiences, but so do perceptual states, experiences of pain, and mental images. Now the members of these two categories are quite different from one another. Thus, occurrent attitudes have propositional objects that have conceptual contents and are individuated by logical forms, but this seems not to be true of states with a proprietary phenomenology (hereafter *P-states*). For example, while it seems to be true that a judgment of the form *if p then q* is different from the corresponding judgment of the form *either not-p or q*, we would not be inclined to say that the contents of perceptual experiences are individuated by logical differences of this sort. Another important difference derives from the phenomenological or qualitative nature of members of the second category. There is reason to doubt that thoughts, volitions, and other occurrent attitudes are individuated phenomenologically. When one is aware of a thought one is necessarily aware of its content, but it seems that one is not aware of any properties like the ones that are usually cited as paradigms of phenomenology – pain, the way it feels to be angry, the way yellow things look, the way oranges taste, and so on. Accordingly, the fact that P-states have a qualitative nature amounts to a metaphysical difference of substantial importance. In view of these differences, it is *prima facie* quite puzzling that we group members of the two categories together under a single concept, calling them all “experiences.” What, if anything, is shared by all of the states that we classify in this way?

There are just three possibilities. It might be that occurrent attitudes and P-states share an intrinsic feature that is revealed by introspection. Perhaps they have a certain halo or phosphorescence. It also might be that we think of them as having similar relationships to the agencies that are responsible for introspective awareness, that we count them all as conscious experiences because we can become conscious *of* them in similar ways, and/or to similar degrees. Finally, it might be that they bear similar relations to a range of high level cognitive faculties – a range that includes the faculties responsible for introspective awareness, but a number of others as well. The thought here is that if a mental state is to count as an experience, it must be available or accessible to several of the faculties in the range, but that it need not be available to all of them. Thus, according to this suggestion, if a mental state occurs in a creature that lacks the

capacity for introspective awareness, it can still have experiential status, provided that it is available to other high level faculties, such as the ones responsible for forming beliefs and desires. I will discuss each of these possibilities in turn.⁴

When you consider an experiential state introspectively, are you aware of an intrinsic phosphorescence that the state shares with all other experiences? When you consider a thought introspectively, for example, are you aware of it as having an intrinsic property that it shares with passing wishes, perceptual states, and experiences of pain? My guess is that your answer to these questions will be “no.” Certainly that is the right answer in my case. When I attend introspectively to a thought, I am aware of content-related properties that distinguish it from other thoughts, and also of properties, such as the property *being a thought*, that it shares with other thoughts, but I cannot discriminate an intrinsic feature that I can also identify as present when I introspectively consider experiences of quite different kinds, such as a perceptual experience of a green leaf or an experience of pain in my right foot. If there is a phosphorescence that all experiences share, it is invisible, hiding shyly from introspection “behind” the more straightforward properties to which we have access. But no such property could give us a *reason* to group all experiences together under a single concept.

What about the second possibility? Might it be true that all experiences enjoy a special relationship with the agencies that are responsible

⁴ There are other ways in which one might try to explain experiential consciousness, but I regard them as being significantly less plausible than the three I cite in the text. One alternative that should at least be mentioned is the view that states with experiential consciousness are so endowed because they involve reflexive awareness – that is, the view that each experientially conscious state provides its subject with awareness of that very state. The idea that consciousness is reflexive is currently enjoying a vogue (see, e.g., Uriah Kriegel and Kenneth Williford (eds.), *Self-Representational Approaches to Consciousness* (Cambridge, MA: MIT Press, 2006)), but I can see no merit in it, especially when, as is usually the case nowadays, it is explained in terms of self-representation. Here is a short argument in support of my view: “Generally speaking, when a mental state *M* represents an item *x*, it is because it is useful for the cognitive agencies that deploy *M* to have information about *x*. Hence, if a mental state *M* represented itself in addition to representing an external state of affairs, it would be because *M* was used by at least two different cognitive agencies – one that performed a world-oriented task (e.g., elaborating plans), and therefore required information about the world, and another that performed a meta-cognitive task (e.g., assessing evidence for judgments), and therefore required information about mental states. We have reason to believe, however, that meta-cognitive agencies have their own proprietary representations – that is, representations that are like *I see a blue sloop* and *I am thinking about New York* in that they have constituents with *explicit* psychological content. (See Chapter 8.) Assuming that this is correct, it is unclear why meta-cognitive agencies should have any use for additional representations of mental states. By the same token, it seems that it would serve no purpose to suppose that the representations that are concerned with the world (that is, representations such as perceptual experiences) also have a second layer of representational content that is self-referential. It seems unlikely that this supposition could do any useful explanatory work.”