FEELINGS[1]

DESPITE the recent popularity of philosophy of mind, the concept of feeling has thus far received short shrift. It has functioned largely as a foil to the more dispositional concepts—desire, intention, belief—that have been in the forefront of attention. Thus several recent authors have been at pains to deny that an emotion like fear or embarrassment is "just a feeling" and have insisted that emotions contain cognitive and motivational elements as well.[2] Again, it has been insisted that to want something is not just to have a certain kind of feeling, and that to perform an action because one wants something is not to have the action caused by a certain kind of feeling.[3] Little is said, however, as to what sort of thing it is from which emotions and wants are being distinguished. It is generally recognized that a feeling is something momentary or "occurrent" rather than dispositional, and that it is the sort of thing to which its possessor has privileged epistemological access. But however true this may be it does nothing to distinguish feelings from thoughts and mental imaginings, of which the same could as well be said.[4] We are left without any answers to such questions as:

[1] An earlier version of this paper was written while the author was a Fellow of the Center for Advanced Study in the Behavioral Sciences, Stanford, California. The author is indebted to Barry Stroud, Alvin Goldman, and Laurence Davis for enlightening comments on earlier versions of the paper.

[2] See Errol Bedford, "Emotions," *Proc. Arist. Soc.* (1956-1957); George Pitcher, "Emotion," *Mind*, LXXIV (1956); Anthony Kenny, *Action, Emotion and Will* (London, 1963), ch. 3; Magda Arnold, *Emotion and Personality* (New York, 1960); S. Schachter and J. E. Singer, "Cognitive, Social, and Physiological Determinants of Emotional State," *Psychological Review* (1962).

[3] See A. I. Melden, *Free Action* (London, 1961); Gilbert Ryle, *The Concept of Mind* (London, 1959); R. S. Peters, *The Concept of Motivation* (London, 1958).

[4] The only recent work by an analytic philosopher that seems to deal with this problem is *The Concept of Mind*. There Ryle says: "By 'feelings' I refer to the sorts of things which people often describe as thrills, twinges, pangs, throbs, wrenches, itches, prickings, chills, glows, loads, qualms, hankerings, curdlings, sinkings, tensions, gnawings, and shocks. Ordinarily, when people report the occurrence of a feeling, they do so in a phrase like 'a throb of

(1) What generally distinguishes feelings from other kinds of conscious states?
(2) How are different kinds of feelings distinguished from each other?
(3) How are feelings related (conceptually) to other kinds of mental states?
(4) How is an intersubjective feeling vocabulary possible?

In addition to the intrinsic interest in answering such questions, and in addition to their importance for the construction of a logical map of the mind, we are hardly in a position to determine, for example, whether the fact that embarrassment contains cognitive and motivational elements prevents it from being "just a feeling," so long as we are left in the dark as to what sort of thing a feeling is; perhaps feelings contain cognitive and motivational elements, too.

The word "feel" is, of course, notoriously a sprawling one. In his article, "Feelings,"[5] Gilbert Ryle distinguishes seven "uses" of the term, while suggesting that there are many more. These include the *perceptual* use (feel the rope around one's neck), the *exploratory* use (feel for the matches in my pocket), the *localized sensation* use (feel a tickle), the *general condition* use (feel sleepy, uneasy, ill), the *propositional* use (feel that a thunderstorm is

compassion,' 'a shock of surprise' or 'a thrill of anticipation' " (p. 84). Again, "Feelings, in any strict sense, are things that come and go or wax and wane in a few seconds; they stab or they grumble; we feel them all over us or else in a particular part" (p. 100). Now, however accurate this may be as a specification of what Ryle refers to by "feelings," it can hardly be taken seriously as an attempt to deal with the full range of feeling terms in ordinary language. Feeling contented, feeling adventurous, feeling ill at ease, or feeling upset are not often described as thrills, twinges, etc. That Ryle is in earnest with his restrictions on the term "feeling" is brought out by the fact that in order to show that to do something from vanity is not to have one's action caused by a certain kind of feeling, he argues that such actions are not invariably or usually accompanied by certain typical thrills, pangs, or prickings (pp. 85-87), and by the fact that he denies that "I feel lazy," "I feel depressed," and "I feel energetic" "report the occurrence of feelings" (since nothing like a thrill, twinge, or throb is necessarily involved) (p. 101). Thus Ryle bypasses the problem of bringing out the distinctive force of just such locutions as "I feel depressed." This will be our concern.

[5] *Philosophical Quarterly*, Vol. I., 1951.

brewing), and the *action readiness* use (feel like taking a walk). (Some of these labels are mine.) In this essay I shall not be concerned with the facts reported by locutions of all these kinds, though I should hope eventually to bring them all within a unified account. For the present I shall restrict myself to what I shall call "adjectival feelings," feelings reported by phrases of the form "feel *f*," where "*f*" is some adjective or adjectival phrase the nominal form of which designates a kind of state of a person.[6] The category of adjectival feelings is a wide one. We might, with no pretensions to either precision or exhaustiveness, dissect it into *emotional* feelings (feel angry, annoyed, indignant, frightened, anxious, disturbed, embarrassed, elated, grief-stricken, depressed), *mood* feelings (feel cheerful, gloomy, tranquil), feelings of *general bodily condition* (feel tense, sleepy, hungry, energetic), and feelings of *behavioral tendency* (feel generous, adventurous, amorous, talkative). This category excludes most of Ryle's "uses"; only his "general condition" use falls within it, or is identical with it as the case may be. It will be my aspiration in this paper to develop a pattern of analysis for adjectival feeling terms.

Some adjectival feelings, particularly the emotional ones, have intentional objects; others—for example, feeling sleepy or tense—do not. Thus one may feel angry *at his brother*, feel embarrassed *over having forgotten to send the letter*, and so on. In these cases one necessary condition for having the feeling in question is that the person have certain beliefs—for example, that he have a brother or that he has forgotten to send the letter. (These may not be the *only* beliefs required.) I shall ignore this aspect of feelings in my discussion and concentrate instead on the question of what else is involved in having a certain feeling, over and above the beliefs that are required for the feeling having a certain intentional object. This means that the schema of analysis at which I shall arrive will have to be supplemented by the requirement of certain beliefs in cases where this is necessary. I shall say no more about intentional objects in this paper.

[6] I shall use the upper-case "*F*" as a dummy for the noun forms ("anger"), the lower-case "*f*" as a dummy for the adjective forms ("angry").

One further preliminary restriction is necessary. Feeling terms are used both for occurrent states and for dispositions. The locutions "I feel sorry for him" and "I feel terribly resentful toward him over that incident" can be, and normally are, so used that they do not report a state of consciousness possessed by the person at the time of utterance or at any other specified time. Rather they are used to report a complex of dispositions. Thus to feel sorry for x is to be disposed to feel sad when one reflects on x's plight, to make efforts to help him if one can think of anything helpful to do and one is not too strongly motivated in opposite directions, and so on. These locutions contrast with "I feel very depressed today" and "When he said that I felt terribly embarrassed," which do imply that the speaker was in a certain state of consciousness at a certain time. Other locutions are ambiguous as between the two interpretations—for example, "I feel uneasy about the outcome of the meeting." In this paper we are concerned with feelings as datable states of consciousness.

I

It is a notable fact, though one too seldom noted, that adjectival feelings are standardly denominated by a phrase that conjoins the word "feel" with a term, the nominal form of which designates some state of a person. Now it is clear that we cannot make a general identification of the feeling with the state. In general they are not even extensionally equivalent; it is possible for one to be tired without feeling tired, and to feel tired without being tired (according to objective indexes of fatigue).[7] The term "f," however, plays an essential role in our specification of what

[7] In some cases it may be doubted that it is possible to feel f without being f. Thus it may be said that if I do feel angry at y, then, whatever is the case otherwise, I am (at least for the moment and at least a little) angry at y, though possibly not with the persistence or depth or intensity that I suppose. This may hold generally for emotional f's—embarrassed, depressed, afraid, excited, etc. On the other hand, even in these cases it seems possible to be f without feeling f. I may be quite angry at x but "not let myself realize it." In any event, the only point I am concerned to make now is that there is no general extensional equivalence between being f and feeling f.

feeling we are talking about when we use a phrase of the form "feel *f*"; this does suggest that the concept of *F* is somehow an ingredient of the concept of feeling *f*, or, otherwise stated, that the concept of feeling *f* is derivative from the concept of *F*. Let us call this the "Dependency Thesis."

It is an equally notable fact that the views on the nature of adjectival feelings that seem most plausible to most people are quite opposed to the Dependency Thesis. The one body of systematic thought in which something like our category of adjectival feeling has figured importantly is the psychology of the nineteenth and early twentieth century. In the work of such men as Wilhelm Wundt and E. B. Titchener, feelings constituted one of the main kinds of elements out of which conscious life is made up. Of course, these thinkers were not directly concerned with the analysis of ordinary feeling terms, but their views on the nature of feeling were partly shaped by, and in turn partly shaped, the proclivities of reflective persons in dealing with feeling terms. When feelings first became prominent in psychology at the end of the eighteenth century, the most natural move for those in the Locke-Hume tradition of "elementaristic" psychology (to whom this all too rapid survey is restricted) was to extend to feelings the treatment that had been accorded to sensations—that is, to suppose that there are a number of unanalyzable phenomenal qualities (qualities the concept of which one can form only from attending to examples in one's own consciousness) that are distinctive of feelings. Thus a feeling would consist of the exemplification of one of these qualities, or some composite of these qualities, in consciousness. We shall call this the "Special Quality View." The most elaborate version is that of Wundt, who distinguished three fundamental qualitative dimensions of feeling—pleasantness-unpleasantness, excitement-depression, and tension-relaxation. He considered all the specific feelings to be distinguished by different patterns of location on these dimensions. Other theorists, such as Carl Stumpf, adopted a more parsimonious view, according to which feelings, instead of being conscious entities of a unique kind, are just complexes of bodily sensations (the "Bodily Sensation View"). Now if we convert these views into proposals for the analysis of adjectival feeling

concepts, they will both be contradictory of the Dependency Thesis. In either case, the concept of a given feeling will be independent of the concept of the state by reference to which it is standardly designated. On the Special Quality View, the concept of feeling ashamed *is* the concept of the exemplification in consciousness of one or more unanalyzable phenomenal qualities; and on the Bodily Sensation View the concept of feeling ashamed *is* the concept of a certain complex of bodily sensations. In neither case does the full concept of the state of being ashamed (which involves "cognitive" elements, like the supposition that one has done something to some degree shameful, and motivational elements, like a tendency to hide oneself from the notice of others) enter into the feeling concept. Let us use the term "Autonomy Thesis" for the general position that the concept of feeling *f* is logically independent of the concept of *F*. The two views just distinguished are the most prominent forms of the Autonomy Thesis. According to the Autonomy Thesis it is only because of the exigencies of interpersonal communication that we name feelings after complex states like anger, fatigue, and distress; this way of designating them is not truly indicative of the character of the feeling concepts. The Autonomy Thesis is deeply involved in many aspects of traditional thinking about the mind, including the pervasive view that the mind can be divided up into three components—cognition, conation, and feeling (or "affection")—each of which is specifiable apart from the others, and none of which is "contained" in any of the others, whatever causal connections there may be between them.

In this essay I shall endeavor to construct a pattern of analysis for adjectival feeling terms along the lines of the Dependency Thesis. But before embarking on this it may be well, in view of the powerful attraction exercised by the Autonomy Thesis, to indicate briefly why I regard both forms as unacceptable, though a thorough discussion is not possible here.

The Bodily Sensation View suffers from being straightforwardly incorrect, at least as a pattern of analysis for feeling concepts. Even if it should be true that a given feeling does in some sense consist solely of bodily sensations, it cannot be the case that our concept of feeling *f is* the concept of a certain complex of bodily

sensations. This is clear from the fact that in most cases we find it impossible to specify in terms of bodily sensations *what* feeling we are talking about. Just what bodily sensations make up feeling elated, relieved, annoyed, cheerful, or frivolous? In some cases we can make more headway. Feeling angry does typically involve sensations of the tightening of certain muscles and feeling sleepy a sensation of heaviness in the eyelids. But even in these cases it seems impossible to devise a description in bodily sensation terms that even looks synonymous with "feels angry" or "feels sleepy." Moreover, in so far as we can analyze a given feeling into components, it does not seem to be the case that the components will always be bodily sensations. A felt impulse to attack would seem to have as much to do with feeling angry as sensations of muscle tensings, and having frequent longing thoughts of home would seem to be as intimately constitutive of feeling homesick as any visceral sensations.

As for the Special Quality View, some forms of it are open to analogous objections, others not.[8] First, consider a position like Wundt's, according to which there are a limited number of basic feeling qualities or "dimensions," from which specific feelings are compounded. Here again we can say that, whatever may be synthetically true of feeling indignant, it cannot be claimed that our concept of feeling indignant *is* the concept of ————, where the blank is filled in by certain degrees of pleasantness, strain, and excitement. Such a specification would not permit us to distinguish, for example, feeling indignant from feeling annoyed, which, at comparable levels of intensity, does

[8] To be sure, one could hold that there are phenomenal qualities of feeling without asserting the Autonomy Thesis. This would just amount to a recognition of the fact that feeling depressed differs intrinsically from ("feels different from") feeling excited, plus an acquiescence in the natural tendency to mark this difference by saying that there is a different quality of feeling in the two cases. But this innocuous way of countenancing feeling qualities is compatible with any otherwise plausible position on the nature of feelings, including forms of the Dependency Thesis. For any view as to what it is to feel *f* could be construed as a statement of *what* quality it is that is distinctive of feeling *f*. If the identification of feeling with the exemplification of phenomenal qualities is to have any distinctive force, it will have to take (at least the basic) feeling qualities to be unanalyzable. It is this view we are combating.

not seem to be differentiated from feeling indignant by occupying different places on these continua. If, on the other hand, we consider a more radical position according to which every distinguishable feeling is constituted by the exemplification in consciousness of a unique quality, we cannot argue in this way, so long as the only way we have of specifying each of these qualities is by saying that it is the one that constitutes feeling *f*. We could, of course, complain that this "analysis" is singularly uninformative as to what distinguishes one feeling from another. But to this it may be replied that since this is the sort of thing a feeling is, we should not expect any more from an analysis.

There are, however, more fundamental objections that apply to all forms of the Special Quality View. The objections most prominent in current philosophical literature are those that stem from Wittgenstein's "private language argument,"[9] according to which it is impossible that one should meaningfully use terms for (have genuine concepts of) unanalyzable phenomenal qualities. But since, for reasons that have to some extent been made explicit in the critical literature,[10] I cannot accept this argument, I shall make no use of it here. Instead I shall marshal some more modest considerations that are directed specifically to the problem of feeling.

(1) In the case of feelings the postulation of unanalyzable phenomenal qualities does not have even the initial plausibility it has in the case of, for example, visual sensations. With the latter we can distinguish various qualities or qualitative dimensions, in respect to which two visual sensations might be more or less similar—shape, hue, brightness, and saturation. And even if these qualities are not exemplified separately in our sensory experience, we can concentrate attention on one or another of them selectively and can notice similarities between two sensations in hue but not shape or brightness, or in shape but not hue or brightness. But it does not seem generally possible

[9] L. Wittgenstein, *Philosophical Investigations*, trans. by G. E. M. Anscombe (New York, 1953), pt. I, secs. 256-272.

[10] See, e.g., Judith Jarvis Thomson, "Private Language," *American Philosophical Quarterly*, I (1964).

to do this for feelings. We can make a start at it here and there. We can say, for example, that feeling angry is like feeling frightened, and unlike feeling relieved, in involving a quality of tenseness. But more often the best we can do is just to say that different feelings feel different; talk of special feeling qualities is thus quite unrealistic phenomenologically.

(2) The unavailability of an adequate store of terms for unanalyzable feeling qualities is reflected in the fact that the devices we in fact employ for characterizing feelings and for distinguishing between one feeling and another are quite different from what the Special Quality View would lead us to expect. What is it like to feel homesick? Well, nothing seems very enjoyable, one often has a sinking sensation in one's stomach, and one often thinks of home with a pang of regret. What is the difference between feeling embarrassed and feeling angry? Well, when one feels embarrassed, one feels a tendency to hide one's head, one feels one's face blushing, one wishes one had never done what one is embarrassed about; while when one feels angry, one feels fists clenching, jaw muscles tightening, and one feels an impulse to hurt the person at whom one is angry. What we are doing in these cases is to distinguish feeling f_1 from feeling f_2 by reference to awarenesses of the components that go to make up the states F_1 and F_2. And it is not that this is just one way of doing the job. There seems to be no alternative. This fact suggests both that our concept of feeling f is *not* a concept of the exemplification of unanalyzable phenomenal qualities, and that it is based on the concept of F.

(3) To be sure, it is hard to draw a sharp line between explanations of a concept and claims as to what is true as a matter of fact, and it is even hard to be confident about any fuzzy lines. Is it analytically true that when one feels homesick one often thinks of home, or is this a "psychological" fact? But if no analysis of a certain concept is clearly dictated by the linguistic data so that we have a choice, we are well advised to opt for that pattern of analysis that will enable us to do more with the concept, that will render it more useful for employment in the description and understanding of the subject matter in question. The Special Quality construal of feeling concepts lies on the low end of the

continuum of usefulness. On this view we can say nothing as to what distinguishes one feeling from another and as to what makes feelings as a class different from other broad classes of conscious states. We get no hints as to the conditions under which one is or is not likely to have a certain feeling, as to how one can tell what another person is feeling, and as to how feeling terms are learned. Even if the view does not make the understanding of these things strictly impossible, as Wittgenstein claims, it makes all of them as difficult as possible. Thus in so far as a choice between analyses is underdetermined by the data, systemic considerations would favor the one that brings the richness of the concept of *F* into the concept of feeling *f*.

II

Spurred on by these considerations, I shall now address myself to the task of exhibiting feeling *f* as conceptually derivative from being *f*. To get started I shall exploit four clues.

Where we can make a clear distinction between being *f* and feeling *f*, one of the salient respects in which they differ is that the feeling is more private, less accessible to outside observers, and correspondingly more completely accessible to its possessor.

(1) A person is a final authority with respect to his feelings; his feeling reports are incorrigible.

That is, for feeling *f*, though not for being *f*, the authority of the subject is at a maximum. No indications available to others can outweigh the sincere asseveration of the subject that he does (not) feel *f*. For the *f*'s that figure, in feeling terms there are various publicly observable indications, though this differs markedly for different *f*'s. There are various ways in which others can tell, or at least compile reasons for believing, that I *am* tired, relaxed, relieved, disturbed, depressed, excited, frivolous, or adventurous, apart from reliance on my report. These indications may become so strong as to outweigh my sincere assertion. "You just think you're not tired." "Maybe you don't realize it, but you are quite disturbed about something." Normally, of course, when a person is (not) in the state *F*, he knows that he is (not),

just because normally when he is (not) *f*, he feels (does not feel) *f*. But in special cases public evidence can override the subject's own judgment that he is (not) *f*. With feeling *f*, by contrast, this possibility does not exist. If a person says at a given moment that he *feels* (*does not feel*) sleepy, disturbed, relieved, at home, embarrassed, cheerful, or frivolous, nothing anyone else can observe will show that he does not (does). Even if public evidence shows that he is not (is) actually sleepy, disturbed, or relieved, and even if the subject accepts that evidence, he can still insist that at least he *feels* (does not feel) sleepy, disturbed, and so forth, and the rest of us are without resources to do anything toward shaking that claim. The concept of feeling *f*, whatever its other characteristics, plays the role of an inviolable final refuge. One is safe from external disconfirmation so long as he restricts himself to saying what he feels. In this respect, feeling *f* seems to be related to being *f*, as visual sensations are related to the external facts of which one claims to have perceptual knowledge. If I say, on the basis of visual perception, that a car is coming (or say that I see a car coming), my assertion is subject to disconfirmation on the basis of evidence available to others; but if I restrict myself to saying that I am having visual sensations of certain sorts, or that my "visual field" contains such and such contents, then no public evidence can outweigh my sincere asseveration.[11]

Although I am unable to go adequately into the epistemological status of feeling-judgments, it may be helpful to forestall a couple of possible misunderstandings of the foregoing. First, I am not saying that it is impossible for a person to be mistaken in a feeling judgment, only that no one else can show it to be mistaken. It may be that this latter is tantamount to the impossibility of

[11] The claim that self-attributions of feelings are incorrigible is a controversial one, of course; and the more general claim that self-attributions of any states of consciousness are incorrigible (suggested on p. 15) is perhaps even more controversial. It is my position that there are no (even possible) cases in which it is clear that someone else has shown that a person is mistaken in attributing a given feeling or other state of consciousness to himself, as opposed, e.g., to having been careless in speaking, or being mistaken about what a given word means. I do not know how to argue for this position, however, except by considering a wide range of putative counterexamples, and there is no room for that here.

being mistaken, but we can leave that problem unresolved. Second, I am not, of course, saying that whenever a person utters a sentence of the form "I feel *f*" or a translation thereof in some other language, we are forced to accept the hypothesis that he does feel *f*. I have said only that we are forced to accept his sincere statement that he feels *f*, and there is many a slip between the former and the latter. An utterance of a sentence of the form "I feel *f*" may not be a sincere statement that the speaker feels *f* because the speaker does not know how to use this sentence or some of its components, because even though he has this mastery he has made a slip in this instance, or because he is lying. In taking his utterance to be a sincere judgment that he feels *f*, we are supposing none of these to be the case.

This epistemological difference between feeling *f* and being *f* is intimately connected with the fact that, unlike being *f*,

(2) Feeling *f* is a kind of state of consciousness (modification of consciousness, way of being conscious).

However we analyze the elusive concept of consciousness, there is no doubt that feelings are among the paradigms of conscious states. To *feel* sleepy, disturbed, or elated is to have one's consciousness modified in a certain distinctive way during the time one is so feeling. It is not (just) a matter of dispositions, of its being the case that certain things *would* happen if certain conditions *were* realized. Whereas being *f*, although the degree of involvement with consciousness varies widely for different *f*'s, is never just or primarily a matter of the momentary character of one's conscious experience; it always heavily involves physical states or processes, behavioral tendencies, and/or bits of knowledge or belief. Thus to *be* tired involves being in a certain describable physiological condition and having a tendency to rest; to *be* relieved involves realizing that something unfortunate that one had expected has not happened; to *be* apprehensive about *x* involves taking *x* to be potentially threatening in some way and having an impulse to avoid *x*.

These two differences are closely related. It is because being *f* straightforwardly involves such things as physical states and behavioral tendencies that public evidence can have a bearing,

and even a decisive bearing, on whether x is f. And the fact that feeling f is a mode of consciousness is closely tied to the incorrigibility of feeling reports. Different positions may be taken on the exact nature of the tie. Traditionally the notion of consciousness has been taken as needing no analysis, and the presence of something in consciousness has been taken as the paradigm of immediate incorrigible knowledge. I, on the other hand, would prefer to analyze the notion of a conscious state in terms of maximally privileged access. But whichever we take to be prior, the connection between being a state of consciousness and being an object of private, externally incorrigible knowledge is unmistakable.

It is important to realize that one can unstintingly accept the fact that feelings are conscious states without swallowing the Special Quality View of what makes a particular conscious state a particular feeling. The very notion of consciousness has been so intertwined with the notion of unanalyzable phenomenal qualities that it is difficult at this time of the day to speak of conscious states without seeming to commit oneself to the Special Quality View. One of the tangential aims of this paper is to demonstrate the possibility of taking seriously the notion of a conscious state without having to suppose that conscious states can be identified and distinguished only by the use of terms that are explainable only through private ostension.

So far we have the suggestion that whereas the "f" in "feel f" denotes a complex state of the person that includes a variety of more or less externally accessible components, the force of adding "feel" is to make the denotatum of the phrase into a state of consciousness to which its possessor has maximally privileged access, and which may exist without the actual presence of the state denoted by "f" and vice versa. We still have to determine how the content of the state, F, contributes to making the feeling of F the particular state of consciousness it is. Our next clues will encourage us to think of feeling f as some kind of consciousness (awareness) *of* being f. The concept of feeling f thus becomes some kind of cognitive derivative of the concept of F, somewhat as the concept of a belief that p is a cognitive derivative of the concept of the fact that P.

(3) The cases in which it is plausible to say that a person is *f*, even though he does not feel *f*, are cases in which the person lacks the kind of knowledge he generally has of being *f*.

When a person is obviously tired, but reports not feeling tired, we suppose that he is so preoccupied with the task at hand that his state of fatigue has not "registered" in his consciousness in the normal way or to the normal extent. Again, when *x* is obviously angry at *y*, but reports not feeling angry at *y*, we think that he is (unconsciously) preventing himself from "taking note" of the internal indications of anger that normally serve to tip him off as to his state.

The most simple-minded way of following out this suggestion would be to identify feeling *f* with one's self-knowledge of being *f*. But there are many reasons why this will not do, the most obvious being the point made earlier that, at least for many *f*'s, it is possible to feel *f* without being *f*, as well as to be *f* without feeling *f*. Since "*x* knows that *p*" implies that *p* is the case, we are thereby prevented from supposing that any state which can exist in the absence of *x*'s being *f* *is* *x*'s knowledge that he is *f*. Full-blown knowledge is too strong for our purpose. But the following considerations suggest that feeling *f* may be identified with something that is on the way to being knowledge that one is *f*.

(4) There is a general presumption that when one feels *f*, one is *f*.

As a result, if I am to maintain the hypothesis that *x* is not *f*, in the face of his sincere insistence that he feels *f*, I must not only have strong external indications of his not being *f*, but I must also provide some explanation of the fact that his feelings do not reflect his actual state in the usual way. This suggests that we might construe feeling *f* as something like a generally reliable, though not infallible, basis for a belief that one is *f*.

This suggestion is reinforced by considering the conditions under which it is most natural to say "I feel *f*," rather than "I am *f*," in reporting my condition. In most cases where I am

telling someone how it is with me, it is much more natural to say, "I am disturbed," "I am sleepy," "I am tired," than it is to say, "I feel disturbed," and so forth. The cases in which the feel-locution is specially called for are those in which there is some doubt as to whether I am really disturbed. (There may be a doubt "in my own mind," a doubt may have been raised by others, or I may suspect that a doubt may be raised by others.) If I think there is some question about the matter, I may say, "Well, I *feel* tired," which is very close to "So far as I can tell from my present state of consciousness I would suppose that I am tired."

In further developing this conception of feeling f as a basis for one's knowledge that one is f, we can make use of a widely accepted schema for the analysis of the concept of knowledge. In my version this schema runs as follows:

> x knows that p $=_{df}$ the conjunction of the following conditions:
>
> 1. x believes that p.
> 2. x is in a position such that one in that position is warranted in holding a belief (making a judgment) that p.[12]
> 3. It is the case that p.[13]

The simplest way of expressing our suggestion in terms of this schema would be to say that to feel f is to satisfy condition (2) for knowledge that one is f. But this would be too simple. The warrant that feeling f confers on the judgment that one is f is one that is subject to being overridden by other features of one's

[12] This is my substitute for the more usual formulation in terms of evidence or reasons. ("x has sufficient evidence for p," or "x has adequate reasons for believing that p.") This more common formulation is too narrow in that it does not allow for immediate knowledge—i.e., knowledge which is not based on other knowledge. It is this kind of knowledge with which we shall be especially concerned.

[13] Recent criticism has shown this schema not to be generally adequate as it stands. See, e.g., E. L. Gettier, "Is Justified True Belief Knowledge?" *Analysis*, 23, 121-123, however, I am using this schema, only to suggest a certain formulation for my analysis of feeling. So long as something of this general sort works, we will not be led astray.

situation. Thus if after three sets of tennis one feels energetic, but realizes on the basis of past experience that after such exertions his energy level is actually rather low, his total situation is not such as to warrant the judgment "I am energetic." What we need is the notion of a prima facie warrant, one which, in the absence of stronger conflicting considerations, will warrant a judgment that one is *f*.[14] When henceforth we speak of feelings as "warrants," this is to be understood as "prima facie warrants" even where the qualification is not made explicit.

It should be noted that warrants differ in strength, and that for some *f*'s, feeling *f* may provide a stronger warrant than for others. Feeling relaxed may more closely "mirror" being relaxed than feeling benevolent does being benevolent. Moreover, as was pointed out in an earlier footnote, there may be some *f*'s such that we cannot really envisage a person's feeling *f* without his being *f* to some degree. This may be true generally of emotional *F*'s, like anger, distress, fear, irritation, or depression. In this case, feeling *f* would provide the highest degree of warrant, one which could not be overborne by conflicting considerations. Our conception is designed to be elastic enough to embrace all these variations. No restriction is put on the strength of warrant involved. As for emotional feelings, if the above suggestion is correct, then the qualifier "prima facie" is not needed, but our general conception will still apply if we read "prima facie" as "at least prima facie," an unqualifiedly sufficient warrant being construed as a limiting case of a prima facie warrant.

Should any other components of one's knowledge that he is *f* be included in the analysis of feeling *f*? We have already given sufficient reason for exluding condition (3); we want to recognize

[14] We could also build in the presupposition that it will be the exception rather than the rule that stronger conflicting considerations will be present. That is, we can make it part of the notion of a prima facie warrant, *W*, for a judgment that *p*, that in general when *W* is present, the person's total epistemically relevant situation will be such that he would be warranted in asserting that *p*. I believe that we would be justified in asserting this of feelings—asserting, i.e., that generally (though not invariably) when one feels *f* his total position is such that he would be warranted in asserting that he is *f*. This additional element is not essential for the analysis, however, and I shall not add it.

that one can feel f without being f. As for (1) we do not want to include it as stated; as the tennis example brings out, one can feel f but realize, on the basis of other considerations, that he is not f. Examples like this are the exception, however; normally when a person feels f he believes that he is f; the feeling carries with it a strong and usually unresisted (though not irresistible) *tendency* to belief.[15] Moreover, it is not clear that this is merely a contingent fact. It seems plausible to suppose that our concept of feeling f would be different if it were the rule that when x feels f he does not believe that he is f. Hence, although I do not consider the indications to point unambiguously in this direction, I am inclined to add to the analysis a weakened version of condition (1)—namely, that for a state to be a case of feeling it must incline the person to believe that he is f, where it is understood that this inclination is again prima facie, one which may be overborne in particular cases by contrary factors.

As so far developed our analysis can be formulated as follows:

> x feels $f =_{df} x$ is in a conscious state, S_f, such that by virtue of being in S_f,
>
> (*a*) x has a prima facie tendency to believe that he is f;
>
> (*b*) x has a prima facie warrant for this belief.

III

The most obvious deficiency of this formulation stems from the fact that if one has external evidence for being f, of the sort another person could have, he could be in S_f, as so far defined, and yet not feel f. Suppose that I have noted various features

[15] In saying that someone has a tendency to do or undergo an M, we are saying that he will do or undergo an M unless strong enough interfering factors are present. What counts as an interfering factor differs for different M's. Where the M is belief acquisition, possible interferences will include having strong enough contrary evidence, and being strongly enough motivated not to have beliefs like this. It is clear that we can apply tendency notions (as opposed to having an abstract understanding of them) only in areas where we are in possession of a (possibly more or less crude) "nomological network" that embodies principles as to what leads to and what is incompatible with what.

of my behavior and thoughts that constitute good reason for supposing that I am angry at *y*. I have noted that I am glaring at him, that I respond coldly to his remarks, that I keep clenching my fists, and that my jaw muscles are tightening. Then by virtue of having noted these things, I would be in a state that gives rise to a tendency to believe that I am angry at *y*, and this state is such (the information contained in the state is such) as to confer a high degree of warrant on that belief. And yet, *ex hypothesi*, I do not feel angry at *y*. This example may strike one as unrealistic; it might be said that if I have noted all these things I *do* feel angry; sufficient conditions for my feeling angry have been satisfied. I do not believe that this is the case; it is conceivable (and clinically documented) that after all these things have been brought to a person's attention and he has not denied their reality one by one, he might still sincerely assert that he does not feel angry. And if what I said above about first-person epistemic authority in the case of feelings is correct, we would, in that case, have to admit that *x* did not feel angry. It is possible, however, to construct an indefinite number of quite unproblematic cases where a person has adequate external evidence for being *f* but obviously does not feel *f*. For one, just modify the above case so that *x*'s reason for supposing that he is angry at *y* is that his psychoanalyst tells him that he is; and let us further stipulate that the analyst is a reliable authority in such matters and is known to *x* to be such. Or consider the case in which *x* can tell that he is tired by noting the decrement in his performance and watching the dials on some instruments attached to his muscles, even though, owing to his absorption in the task at hand, he does not feel tired.

In order to deal with these counterexamples we must make some further restrictions on the kind of warrant the feeling confers on the belief that one is *f* or on the way it does so. The counterexamples were all cases in which what put *x* into a position to make a reliable judgment that he is *f* was some other knowledge (or warranted beliefs) that he had. The position in each case consisted of a set of facts of the form "*x* knows that *p*," where the *p* in each case was not "*x* is *f*" but something the truth of which counts in favor of the truth of "*x* is *f*." If we could maintain

that *x*'s *feeling f* performs its epistemic function not by containing knowledge of facts that serve as evidence for *x*'s being *f*, then we could in this way distinguish it from states of the kind mentioned in our counterexamples.

It seems to me that we can. I take it that when a person feels disturbed, he is thereby in a position (*ceteris paribus*), to make a reliable judgment that he is disturbed, regardless of whether he knows anything that could serve as evidence for the judgment. It is a fundamental fact about human beings that when they believe themselves to be *f*, while feeling *f*, their belief is prima facie very likely to be correct. If I know that *x* feels disturbed, then, in the absence of stronger conflicting considerations, that fact alone will warrant my accepting *x*'s claim to know that he is disturbed, quite apart from any question as to whether *x* has any evidence to support his claim. It is like my knowing that *x* is having visual sensations of a certain sort as a result of the stimulation of his visual sense organs; that in itself warrants me in crediting *x*'s claim to know that there is a tree in front of him, quite apart from any question as to whether *x* has any evidence to support his claim. It is just a fact about human beings that they are sometimes able to make reliable judgments about certain kinds of things without having any evidence on which to base those judgments. When the kind of thing in question is a state of the individual that is not itself a conscious state, we call the condition that enables him to do this a "feeling." A feeling of *F* is what normally enables one to have *immediate* knowledge that one is *f*, where by "immediate knowledge" is meant knowledge that is not based on other knowledge.

To be sure, it can happen that when I feel *f* and know that I am *f* I do have other knowledge that can serve as a basis for my belief that I am *f*, and with adult human beings this may be the usual case. This other knowledge can be of two sorts. First there is the knowledge that I feel *f*. If, when I claim to be tired, I know that I feel tired, then it would be quite natural and proper for me to cite the fact that I *feel* tired in support of my original assertion, if that were challenged. Second, there is the knowledge that various components of *F* are present. Let us think of the state of being angry at *y* as consisting of such

components as seeing *y* as frustrating one in some way, having an impulse to attack *y*, and such involuntary physiological reactions as the tightening of jaw muscles and the rushing of blood to the face. Now it seems that often when one feels angry at *y* and knows that he is angry at *y*, he also knows that he has an impulse to attack *y*, that the blood is rushing to his face, and so forth. And he might naturally and properly cite facts of these sorts in support of his assertion that he is angry at *y* if that assertion were challenged. How then can we claim that feeling *f* is what puts one in a position to have *immediate* knowledge that one is *f*?

First appearances to the contrary, this is no real objection to our thesis. We are not maintaining that when *x* feels *f*, his knowledge that he is *f* (when he has such knowledge) will always or even usually be only immediate knowledge. We are only saying that feeling *f* is *enough* (in the absence of stronger conflicting considerations) to satisfy condition (2) for knowledge that one is *f*, apart from any evidence the person may have. But this is quite compatible with the person *also* having other ways of satisfying condition (2). Overdetermination is as important in epistemology as in psychology.

In fact, our thesis is logically compatible with *x*'s always having evidence for his belief that he is *f* when he feels *f*. To be sure, if that were the case our thesis would be most implausible; the most reasonable position would be that the feeling performs its epistemic function by providing evidence for the belief that one is *f*. But it is not the case that evidence for being *f* always accompanies feeling *f*. As for the first kind of supporting knowledge, knowledge that one feels *f*, there are stages of development in which this support is not available. It seems reasonable to suppose that the concept of feeling *f* is more sophisticated than the concept of *F*, that one does not acquire the former until he has learned that sometimes when he is inclined to say that he is *f* (apart from evidence) he is not in fact *f*. Having learned this, having learned to be critical about his judgments, he can then learn to make a more noncommittal claim by saying "I feel *f*," rather than "I am *f*," and when that has come to pass he will have acquired the concept of adjectival feeling. But there will be a period in which a child has the concept of, for example,

being tired, but has not yet acquired the concept of feeling tired. And during that period there will be occasions on which he will know that he is tired because he feels tired, but where this knowledge will not be based on his knowledge that he feels tired, this latter knowledge being unavailable to him by reason of the absence of the concept.

As for the second kind of evidence, that concerning the components of *F*, it takes a relatively high order of analytic skill to dissect a state like fatigue, relief, or depression into its components; the more primitive situation is that in which the person, through his linguistic training, has learned to describe himself as tired, depressed, and so forth, when in a certain state, without being able to specify the features of that state by which he recognized it as a state of being tired, depressed, and so forth. Moreover, different states, as well as different persons, differ in the extent to which a dissection of the state into its components is a real possibility. I can go a considerable distance in saying what sensations, beliefs, and/or impulses are typically present when I am angry, anxious, sleepy, or relieved. But I am hard put to say what makes up being disturbed, even after a lot of hard thinking about the analysis of such concepts. If someone asked me how I knew I was disturbed, an answer specifying components of the state would not be readily forthcoming.

These considerations bring out the fact that feeling *f* is the sort of thing that puts one in a position to make a reliable judgment that one is *f*, even in the absence of other knowledge that will count as evidence.

In modifying our analysis along these lines, we have a choice between either specifying that S_f contains no knowledge of any sort, or leaving that question open and making the weaker stipulation that if it does contain any knowledge, that knowledge has nothing to do with its epistemic function. The first clause of the analysans would be rewritten for each of these alternatives as follows:

(1) x is in a conscious noncognitive (one that contains no knowledge) state, S_f, such that by virtue of being in S_f:
(2) x is in a conscious state, S_f, such that by virtue of being

in S_f, but not by virtue of any knowledge contained in S_f:

A desire to stick our necks out as little as possible in order to get the job done would lead us to opt for (2), but I find myself inclined to adopt the stronger condition. Feeling has traditionally been regarded as something noncognitive, as existing at a level of structure too simple, too undifferentiated to allow for belief, judgment, or even conception.[16] This conviction is embodied in the common practice of regarding feeling as one of the fundamental divisions of the mind, coordinate with cognition.[17] Furthermore, I believe that my inclination is backed by something more solid than an appeal to tradition. If we were to think of any knowledge as included within, for example, feeling angry at y, it would be knowledge of the various things that go to make up the state of anger—muscle tightenings, an impulse to attack, and so forth—or perhaps knowledge of feeling angry at y. But these are all pieces of knowledge that are highly usable as evidence for the claim that one *is* angry at y. Hence if feeling angry at y were to contain knowledge, that knowledge could hardly fail to enter into the capacity of the feeling to satisfy condition (2) for knowledge that one is angry at y. In other words, there seems to be no "neutral" knowledge that it would be plausible to suppose is contained in feeling f but that would have nothing to do with the epistemic function of the feeling. Hence I feel justified in embracing alternative (1).

As the analysis is now formulated it is a necessary condition of x's feeling f that he have a prima facie tendency to believe that he is f. But we commonly attribute feelings to creatures that do not have the concepts of the corresponding F's and hence

[16] In the Hegelian tradition feeling is taken to be noncognitive because it contains no "subject-object distinction." It is a seamless whole, though it contains within it seeds of knowledge that will flower when its implicit content is brought to light. I take it that our analysis brings out in a more intelligible form the insight that lies behind these dark sayings.

[17] It should be clear that we are not wholly going along with this practice. Although we are going to take a feeling not to *contain* any knowledge or belief, we are departing from the tradition in holding that a given feeling *concept* does contain cognitive *concepts*.

cannot believe that they are *f*. For example, we say of dogs and infants that they feel upset, depressed, angry, excited, and contented, but there is no reason to suppose that dogs and infants have a concept of being upset, angry, contented, and so forth.[18] And if they do not have the conceptual equipment required for believing that they are *f*, how can we attribute to them a tendency, prima facie or otherwise, to have such a belief? We can handle this complexity by the simple device of adding the possession of the relevant concepts as a condition either for the possession of the tendency or for its actualization. The former seems the more reasonable course, since a tendency to *r* on condition *C*, where *C* is not possible, could not have any psychological reality. Thus, even though my dog is not conceptually sophisticated enough to believe that he is depressed, still when I say of him that he feels depressed I am attributing to him the kind of condition that would issue in such a belief, or rather would tend to do so, if he did have the requisite concepts. With these latest modifications our analysis becomes:

> *x* feels *f* (where "*f*" is some adjective that denotes a state of the sort possessed by persons) $=_{df}$ *x* is in a conscious, noncognitive state, S_f, such that by virtue of being in S_f:
>
> (*a*) If *x* has the concept of *F*, he has a prima facie tendency to believe that he is *f*.
>
> (*b*) *x* has a prima facie warrant for this belief.

For abbreviatory purposes I shall henceforth sum up the analysans of the above as "*x* is immediately aware of being *f*." I choose the term "immediately aware" because it is the least

[18] This is, of course, a controversial point, partly because of an insufficiency of evidence, but more importantly because of a lack of clarity as to what it takes to have a given concept. If a dog could be trained to raise his left front paw whenever he is excited, would that show that he has acquired the concept of being excited? If these problems are resolved in such a way as to imply that it is impossible for a creature to feel *f* unless he has the concept of being *f*, then no modification of the analysis is called for. The ensuing discussion is designed to indicate how the analysis will have to be modified if the issue is resolved in such a way as to allow for the possibility of a creature feeling *f* without having the concept of being *f*.

unsuitable one available; the reader is hereby warned not to attach to it any meaning beyond what is spelled out in the foregoing. It may be that as ordinarily employed "*x* is immediately aware that *p*" entails "*x* knows that *p*," and I certainly do not want that. Moreover, the term "immediate awareness" is encrusted with a complex of philosophical views from which I am anxious to dissociate myself—for example, the view that there are certain objects of knowledge that are "self-certifying," that are "given" in some absolute sense, and so forth. None of this is involved in the term "immediately aware" as used here.

We will not be able to apply this schema mechanically to every adjectival feeling term; our language is not that simple. It will often be necessary to make adjustments in individual cases. For example, with respect to what I earlier called "feelings of behavioral tendencies"—feeling generous, adventurous, and so forth—the "*f*" in the "feel *f*" is often used by itself to attribute a personality trait to the person rather than a temporary state, and so is not suited to specify that of which the feeling is an immediate awareness. To be generous is to be a certain kind of person, a person who typically acts generously when he has the opportunity to do so. To specify accurately that of which feeling generous is an immediate awareness we need some phrase like "being in a generous mood" or "being disposed to act generously now." Again, many *f*-terms take on figurative senses when they appear in feel-*f* constructions. To feel rooted to the spot is to have a warrant for believing that one is rooted to the spot only in a figurative sense, and the same is true for feeling light as a feather or feeling weighted down. It falls beyond the province of this paper to work out such detailed problems of application.

IV

Let us sum up the salient facts or convictions about feelings that are reflected in our analysis.

1. Feelings are states of consciousness concerning which their possessors are final authorities.

2. We standardly name feelings after states that are not feelings.
3. We distinguish feelings from each other in terms of the states after which they are named, including various components of these states.
4. Generally, but not invariably, the feeling and the state after which it is named are present together.
5. Feeling *f* is in some way a source of knowledge that one is *f*.
6. Feelings are not themselves cognitions, but they contain, so to say, seeds of cognition.

Traditionally, concentration on (1) has led to the Special Quality View, with the result that (*a*) the other facts were left unaccounted for, (*b*) the generic character of feelings was not specified, and (*c*) it remained mysterious how we can have, as we seem to have, an intersubjectively shared vocabulary for feelings. By taking our initial clue from (2), we have developed a conception that remedies these defects. By construing a feeling as an immediate awareness (in our special sense) of a state *F* that is not itself a state of consciousness, we take full account of the conceptual entanglement of *f* in *feeling f*, and we take full account of the quasi-cognitive nature of the relationship between feeling *f* and being *f* (brought out in statements 4-6). At the same time we can recognize the privileged access feature (1) that has been the mainstay of Special Quality Views. For just as it is understandable that the reduction in objective claims makes one unassailable in reports of visual sensations, in comparison with reports of actually seeing a tree, so a parallel reduction in objective claims makes one unassailable in reports of feeling *f*, in comparison with reports of actually being *f*. Moreover, we now have a way of bringing out what is distinctive of feelings, vis-à-vis other kinds of conscious states. Feelings are conscious states that are in the special relation we have dubbed "immediate awareness" to other states of the person that are *not* conscious states. This gives us a way of generally distinguishing feelings from, for example, thought and mental imagings. Finally, this account makes clear how it is that a feeling term can have a publicly

shared meaning; at least it shows that any problems that remain here are not peculiar to feeling concepts. For the schema provides a device for defining feeling terms by the use of nonconscious state terms like "believe," "disturbed," "angry," "sleepy," and "adventurous," plus epistemic and other "topic-neutral" terms like "tendency" and "reliable." If terms such as these can have shared meanings, our analysis would imply that feeling terms also can.[19]

V

I think it cannot be denied that our analysis nicely embodies the items on the above list. Nevertheless, it must be admitted that there are other, hitherto suppressed, principles that are widely regarded as equally fundamental facts about feeling. And it may be argued that these principles show our analysis to be defective in not taking full account of the privacy of feelings. The principles are as follows:

7. If one has never had the experience of feeling *f* he cannot have the (full-blooded) concept of feeling *f*.

If I have never felt sleepy, frightened, or indignant, then it would commonly be said that I do not know *what it is* to feel sleepy, frightened, or indignant. We say of someone, "He doesn't know what fear is." And one might say, "I never knew what it was like to feel sorry for someone until. . . ." This makes it sound as if actually having felt *f* is a necessary condition of having the concept.

8. If a state is not markedly similar phenomenally to (does not "feel like") the states *I* call "feeling *f*," then it is not a case of what *I* mean by "feeling *f*."

The strongest argument for this principle involves reflection on various kinds of (apparently) conceivable reversals of psycho-

[19] Of course the analysans also contains the generic term "conscious state." But if, as was suggested above, this term can be explicated in terms of maximally privileged access, we can see how it too can acquire an intersubjective meaning.

logical regularities that now (presumably) hold. What if I changed in such a way that when I am immediately aware of being disturbed, my state is markedly similar phenomenally to what I now call feeling tense (or refreshed or sleepy), but not to what I now call feeling disturbed? Under those conditions I would not call a case of immediate awareness of being disturbed "feeling disturbed," or if I did, the sense of the phrase would have changed. It would not be feeling disturbed, in the sense in which I use that phrase at present. Again, take the kind of interpersonal reversal that has given rise to so much concern over the other-minds problem. Suppose it were the case (and for all I know might it not be the case?) that when another person, Jones, is in a state that conforms to our analysans for feeling disturbed, his state is phenomenally very similar to states of mine that satisfy our analysans for feeling sleepy (tense, relieved). If that were the case (and I realized it was the case) I would not admit that Jones is feeling disturbed when he is immediately aware of being disturbed, at least not in the sense in which I predicate "feeling disturbed" of myself.

(8) directly suggests that I should add another clause to the analysis of "feel f" as used by me:

(c) S_f is markedly similar phenomenally to $x, y, \ldots$

where "x," "y," and so forth refer to samples or "paradigms" of feeling f in my experience. For if similarity to my private paradigms is a logically necessary condition of something's being a case of feeling f, as I use that term, it must be that such similarity is a part of my concept of feeling f. This addition would also enable us to take (7) into account. We could then say that the impossibility of having the concept of feeling f without having felt f stems from the fact that in that case I would have no private paradigms of feeling f and so would not be able to form the conception of similarity to such paradigms. Let us call the proposal to add (c) the Privatist Proposal.

It is clear that the original analysis, just because it is entirely in "public" terms, cannot accommodate (7) and (8), and in fact implies their negations. It implies the negation of (7) because every term in the original analysis of "x feels sleepy" could be

understood by one who had never felt sleepy; one could have learned that to be sleepy is to have a strong tendency to go to sleep, without ever having learned to talk about feeling at all; likewise one could have acquired the concepts of a conscious state, reliability, and so forth, in areas quite remote from feeling, and certainly without ever having felt sleepy. On the basis of all this one could then acquire the concept of feeling sleepy by being given a verbal definition in these terms. If the original analysis (minus [*c*]) is correct, actually feeling sleepy is no more necessary for having the concept of feeling sleepy than actually being President is necessary for having the concept of being President. As for (8), the argument given for that principle presupposed that a state could conform to the original analysis of "feels disturbed" without being phenomenally similar to my private paradigms of feeling disturbed.

Two points are to be noted about the Privatist Proposal. First, although principles like (7) and (8) have often been adduced in support of Special Quality Views, I have not formulated the proposed addition in those terms. This is because it seems to me that (7) and (8) do not really support talk about unanalyzable phenomenal qualities. Such talk does not add anything to the statement in terms of similarity to paradigms, except the misleading suggestion that something is added. No doubt, if an experience of mine is similar to another experience of mine it must be similar in some respect(s) or other, and the talk of phenomenal qualities reflects this ontologico-logical truth. But it also suggests that we have resources for specifying, apart from the notion of similarity to a paradigm, what these respects are; and, as far as feelings are concerned, as was pointed out above, this is simply not the case.

Secondly, I have been assuming that the most one could justifiably claim for (7) and (8) is that they show the necessity of *adding* (*c*) to our analysis. It has often been claimed, on the basis of (7) and (8), that (*c*) can do the whole job, that my concept of feeling *f* is just the concept of something that is phenomenally similar to certain private paradigms. Suffice it to say that all the considerations presented earlier in support of our analysis can be marshaled against *that* claim. If we try to analyze

"*x* feels *f*" in terms of (*c*) alone, we leave it completely mysterious what makes different feelings fall into a general category of feeling, why we identify feelings in terms of states that are not feelings, why it is odd to suppose that people never are *f* when they feel *f*, and so on.

If condition (*c*) is added to the analysis, it follows that no two people have exactly the same concept of feeling *f*. For the paradigms referred to are necessarily different for each person. I cannot exhibit my paradigms to you, so as to enable you to base this part of your concept on them, and you suffer from a corresponding disability. The only sense in which I can "exhibit" my paradigms to you is to tell you or show you the publicly specifiable conditions under which I have them. Thus I can say, "Now I'm having one," and leave it up to you to make an identification in terms of publicly observable circumstances. Or I can tell you that my paradigms are experiences that I typically have when. . . . But these ways of bringing in paradigms fail to carry out the spirit of condition (*c*). For they do not tell the other person how the paradigms feel, what they are like phenomenally. And the force of adding (*c*) to the analysis is that I mean something more by "feel *f*" than is specified by any identification of the state in terms of its external relations, accompaniments, causes, or potentialities.

Despite the implication of the unsharability of feeling concepts, the Privatist is not committed to holding that one person cannot tell whether his feeling concepts apply to the feelings of other persons (and when they do so). For he can, compatibly with his position, accept some form of the argument from analogy and, on that basis, claim to know that in a given instance someone else is feeling disturbed in his sense of that term. It would still remain true that his *concept* of feeling disturbed could not be possessed by anyone else; for *what* he would be claiming to know in this instance is, *inter alia*, that the other person's state of consciousness is markedly similar to *his* private paradigms.

To be sure, one may resist the Privatist Proposal. This resistance may take either of two forms. First one may, along the lines of Wittgenstein's "private language argument," reject the whole notion of similarity to paradigms in one's conscious experience

as meaningless or otherwise not in order. Since, the argument runs, it is impossible for anyone else to check up on my supposition that a given experience is or is not similar to a private paradigm, I am left without any resources for distinguishing between its seeming to me that x is similar to the paradigm, and x's being similar to the paradigm. But that means that I do not really have a concept of similarity to the paradigm, for it is essential to any concept W that we be able to distinguish between "is W" and "seems to be W but is not." If we accept this argument, we will feel relieved of any obligation to include anything in any concept other than conditions for the application of which there are public tests. Even if the line of argument, leading through (7) and (8) to (*c*), sounds plausible, we will resist it on the grounds that the conclusion makes no sense.

The second form of resistance does not depend on branding the notion of a private paradigm as meaningless. We may agree that one ordinarily supposes that the states to which he applies the term "feels disturbed" are phenomenally similar to certain paradigms in one's own experience, but deny that this supposition forms any part of the *concept* of feeling disturbed. On this view a proposition of the form "If A feels disturbed, A is in a conscious state that is phenomenally similar to $x, y, \ldots$" is, if true, a contingent rather than an analytic truth. If it were to turn out that things had altered so that the states of consciousness that put one in a position to make a reliable judgment that he is disturbed are much more phenomenally similar to my present paradigms of feeling sleepy than to my present paradigms of feeling disturbed, this view would imply that I would (should) unhesitatingly apply the term "feel disturbed" to these states in just the same sense in which I now use that term, although I would undoubtedly be very surprised that I now feel this way when I feel disturbed.

I shall make no decision on the Privatist Proposal in this paper. To do so would require a thorough treatment of such thorny problems as the cogency of the private-language argument and the argument from analogy, the status of the verifiability criterion of meaningfulness, and the conditions under which we are justified in incorporating firmly held beliefs about C's into the

analysis of the concept of *C*. I shall content myself with pointing out the following. If (*c*) is rejected, for whichever reason, our analysis stands as originally set forth. If (*c*) is added to the analysis, the original schema, minus (*c*), can still claim to constitute an adequate account of the public aspect of feeling concepts, and as such it retains all the virtues imputed to it on page 27. It will still make explicit the communicable framework of feeling concepts, which is then filled in differently by everyone who has private paradigms with which to do so. Too much recent philosophy of mind has been imbued with the implicit assumption that concepts of conscious states must be either public or private. The possibility of a judicious mixture should be taken more seriously.

VI

One important implication of our analysis, in contrast to the Autonomy Thesis, is that feelings are ill fitted to function as ultimate termini of analysis. It is generally recognized nowadays that a desire, attitude, or emotion is of a high degree of logical complexity, consisting of a (perhaps indefinite) variety of dispositions of disparate sorts. But it is often implicitly supposed that in so far as we can spell out the various dispositions involved in a desire or attitude we can exhibit the concept as a construct out of concepts which do not have this degree of logical complexity. Thus it is correctly pointed out that to have a desire to go to Europe is to be disposed to *do* certain things under certain circumstances (for example, to be disposed to ask about travel costs if one believes that there is any chance of going), and it is to be disposed to have certain feelings under certain circumstances (for example, to be disposed to feel elated if one suddenly learns that one will be able to go to Europe).[20] But in exhibiting these features of the concept of desire, is one reducing it (even in part) to a lower level of complexity? Are we

[20] See, e.g., R. B. Brandt and Jaegwon Kim, "Wants as Explanations of Actions," *Journal of Philosophy*, 60 (1963); and W. P. Alston, "Motives and Motivation," *Encyclopedia of Philosophy*, ed. by Paul Edwards (New York, 1967).

getting down to logically simpler components when we pass from desires to feelings and actions? It has been repeatedly pointed out recently that this is not the case with actions, and the considerations of this paper show that the same holds for feelings. The concept of a given feeling is based on the concept of that of which it is a feeling. Thus the concept of feeling f is *more* complex than the concept of F. And in general concepts of felt states are more complex than the concept of a desire, since they contain desires among their conditions. (Part of what it is to be angry at x is to want to attack x; part of what it is to be tired is to want to rest.) Hence, although we undoubtedly contribute to our understanding of the nature of desire by showing the ways in which to have a given desire is to be disposed to have certain feelings under certain circumstances, it would be a mistake to suppose that we are making that contribution via exhibiting simple (or even simpler) components of the concept. Having rejected, for feelings as well as for emotions, desires, and attitudes, the Humean tendency to construe mental states as simple unanalyzable modifications of consciousness, we must re-examine the relationships of all these concepts.

WILLIAM P. ALSTON

University of Michigan