

# 1 The Fundamental Intentional State

Two crucial aspects of the human mind are phenomenal experience and intentionality. Under the influence of behaviorism, functionalism, and early identity theories of the twentieth century, accounts of these two aspects have been advanced that are independent of each other; certainly the two have received separate treatments by these theories.<sup>1</sup> Moreover, insofar as consciousness has been considered at all, it has typically been associated with phenomenal experience. The separation of the two and the alignment of consciousness with just phenomenal experience are fundamental errors resulting in distorted accounts of both aspects, especially of intentionality. I argue that the major source of these errors is the reliance on a strictly third-person methodology.

Both phenomenal experience and intentionality are conscious phenomena, and a proper and unified understanding of them requires a first-person methodology. Although others also emphasize the importance of first-person methodologies, they typically apply it only to phenomenal states. Recently this has changed. A few now exploit the first-person perspective in an attempt to expose phenomenal or phenomenological features of intentional states.<sup>2</sup> To the extent that the latter efforts are designed to bring intentionality back under the umbrella of consciousness I am sympathetic to them; nevertheless, I am dubious of the means used to achieve this end.<sup>3</sup>

My theory differs from all such efforts. I deploy a first-person methodology to uncover a unique kind of non-phenomenal narrow content, which I call *minimal content*. This novel content plays a fundamental and crucial role in my analysis of intentionality. It makes clear that consciousness is as strongly implicated in intentional states as it is in phenomenal ones. Indeed, I will argue below that minimal content gives rise to what I shall

call *the fundamental intentional state*. This intentional state has a unique logical structure, distinct from all other intentional states, but all others are rooted in it. Minimal content plays a similar foundational role in phenomenal states too, as I will explain in chapters 3 and 6. Minimal content, therefore, is a fundamental and unifying concept for a theory of the mind. I will argue in chapter 8 that it is fundamental in the philosophy of language too.

### Some Preliminaries

Any theory of intentionality must confront at least two striking and apparently contradictory features:

- (I) There is an asymmetry between oneself and others with regard to access to the contents of thoughts.<sup>4</sup>
- (II) An individual may neither know nor be in a better position than someone else to ascertain what his own thought is about.

The resolution of this apparent conflict turns on the recognition that a correct analysis of intentional states involves not only two kinds of content but also two kinds of methodology. I will argue that the sense of 'content' that preserves the first feature requires a first-person analysis, for it is invisible from the third-person perspective. The sense of 'content' that saves the second feature requires a third-person analysis. Both are required. An exclusive use of either type of analysis must fall short of an adequate theory of intentionality.

Advocates of a strictly third-person analysis of content abound, but in order to be successful they must either provide for the first feature or show that it is a mere appearance. Including the first-person perspective in the study of the philosophy of mind runs counter to the views held by a formidable array of contemporary philosophers. For example, consider Daniel Dennett's unequivocal rejection of the first-person perspective: "I declare my starting point to be the objective, materialistic, third-person world of the physical sciences." (1987, p. 5) Dennett himself characterizes this starting point as "the orthodox choice today in the English-speaking philosophical world" (ibid.). And later in the same work he says: "I propose to see just what the mind looks like from the third-person, materialistic perspective of contemporary science. I bet we can see more and better if we start here, now, than if we try some other tack." (ibid., p. 7) In between these

quoted passages, Dennett disparagingly cites a notable exception to the orthodoxy: Thomas Nagel's insistence on the importance of the first-person perspective, or the subjective, for a proper understanding of a number of things about humans and the world. I think Nagel is right. However, in this chapter my employment of the first-person perspective will not consider the qualitative or phenomenal aspects of mental states often associated with Nagel's work. (I will turn to them in chapter 3, and to sensory phenomena in chapter 6.) Instead, I will argue that the first-person perspective plays an indispensable role in uncovering a non-qualitative kind of content, one that plays a crucial role in studies of intentionality. Keeping faith with the orthodoxy, rather than enabling us to "see more and better," runs a significant risk of derailing the whole endeavor, since the content I identify is invisible to the orthodox methodology. There is no need to choose between the exclusive use of one or the other methodology in studies of the mind. Excluding either one is a mistake.

In what follows, I ask my readers to realize that at times they must adopt the perspectives of the individual thinkers I consider, and to think of the situation described as though they themselves were in the situation. Where I am not explicit, I let the context of my discussion serve as clues as to whether the first-person or the third-person perspective is the appropriate perspective to adopt. This is necessary because if I am right about minimal content, it is invisible from the third-person perspective; thus, my interlocutor must at least temporarily assume the first-person perspective if she is to fairly criticize what I say. My arguments for awareness of minimal content cannot even be comprehended unless one assumes that perspective on the cases I develop. This is a harmless request, since temporary adoption of the first-person perspective should not in itself beg any questions against opposing views.

Some possible misunderstandings of what follows may be avoided if I unequivocally state at the outset that I am not concerned with the ontology of mental states, whether they are intentional or phenomenal mental states. My interest lies with the *analysis of statements* concerning such states. The analysis is constrained by various simple statements that we find difficult to reject upon considering our experiences, notably the apparently conflicting statements ((I) and (II) above) regarding features of intentional states. Of course, there is also the pivotal constraint that the analysis provides an account of the fact that intentional states are "about" something

or some state of affairs, that they are “directed.” Giving an account of this “aboutness” that can be reconciled with (I) and (II) is my central task. Much if not all of what else we can correctly say about intentionality will follow from this analysis. The analysis is neutral regarding ontological commitments. Others, if they wish, may attempt to draw ontological conclusions from my analysis, but they are not my conclusions, nor are they forced by what I say. I will return to why ontological claims should be avoided from time to time in this and subsequent chapters, and in chapter 9 I will offer a more systematic discussion concerning my general conclusion regarding the idleness or vacuity of ontological claims.

### **Privileged Access and Minimal Content**

The asymmetrical access we have to our own thoughts is sometimes referred to as *privileged access*. Many different conceptions of the latter have been advanced, but two general reasons why privileged access is held in disrepute by some are that exaggerated claims have been made on its behalf, and it has kept company with dualism. However, privileged access is not necessarily connected to dualism, and certain alleged features (such as complete transparency and incorrigibility) may be dropped while still preserving an important point to the special access we have to some of the content of some of our states. The special access at issue amounts to no more than one’s ability to non-inferentially know the content, in some sense, of at least some of one’s thoughts. It is in some such limited guise that privileged access remains a compelling doctrine.<sup>5</sup> Although the asymmetrical access amounts to no more than this, we shall see the consequences are great and have not been adequately recognized.<sup>6</sup>

Let no one worry that if this privileged access is allowed all the work necessary for understanding intentionality is (mysteriously) done. Frankly, the access that the first-person perspective provides does not explain anything; still, it is what exposes the content to which we have special access and which is invisible to a strictly third-person methodology. Privileged access and this content are central features that must be explained or shown to be mere appearance. They cannot be ignored.

As an illustration of the kind of access in question, consider an example introduced by John Heil (1988). Suppose you ask me to form an image of my grandmother. On informing you that I have done so, you inquire how

I know it is an image of her and not of someone else. For such an inquiry to make sense, it must be possible that I could be mistaken when I think the image I formed is an image of my grandmother. But insofar as I formed it expressly to be of my grandmother, such a possibility must be ruled out: *It is constitutive of the forming of the image that it is of her.*

There is room for error on the agent's part of a sort, one that is harmless to the point here. For example, the woman whom I have come to think of as my grandmother may really be an impostor. In that case, however, I would not be making a mistake about whom my image is of; rather, it is a mistake about my blood relationship to her. Another kind of case that may be put aside is one where some image randomly comes to my mind. Here, although the image is mine, I would be in no privileged position to ascertain of whom it is; indeed, it would appear to be no more *of anything* than are the "stars" I experience on receiving a blow to the head. If one were to maintain that the image was of someone in particular, the criteria for deciding would be at best unclear. In any case, I certainly would not be in any privileged position to know this in such a case.

Thus, although I can be wrong about images of mine in some ways, I cannot err in identifying whom my image is of when I deliberately form it to be of some particular individual. I cannot err in the latter simply because the possibility of error in these circumstances does not make any sense, not because I have some special mental powers or because I am cognizant of a special kind of entity. That an image is of the particular individual in question is a *constitutive* element of the very act of forming the image. It could not be *that act* if it were not *of that individual*. Given this, plus the fact that the content at issue is only part of the content of a thought, such "infallibility" is not to be confused with the Cartesian kind.

Someone who holds a "resemblance" criterion for what an image is an image of might think there is a possibility of error here. Suppose the image I formed bears a rather poor resemblance to my grandmother. The image itself may even be an excellent resemblance of someone other than my grandmother; nevertheless, the degree of visual faithfulness to my grandmother is irrelevant. It is, after all, the image I deliberately formed to be of her. The criterion determining who or what it is an image of cannot be based on what it is the most (visually) similar to. It is not as if I conjured up an image and then began to wonder who it depicts; the image was conjured precisely to be an image of my grandmother.

Resemblance could be relevant in a very different kind of case: Were I to find a photograph of someone, I might well wonder who is depicted in the photo and use resemblance as one criterion for deciding. In this case resemblance would be appropriate because of the different circumstances and causal relations involved in producing the photo. The relevant circumstances and causal relations are radically different in the case of a deliberately formed image, however; since here resemblance is irrelevant as a factor in the determination of the individual represented by the image. The formed image is a direct result of my act to produce not simply an image but an image of a certain individual. That being a constitutive element of the act, I cannot perform the act without the result being of that particular individual; otherwise, it would be a different act.

The same point could be made in regard to a sketch, which has the advantage of being publicly observable. If you ask me to sketch my grandmother, the result may indeed look more like your neighbor than my grandmother, and we may even agree on this. But a poorly drawn sketch of my grandmother is still a sketch of her. That is why it is said to be *poorly* drawn; it is not said to be a (well-drawn) sketch of whomever it most closely resembles. In parallel to the imagery, and in contrast to the photo case, there is neither a possibility of error on my part as to whom my deliberately drawn sketch is of nor any need for me to make any inferences to determine this. There is a possibility of error in your judging whom my sketch is of, and your judgment will be based on inference and may rely on resemblance, but my judgment does neither. That difference is just a manifestation of the asymmetry of access.

I examine another example<sup>7</sup> to illustrate further both the kind of privileged access that is at issue and the kind of content to which we have this special access. Suppose I make a diagram while lecturing on the battle of Borodino. I make Xs to mark the locations of Napoleon's troops and Os to mark the locations of Kutuzov's. Though there are countless errors I may make in my lecture and in the accuracy of my diagram, it makes no sense to ask me how I know that the Xs represent Napoleon's troops rather than Kutuzov's. Since the diagram is mine, the Xs cannot fail to represent what I intend them to represent. Suppose that, upon looking at my diagram, I have the thought that Napoleon had too heavy a concentration of troops in the northeast. On having this thought, I non-inferentially know that an X represents (a certain number of) Napoleon's troops. I know this straight

out, without recourse to a thought about this thought. It would make no sense for me to look for evidence of this, to puzzle whether an X really represented (a certain number of) Kutuzov's troops. It does make sense that you might be puzzled about such a matter; you may have to infer on the basis of evidence, which could be as simple as the fact that I told you, what my Xs represent. Although I may be wrong in thinking there were too many troops in the northeast (as a matter of military strategy), I cannot be wrong in thinking that one of my Xs represents (a certain number of) Napoleon's troops.

Now consider a new twist on the case just presented. I may wonder what Kutuzov would have done had he been in Napoleon's position, my speculations being based on my knowledge of his psychology. But now the Xs represent (a certain number of) Kutuzov's troops. When I consider this situation from my first-person perspective, it is evident that here too I know this content of my thought straight out. My having this non-inferential knowledge clearly does not require a second-order thought, as is commonly supposed. This is in marked contrast to another's acquiring knowledge of the shift in my thought content. Someone else would have to consider a second-order thought in order to ascertain what my thought is about, and would also have to make inferences from evidence. (Further support of these contentions is offered in chapter 2.)

Not only do the cases just presented illustrate a relevant asymmetry of access, a kind of privileged access; they also help to identify a restricted kind of content: What each of us has non-inferential knowledge of in these examples is basically the "subject of one's thought." Just what is meant by this locution is unclear at this stage, however. It is a central contention of mine that this locution requires two distinct decompositions. I will call one *minimal content* and the other *objective content*.

*Minimal content represents the subject<sup>s</sup> of the intentional state as the agent conceives it.*

*Objective content indicates the subject an objective observer of the agent would ascribe as the subject of the agent's intentional state.*

I will argue that both minimal content and objective content are required in the analysis of what a thought is about, and that each signals what the thought is about in a different way. Neglect of either is neglect of a critical element required for a complete analysis of intentionality. As the examples

suggest (and I will develop the point further below), the thinker plays a *constitutive* role in determining the minimal contents of her thoughts.

The privileged access in the above examples was to the minimal content, the subject of the agent's thought *as she conceived it*. It is this that was non-inferentially known in each of the above examples.<sup>9</sup> It is worth pointing out here, while the cases are "fresh," that one's special access to what is represented in such cases is known *straight out* by the thinker. One may non-inferentially know the subject of one's thought in the sense of minimal content and do so without having to have a further thought about the initial thought. I take it that the above cases demonstrate this—but they do this only if they are considered from the first-person perspective.

### Intentional States and Minimal Content

A peculiar feature of an intentional state is that it is directed at an object, the intentional object. The explanation of this directedness and what the intentional object consists in is difficult; different authors provide considerably different accounts. These differences have serious consequences. My theory builds upon John Searle's theory of intentionality, so I will explain the basics of his theory first.<sup>10</sup>

For Searle, the intentional object just is the actual object or state of affairs to which the intentional state is directed. For example, if someone loves Sally and also believes that it is currently sleeting outside, then the intentional objects for these two states are the flesh-and-blood woman Sally and the coming down of freezing rain, respectively. But these states also have representative contents, which represent, respectively, the woman and the relevant state of affairs. (Searle sometimes calls the representative content *intentional content*.) On the other hand, there are intentional states that have no intentional object, though they must still have a representative content. For example, Ponce de Leon, in searching for the fountain of youth, was in an intentional state that had a representative content (consisting in part of a representation of what he was seeking), but his intentional state did not have an intentional object. Thus, on Searle's view, every intentional state has a representative (intentional) content, and it is by virtue of this content that the intentional state is directed at an intentional object, but not all intentional states have intentional objects.

It is of paramount importance to keep the intentional object and the representative content distinct. In particular, in those cases that Searle would describe as failing to have an intentional object, one must resist the temptation of identifying the representational content with the intentional object. The temptation to do so is abetted by the consideration that the thought must be about something—it is not about nothing. Since there is no relevant actual object or state of affairs for it to be about, the reasoning continues, it must be about the representative content.<sup>11</sup> The idea behind such views is that, if one is to avoid talking or thinking about “nothing,” there always must be an intentional object of some sort.<sup>12</sup>

Sometimes the representative content is propositional in form, as when one believes that a certain state of affairs obtains; sometimes it is not, as when one desires a certain object. In either case, however, an object is “signaled” by the representative content. I endorse everything in Searle’s view of intentionality that I have presented to this point, but central to my theory is a specific addition. It is my contention that in the *analysis* of any intentional state there are *two* distinct but correct ways of characterizing the object signaled: minimal content and objective content.

Both minimal content and objective content are restricted to the subject of a thought; neither includes what may be thought of the subject, what may be attributed to it. Just as minimal content and objective content must be kept distinct from one another, neither is to be identified with the intentional object. The latter is an actual object or state of affairs, whereas minimal content and objective content are concepts employed in the analysis of intentional states. They correspond to first-person and third-person descriptions, respectively, of what is the subject of the agent’s intentional states.

Minimal content and objective content must not be identified for several reasons: (1) They play different roles in the analysis of intentional states. (2) They may signal different objects. (3) The agent’s access to these two contents is importantly different. Because of this difference in access, I spoke above, and I will continue to speak, of minimal content’s *representing* and of objective content’s *indicating* what the thought is about. Though these terms mark a difference, I do not claim to explain the difference here. When a generic term is required, I will speak, as I did above, of *signaling* what the thought is about.

The situation may be depicted as follows: The schema for an intentional state is  $\Psi(R)$ , where  $\Psi$  is some psychological mode (e.g. believing or desiring) and  $R$  is the representative (or intentional) content and is distinct from the intentional object, as defined by Searle.<sup>13</sup> If I am right, any such schema also requires a twofold decomposition:

(i)  $\Psi(\Phi(m))$

and

(ii)  $\Psi(\Phi(o))$ ,

where  $m$  is the minimal content,  $o$  is the objective content, and  $\Phi$  is what is attributed to “the subject of one’s thought.”<sup>14</sup>

That one of these contents is subjective and the other objective follows directly from the thinker’s different relations to them. This is crucial in explaining why (I) and (II) in the first section of this chapter are both true despite the appearance of conflict. The relevant contents in (I) and (II) are different, corresponding to minimal content and objective content, respectively.<sup>15</sup> Since the contents are different, any apparent conflict between (I) and (II) disappears.

Those who exclude a first-person methodology in the analysis of intentional states cannot countenance minimal content. Hence, the employment of a strictly third-person methodology encourages, if it does not imply, the extremely counterintuitive rejection of (I) for the sake of (II). Since the orthodox methodologies in contemporary analytic philosophy of mind employ a strictly third-person methodology, it is not surprising that minimal content has gone unnoticed and that (I) and (II) have appeared to be in conflict.

In maintaining the importance of a subjective first-person perspective in the analysis of intentionality, I am not denying the possibility of an objective account of it. Indeed, my introduction of the notion of minimal content and my analysis of any intentional state into two characterizations,  $\Psi(\Phi(m))$  and  $\Psi(\Phi(o))$ , are steps in this direction. (Further steps are advanced in chapter 3.)

### **Illustrative Applications of the Concepts**

To further clarify these different concepts of content and to illustrate the fruitfulness of deploying both, consider the following example. A mid-

eighteenth-century chemist might have thought that phlogiston was abundant in charcoal. The objective content of such a thought could not indicate phlogiston, since there is no such thing; nevertheless, the minimal content of the chemist's thought, the subject of her thought as she conceived it, did represent phlogiston. While the chemist had a special access to her minimal content, she certainly had no special access to the objective content or what it indicates. This absence of special access to the objective content and what it indicates is not simply because, as is the case in the present example, that there is no such thing. Even when the indicated object exists and is co-extensive with what the chemist's minimal content represents, there is still no privileged access to *it*, the existing intentional object. This last point is simply the familiar one that whatever privileged access there is, there is none that we have to actual objects in the world.<sup>16</sup>

Though the objective content and one's minimal content may signal different objects, they need not. Suppose I am correct in thinking that a certain woman is my grandmother, but I have many false beliefs about her, and there are a number of important things true of her of which I am ignorant. Suppose I entertain the thought that my grandmother was a gracious woman. The correctness of that judgment is a function of what she actually did or would do in certain circumstances. My judgment will be a function of what I believe, rightly or wrongly, about her. I may harbor so many false beliefs pertaining to my grandmother that another, someone who has a more accurate view of her, would say of me, in a colloquial vein, "He does not know her at all." Despite my lamentable epistemological status, any of my intentional states concerning my grandmother will have a minimal content that represents her (the right woman), and I can know this non-inferentially. Thus, both the minimal content and the objective content can signal the same object, in spite of all the misinformation I may have regarding the intentional object.

My grandmother is relatively easily to individuate, even if I labor under a vast amount of misinformation about her. Lest it be thought that that is why the minimal content and the objective content can signal the same subject, consider this: Suppose I entertain the thought that quarks are difficult to experimentally detect. Suppose also that the physicists who theorize about such things are right in thinking that there are quarks. Although I have barely the foggiest idea of what quarks are, it is still plausible that my minimal content and the objective content of my thought signal the same

object. For, at the very least, had the physicists not developed their theory of quarks, I would not have my (foggy) thought about quarks. Because I have so little knowledge of quarks, the subject of my thought, as I conceive it, is intended to be whatever the physicists are talking about. I do not add to it in any way that could make what is signaled deviate from the objective content from what those in the know take it to be. (Here, perhaps, a germ of truth in direct theories of reference is operative.)

The above is not to say that to be meaningful my thoughts require that the objects signaled by them exist (as with Russell's logically proper names), as the phlogiston example should make clear. Quite the contrary, it is the claim that minimal content represents something to me—whether what is represented exists or not, whether or not my representation of it is accurate or not, whether I have or do not have correct collateral information pertaining to it—and that it does so on pain of nonsense. Constitutive of my entertaining some particular thought is that its minimal content be what I conceive it to be. *It would not be the thought that it is if it did not have the subject that I conceive it to have.* For, at the very least, I must have *some conception* of what I am “thinking about”; without this, the very meaningfulness of the claim that *my thought*—as opposed to some sentence I might utter—is about anything is brought into question. It is for reasons such as these (more support will come later) that the thinker plays a constitutive role in determining her minimal content. It is also why I call this kind of content minimal.

When my minimal content does match the objective content, but I labor under a number of false beliefs regarding it, I clearly have limited understanding of the object represented, and my ability to explicate my thoughts pertaining to it is thereby limited. I explain below the contrast between *being aware of* and *understanding* what is represented by one's minimal content. For now, I simply point out that to have a thought about something (e.g. quarks) I need have no great understanding of them; indeed, they need not even exist. When my understanding is slight, I will, no doubt, be unable to explain my thought adequately. Still, to have such a thought at all, I must be able to somehow represent quarks in thought and, therefore, to individuate them to some degree.

### **Partial Recapitulation**

The analysis of a single intentional state requires both a minimal content and an objective content, because what the thinker takes her thought to be

about may not in fact exist, or may be seriously misconceived by her, or may be entirely different from what it is about objectively. The occurrence of any or all of these mishaps in no way vitiates the thinker's awareness of the subject of her thought as she conceives it, though they may mislead another thinker who is not adopting her point of view. Ideally, what is signaled by minimal content and objective content is one and the same, but this cannot be guaranteed. To allow for the possibility of either a match or a divergence in what is signaled by these two contents, a correct analysis of an intentional state requires both kinds of content; but one has privileged access to only one of them: one's own minimal contents.

### Further Development of Some of the Fundamental Ideas Presented

I now introduce a series of cases in order to clarify a number of central ideas:

1. minimal content
2. what it is for an individual to be aware of minimal content
3. the character of an individual's privileged access to her minimal contents
4. the difference between being aware and understanding one's own minimal contents.

The words 'awareness' and 'understanding' are widely used but vague. It is certainly questionable whether they have univocal uses in their wide applications. In this section I will contrast the two terms and show when each is preferred. (I do not claim that the contrast I have developed is or should be universally used, but it does mark, I think, an important difference.) *Awareness*, when taken as non-inferential knowledge, is preferred in discussions pertaining to the individuation of the subject of a thought, construed as minimal content. *Understanding*, on the other hand, comes to bear when the explication of a thought or the determination of its subject, construed as the objective content, is at issue. I will clarify this by reminding you of some familiar truths about formal theories.

Consider a formal theory that has as one of its models the natural numbers. It is well known that such a formal theory has other models—for example, a model that has sets, and not numbers, in its domain. So the theorems may be viewed as truths about numbers or as truths about sets. Now consider two individuals, A and B. Individual A learned the theory as axioms about natural numbers, and only them; she doesn't even know that

the theory has other models. Individual B learned the same formal theory as axioms exclusively about sets.

What seems quite uncontroversial is that each of these individuals can increase her respective understanding of the theory as she explores ever further the consequences of the axioms. The increase in understanding may be measured by the extent to which A and B can prove theorems and explain their proofs. Suppose, somewhat artificially, that A and B are identical in their symbol-manipulating abilities—that A can go on in fascinating detail with truths about numbers, and B can go on in equally fascinating detail about sets. The point, of course, is that their capacities to respond differentially and appropriately to the symbols are, by hypothesis, identical. There is absolutely nothing that would differentiate their responses vis à vis the manipulation and concatenation of the symbols of the theory. Yet A and B take themselves to be proving truths about different things; their respective minimal contents are different and represent different objects.

We, too, may acknowledge that what A and B take their respective contents to be are quite distinct, for the one it is numbers, for the other it is sets, and we may do so without our making any ontological commitment to the nature of either numbers or sets.<sup>17</sup> Indeed, we may acknowledge that A and B conceive themselves to be thinking about different objects even while *we* hold that they are ultimately not different (the numbers just are certain sets), or that neither numbers nor sets are actual objects at all (say, along nominalistic lines). In short whether our ontological commitments coincide with A's or B's or neither, it will have no bearing on our discussion of what *their* minimal contents are. Thus, while we acknowledge what are clearly the respective minimal contents of our theorists' thoughts, it might well be that the objective content of one or the other's thought indicates something other than what is represented by her minimal content.

The above is intended to bring out, in a preliminary way, the importance of keeping the discussion of issues surrounding symbolic manipulation, understanding, and ontology separate from the discussion of issues surrounding minimal content. This follows from the fact that A and B are identical with respect to their symbol manipulation and understanding, yet they have different minimal contents, and that all of this is consistent with several different ontologies. Furthermore, though A's minimal contents

represent numbers, as A becomes more sophisticated she may conclude that there are no numbers but only, say, sets. A's minimal content could on occasion still represent numbers, though she does not believe that there are any numbers, just as we may entertain thoughts of mermaids without believing that any mermaids exist.<sup>18</sup> Thus, symbol manipulation, understanding, minimal content, and ontology are independent in at least the ways indicated here. For the next few pages I will focus on how awareness of minimal content is largely independent of both understanding and symbol manipulation. I will also touch on the relation between the latter two. I will not address any ontological questions as such.

Recall that understanding involves the ability to determine a thought's objective content and to be able to explicate the thought in such a way that there can be inter-subjective agreement among qualified experts. It is undoubtedly the case that a significant factor in the measure of one's degree of understanding is the extent to which there are appropriate and differential responses to a relevant set of symbols. Such differential responses probably are extremely relevant to accounts of understanding. At the very least, they would be essential as a *measure* of one kind of understanding. However, whatever bearing differential symbol-manipulating capacities have, or do not have, vis à vis the question of understanding, they do not have much to do with the concept of minimal content or with the fact that an individual is sometimes aware of her minimal content.

When I hold that one's minimal content is *independent* of symbol manipulation, I am not holding that one can have and be aware of minimal content *without* being able to engage in some relevant symbol manipulation. The claim of independence results from my argument that the symbol manipulation, no matter how sophisticated, may remain constant while the minimal content (i.e. sets or numbers) varies. Qua minimal contents, A's and B's understandings differ, despite their identical symbol-manipulating abilities. A's minimal content represents numbers and B's represents sets, and this is so regardless of what the true ontology includes (numbers but no sets, sets but no numbers, both, or neither). This is not to say that either A or B would have any understanding if she were unable to perform some relevant symbol manipulations, only that symbol manipulation is not sufficient to determine understanding since, at least in this case, it is not sufficient to determine either the thinker's minimal content or the objective content of her thought.

Furthermore, as the consideration below of the theorist C will show, one may be aware of the same minimal content as A, even though one's symbol-manipulating ability is considerably inferior to A's. Although such inferior symbol-manipulating ability may well be a sign of a poorer understanding of numbers, it in no way diminishes the fact that the agent's thoughts are about numbers. Just how much symbol manipulation is required for awareness of minimal content is not clear. What is clear is that *different* minimal contents are consistent with identical high symbol-manipulation abilities (e.g. the theorists A and B) and that the *same* minimal content is consistent with radically different levels of symbol-manipulation abilities (e.g. the theorists A and C). Thus, rapid and smooth symbol manipulation is not necessary for one to be aware of one's minimal content, and very little is sometimes enough. (Compare my discussion of Van Gulick in the next chapter.) To this extent, and in this sense, awareness of minimal content is independent of symbol manipulation. Let me explain further.

Consider a third individual, C, who, like A, learned the formal theory strictly as a theory about numbers, but is not nearly as adept at proving theorems as is A. On the basis of the disparity in their symbol-manipulating skills, we would say that A has a better understanding of the theory of natural numbers than does C. Yet C, though stumbling in her symbolic manipulations, still knows without inference the minimal content that those symbols or her corresponding thoughts have for her. Her difficulty consists in providing proofs of particular truths about numbers. In spite of this, whether it is a statement (thought) that she can prove easily or one that she cannot prove at all, she is equally aware that it is a statement (thought) about numbers (whatever their ultimate ontological status). What is important here about minimal content is that C is directly aware of the same content in both cases and that in either case it is correct to say that C has the same minimal content as does A, in spite of A's greater facility at theorem proving. (The latter is indicative that A's understanding of number is greater than C's.) In contrast, B's minimal content differs from both A's and C's, despite the fact that B's symbol-manipulation ability is identical to A's.

The results drawn from our consideration of the three theorists depended on our projecting ourselves into their respective circumstances—that is, adopting their first-person perspective. When we do so, it is evident that, while each can be aware of the minimal content of her thoughts as she

manipulates symbols, the appropriate differential responses themselves have very little, if anything, to do with either what that minimal content is or the agent's awareness of it. What they conceive the objects to be is the same or different independent of those responses—independent in the sense explained above.

C's limitations in her theorem proving may well indicate limitations in her understanding of number (or, perhaps more accurate, limitations in her understanding of number theory), though this limitation affects neither what her minimal content represents nor her awareness of it.<sup>19</sup> Any proper account of understanding, undoubtedly, will be one that admits of degrees of understanding. Indeed, the varying degrees of understanding may well be reflected, if not partly constituted by, the varying degrees of successful symbol manipulation. In important contrast, *awareness of minimal content does not admit of degrees*. Either our theorist is aware that she is proving (or attempting to prove) theorems about numbers or she isn't. The relative ease with which she is able to construct such proofs appears irrelevant to her awareness of her minimal content. She may prove, say, "2 is a prime number" or "There is no greatest even number" with ease, but despite her best efforts she may be unable to prove, say, "Every number has a prime factorization." In *any* of these cases, however, C is *equally* aware that her minimal content represents number, be it a specific one, as in the first sentence, or number in general, as in the latter two.

To repeat, it is not the theorist's symbol-manipulating ability that determines which minimal content she has, nor does it determine her awareness of it. Clearly, then, no matter what the full account of understanding is, it is importantly different from one's awareness of what the minimal content of one's thought represents. Our understanding admits of degrees; our awareness of minimal content does not. Though we may well have more or less understanding of, say, numbers (as evidenced by how adequately and thoroughly we can explain or prove that numbers have certain properties), it is nonsense to say that one's awareness of numbers has increased or decreased (or is more or less than someone else's awareness of numbers). C, our weak number theorist, knows that her thoughts are about numbers. Period. It does not matter whether it is a thought that she can easily prove or explain, or not. Nor does it matter that *generally* she is not very good at explaining the properties of numbers, or at proving theorems about them. These abilities bear on the strength of her understanding and give rise to

talk of degrees of understanding. Such abilities and the resulting degrees are not transferable to awareness of minimal content.

I have exploited the familiar idea of multiple models for a single formal theory in my discussion of the theorists A–C. Additional light may be cast on the idea of one's being aware of one's minimal content by considering a converse relation: cases where the same object is characterized by different formal theories. Rota et al. (1989) argue that such situations are quite common in the practice of mathematics. Among other examples, they discuss the real number line and groups as cases where the practicing mathematician has a pre-axiomatic grasp, an understanding of these objects that is free from any particular axiomatization.<sup>20</sup> This grasping of the object is critical, Rota et al. claim, to identifying different axiomatizations as being axiomatizations of the same object.<sup>21</sup>

Rota et al. recognize that a student learning a theory will sometimes be unavoidably dependent on some particular axiomatization of that theory. This dependency is especially strong in a case such as group theory, much less so in other cases, e.g., the real number line. But, in any case, it is a dependency that is overcome once a student becomes familiar with the theory. Becoming familiar with the theory is becoming aware of, say, groups. In my terms, it is becoming aware of a minimal content that represents groups in a way that transcends the particular axiomatization by which the student was introduced to groups. Rota et al. argue that “to the mathematician, an axiom system is a new window through which the object, be it a group, a topological space or the real line, can be viewed from a new and different angle that will reveal heretofore unsuspected possibilities” (1989, p. 382). In holding this, Rota et al. are neither arguing nor claiming that these “grasped” objects have some special ontological status: In particular, they insist that they are not arguing for Platonism, and that they are just “acknowledging the actual practice of mathematics” (ibid.). Nor are they arguing that the mathematician is exercising any special faculty, mental or otherwise, when grasping these objects. They take a neutral stand with respect to both the ontological status of these objects and the character of the faculty by which one is aware of them. Their analysis is developed from the first-person perspective of a practicing mathematician, not from a third-person perspective. Thus, their approach is methodologically similar to my own, and it is appropriate to characterize the mathematician's grasping of the objects of his own thoughts as another

characterization of one's being aware of the minimal content of one's thought.

One of the distinct advantages of the concept of minimal content is precisely that it permits us to avoid the troublesome issues surrounding the concept of understanding while still allowing an advance in the analysis of intentionality. The scope of the word 'understanding' is clearly broader, if not also vaguer, than is one's non-inferential knowledge, awareness, of one's own minimal content.

### Understanding, Explicating, and Individuating

Near the end of his 1988 paper "Individualism and Self-Knowledge," Tyler Burge draws a distinction related to the one I have drawn between awareness and understanding. He distinguishes thinking a thought and explicating it. It is correct to do so. I can think the thought that mercury is one of the elements (Burge's example) without being able to give a proper explication of that thought. When I am unable to explicate my thought, presumably it is because I do not adequately understand it. I may know very little about the periodic table and how mercury fits into it, its atomic structure, etc. But Burge mistakenly conflates *explicating* with *individuating*: "One clearly does not have first-person authority about whether one of one's thoughts is to be explicated or individuated in such and such a way." (1988, p. 662) I agree that one does not have first-person authority over the *explication* of one's thought, say, that mercury is an element, but *to have the thought at all* one must be able to somehow *individuate* mercury as what is represented by the minimal content of that thought, no matter how ignorant one might be or how many misconceptions one might have of mercury. Explicating is strongly related to understanding, but individuating a subject of a thought, in the sense of being aware of one's minimal content, is quite independent of both explicating and understanding.

There are two problems here. The first is the problem, just mentioned, of conflating individuating and explicating. The second turns on the dual analysis of the reference of the subject of a thought. Take the latter first. Earlier I indicated a schematization of a thought as  $\Psi(R)$  and stated that it must be decomposed into two schemes:  $\Psi(\Phi(m))$  and  $\Psi(\Phi(o))$  (where  $\Psi$  is some psychological mode,  $\Phi$  is some attribute,  $m$  is the minimal content, and  $o$  is the objective content of the thought). Now, insofar as the

quotation from Burge pertains to a thought in the sense of  $\Psi(\Phi(o))$ , I have no disagreement with him. There is no first-person authority over one's thought in this sense of thought. But things are quite different when the thought is construed as  $\Psi(\Phi(m))$ . Here the conflation of individuating and explicating comes to bear.

To have a thought at all, in the sense of entertaining it, one must be able to individuate the thought, even when one is unable to explicate it. An inability to explicate one's thought manifests one's lack of understanding, but still the thinker would have first-person authority with regard to how the thought and its minimal content are individuated. Individuating a thought or a minimal content is different from explicating either of them, as I now will explain.

To individuate a thought  $\Psi(\Phi(m))$ , one must be able to differentiate it from other thoughts—either because the minimal content is different or something different is attributed to the same minimal content. Obviously, one must have already individuated the minimal content to do either. And one must be able, when the occasion arises, to re-identify it as the same thought as one had before. To explicate a thought, on the other hand, one must explain how that thought relates to other thoughts. One must state a good number of properties the represented object is to have and, where appropriate, show that it possesses those properties. In the case of numbers, this amounts to providing proofs of and commentaries on theorems. Again, our not-very-competent theorist C is unable to do this to any great extent, and so C cannot explicate many of her number thoughts. Her shortcomings, however, do not interfere with her ability to individuate her thoughts and their minimal contents. But her inability to explicate her thoughts does provide reason for our saying that she does not understand them, or, at least, that she does not understand them well.

One can be aware of one's thought  $\Psi(\Phi(m))$  and its minimal content  $m$  without being able to do much explicating; nevertheless, one cannot even begin to explicate a thought unless one has already individuated it. This is true no matter how inadequate one's explications may be. Understanding entails the ability to explicate a thought, and doing this presupposes that one has individuated the thought. But individuating one's thought requires individuating its minimal content, regardless of whether the minimal content matches the objective content of the thought. Therefore, understanding presupposes minimal content.

In saying this I am *not* holding that either the identities of concepts or the meanings of terms are ultimately determined by the individual. Thus, when Burge elsewhere maintains that “the meanings of many terms . . . and the identity of many concepts . . . are what they are even though what the individual knows about the meaning or concept may be insufficient to determine it uniquely” (1992, p. 46), there is—in one sense—no disagreement with what I say. Individual ignorance is varied and extensive. Such ignorance pertains to the objective content and what may be correctly attributed to it, and, thus, to how much one’s understanding of one’s minimal content may diverge from a correct understanding of the thought’s objective content. Still, in another sense, I do disagree with Burge: There is first-person authority over what one’s thought is about in the sense of minimal content, the subject of the thought as conceived by the thinker, regardless of the extent of the thinker’s misconceptions pertaining to the thought’s objective content.<sup>22</sup> The greater one’s misconceptions, the less one is able to correctly explicate one’s own thought, the less one understands it. One neither has first-person privilege over the objective content nor over what counts as understanding, but this is different from individuating or knowing the minimal content of one’s thought. Here, one does have first-person privilege.

If I read Burge’s comments as pertaining to the objective content, I am in complete agreement with him. Specifically, there is no privileged access to the objective content of one’s own thought. But I have argued that any analysis of intentional states must countenance minimal content in *addition* to the objective content, and the situation is quite otherwise with respect to minimal content, as I have tried to show.

It is precisely because minimal content and objective content play distinct roles in the analysis that we are able to avoid talk of understanding and still have something significant and new to say about intentionality or representation. Once we distinguish awareness of minimal content from knowledge of the objective content or understanding the thought objectively, we can go further than we have in showing that the capacity for responding differentially to symbols is not sufficient for determining awareness of minimal content. To see this we can collapse our two individuals A and B into one, dropping the restriction that the individual doesn’t know about alternative models. Call our new theorist D. The syntactic strings that D manipulates are identical whether the theorem she is proving is “about”

numbers or sets. Yet on some occasions D's minimal contents represent numbers and on others they represent sets; she may alter what the symbols represent to her at will. (This is similar to the "new twist" on the Kutuzov example discussed earlier.) Surely it is wrong to say that D's differential symbol manipulations account for, or in any way determine, the different contents that she can be aware of, for these are identical though the minimal contents are different. Nor, for the same reason, does the converse relation hold. Thus, appropriate differential responses to symbols, on the one hand, and awareness of one's minimal content, on the other, are once again seen to be utterly independent in the sense described above.

### The Fundamental Intentional State

I am now in position to make a rather singular and remarkable claim:

*The subjective constitution by an agent of her minimal content makes it a different and unique intentional state—one that is presupposed by all other intentional states.*

This subjective constitution is what characterizes the fundamental intentional state. Here is how it differs from all other, "normal," intentional states. Generally, intentional states involve an agent S having a certain psychological attitude  $\Psi$  (believing, desiring, and so forth) toward a representational content R. Schematically,  $S(\Psi(R))$ , where R may represent a thing or a state of affairs. Call states having this structure *normal intentional states*. It is fundamental on my theory that R requires *both* an objective (o or  $\Phi(o)$ ) and a subjective (m or the  $\Phi(m)$ ) characterization. Now, the minimal content, m, being *about* something, is itself "directed," but, crucially, not in the way that representational content as objectively characterized is directed. On the objective reading of R, R is not constituted by the agent. In contrast, on the subjective reading of R, as m or  $\Phi(m)$ , it or a component of it is subjectively constituted. Thus, while one's having a minimal content is itself an intentional state, it is a unique one. Its logical structure differs from that of normal intentional states. The having of minimal content *cannot* be characterized as a *psychological attitude directed at something else*, as normal intentional states are.

Unlike all other contents, the agent constitutes minimal content. Minimal content does not merely represent; it has the uniquely singular property of *representing in virtue of the agent's constituting the content*. The

minimal content of the act and the act itself—the constituting—are *not logically separable*. Thus, if one were to attempt to construe the constituting act as a psychological attitude, it would be distinctively different from the usual psychological attitudes, since all of those are logically separable from their contents. Therefore, the *aboutness* relation involved with minimal content is unique.<sup>23</sup> This is a wonderful consequence, since if it had the same logical structure as that of normal intentional states there would be a danger of infinite regress.<sup>24</sup>

Minimal content is at once different from and presupposed by all normal intentional states. It is *the fundamental intentional state*. I call it fundamental because there can be no normal intentional states without it. I call it intentional not because it has the same logical structure as normal intentional states (it does not) but because it has the central feature of such states (“aboutness,” though, as I have emphasized, it is about things in way peculiar to it)—unlike normal intentional states, it involves a constituting relation that is logically inseparable from the content it constitutes.<sup>25</sup> This difference goes to the very heart of the ambiguity in the locution ‘what a thought is about’, an ambiguity that is accommodated by the introduction of the distinction between minimal content and objective content. If I am right about minimal content, however, the fundamental intentional state is not similarly ambiguous. It follows that the subjective side of the ambiguity is logically prior to its objective side.

### The New Problems of Absent and Inverted (Non-Phenomenal) Content

The concept of minimal content has other interesting consequences. I argued in my discussion of the theorists A–D that, no matter how correct and elaborate the symbol manipulations are, *such manipulations are not sufficient to individuate one minimal content from another; still, the content is easily and directly individuated and differentiated by the agent, she is aware of what it is, and all this is achieved from her first-person perspective*. Theorists A, C, and D differentiate numbers (or particular numbers) and are able to distinguish them from sets (or particular sets). This ability and their awareness of minimal content do not require that they make an inference, though an observer of them must make inferences regarding their contents.

I have also argued that a difference in minimal content for an individual does not require a difference in symbolic manipulation. (Compare theorist

D.) Thus, the symbol manipulation is neither necessary nor sufficient for determining different minimal contents. Since the symbol manipulation is not sufficient to determine either A's or B's minimal content, there is here a sort of indeterminacy of minimal content from the third-person perspective, though each theorist has a definite minimal content that is different from the other's and is directly accessible to the theorist who has it.<sup>26</sup> As a result of this, we have an analogue to the problem of the inverted spectrum. Given the behavior of A and B as observed from a third-person perspective, we can attribute thoughts about numbers or sets to either A or B, but we cannot determine on that basis which minimal content the thinker has—an inverted minimal content problem.

We can now go further. Such manipulations, which are accessible from a third-person perspective, are not sufficient to say that the symbols have *any* content for the processor. Programs for proving theorems can and have been developed. Some are able to prove more difficult theorems than others, and to do it more efficiently.<sup>27</sup> Still, we have no reason to think that the symbols have any content for the program or for the machine running the program.<sup>28</sup> We could even build into the program the "disposition" to display 'number', or '2', when asked appropriate questions. The incorporation of this little programming task yields no basis for holding that the computer running the program has special access to its minimal content, not even that it has minimal content. None of this would provide reason to hold that the symbols it manipulates have content *for it*.<sup>29</sup>

Indeed, the same result that applied to a computer running a program could be achieved with a person. Let E be yet another theorist who learns the same formal theory as did A and B, but learns no semantics for it beyond that for the logical constants. E may still become quite adept at manipulating the symbols in accordance with the *formal* rules and *uninterpreted* axioms of the theory, as adept as A and B. But with E we go beyond this indeterminacy of minimal content relative to the third-person perspective. Here we have *absent content*. All the relevant symbol manipulations may be realized in E's activities, but E has *no* minimal content relative to these symbols or their manipulations.

The expression 'absent content' is obviously chosen because my point parallels one that Ned Block (1978) raised against functionalism. The general idea Block argued for, in opposition to functionalism, was that, though all the functional roles that a functionalist might require could be in place

in some system, it would still be plausible in some such cases to deny that the system has any qualia. My arguments establish that all the relevant functional roles could be in place, yet a *non-phenomenal* narrow content—minimal content—would not be thereby determined or is in fact absent. Theorists A and B exemplify the minimal-content analogue to the inverted spectrum. Theorist E is the minimal-content analogue to the absent-qualia problem and is also a special case of the intentionality analogue to the “zombie case.” However, my support for absent and inverted minimal content is stronger than the support offered by others for the inverted spectrum, absent qualia, and zombie problems. My arguments, unlike the others, *do not depend on mere conjectures of possibilities—they are verifiable from the first-person perspective.*

### **Constituting and Grasping Minimal Content: Fregean, Cartesian, and Searlean Comparisons**

We have seen that from the first-person perspective that there is a kind of direct “grasping” of content that does not admit of degrees and is distinct from understanding the content. This grasping of minimal content, our non-inferential awareness of it, is largely a function of the fact that we subjectively constitute it. I make no claim as to just how we are able to constitute minimal content, but the various cases presented (forming an image, the battle diagram, with and without the “twist,” the various theorists recently considered) make evident that we do constitute minimal content. Because of the subjective character of such constituting, this is evident only if one considers these cases from the first-person perspective. The crucial subjective features thus revealed are opaque from a strictly third-person perspective. Saying this, however, no more makes the act of constituting minimal content “spooky” or mysterious than is any other empirical phenomena *that* we know occurs without yet knowing *how* it occurs.

#### **Frege**

The minimal content of a thought is the subject of a thought as conceived by the thinker. Some may be tempted to assimilate minimal content to a Fregean sense, understood as a “mode of presentation,” for it is natural to take the expression ‘as conceived by’ as though it were under a certain mode of presentation, or the subject of the thought under a certain “guise.”

Such temptations must be resisted, for such a construal would be seriously misleading. My view of minimal content departs significantly from Frege's concept of sense. First a minor difference: By definition, both minimal content and objective content can only signal an object, whereas Fregean senses may also signal states of affairs. A much more critical difference is that I hold that a *single thought* is properly analyzed as having *both* a minimal content and an objective content; for Frege a thought has only one sense. Finally, and most important, minimal contents, depending as they do on how the *individual* conceives things, are *subjective* in a way in which Frege insisted senses are not.<sup>30</sup>

The last point is of great importance when making a comparative evaluation of Frege's claim that we grasp senses and my claim that we grasp minimal contents. The idea of grasping plays a central role in both of our views. But since minimal contents are unabashedly subjective and are in part constituted by the thinker, the idea that we grasp our own minimal contents seems clearly right. In contrast, just how a thinker is to grasp an objective Fregean sense, a sense that the thinker does not even in part constitute, does appear a bit mysterious; it is not addressed by Frege. The concept of minimal content clearly differs from and has a clear advantage over the concept of Fregean sense, at least in this respect. From the first-person perspective, it is *constitutive* of my entertaining some particular thought that it has the minimal content that it has. It simply would be a different thought if it did not have the subject that I conceive it to have. The grasping of minimal content by the thinker is unproblematic because the act of thinking that particular thought is a subjective act constitutive of its minimal content.

### Descartes

Does our awareness of our minimal content amount to a privileged access of the Cartesian sort? The latter is typically portrayed so that the contents of one's mind are completely and infallibly transparent to oneself. I am not convinced that this is a fair portrayal of Descartes, but I will not attempt here a scholarly defense of a different reading of him.<sup>31</sup> The common interpretation of Descartes' view, whether it is his view or not, has been widely and rightly criticized. Certainly, the objective content is not completely and infallibly transparent to the agent. I do not challenge this. So long as 'content' is construed as the objective content there is no privileged access to it, none whatsoever. But this is no reason for *all* privileged access to *any kind*

of content of one's thoughts to fall into disrepute. Minimal content is a distinct content from either the objective content or the representative content.<sup>32</sup> The privileged access that I endorse applies only on certain occasions and is severely restricted to minimal content, *m*. Moreover, any infallibility with regard to what the minimal content represents which results from this privileged access turns on the *senselessness* of an attribution of error, not upon our having some special ability or faculty. Thus, a limited, non-Cartesian form of privileged access is preserved.

### Searle

My view is most sympathetic to Searle's; indeed, it builds upon it and so shares a number of its features. Importantly, the capacity for awareness of minimal content, being a necessary condition for an agent's having intentional states, is a commitment to intrinsic intentionality. Searle emphasizes intrinsic intentionality, as distinct from derivative or metaphorical intentionality, though he does not identify minimal content as a critical component of it.

Searle's idea of intrinsic intentionality is famously (or infamously) presented in his Chinese Room thought experiment.<sup>33</sup> One might wonder whether I have presented a remodeled Chinese Room. In Searle's thought experiment, the central concept is that of *understanding*; whether the room system has subjective phenomenal states is not at issue (pace Van Gulick—see chapter 3). Though Searle does not speak explicitly of content in that work, clearly it is involved; the question "Does the system understand?" may fairly be put as the question whether the strings of marks have any content for the agent (or system) manipulating them. Though I fundamentally agree with Searle regarding his views of intrinsic intentionality and the Chinese Room, casting the associated issues in terms of awareness of minimal content has several advantages over talk of understanding and natural languages.

First, rather than trying to contrast the Chinese Room with a genuine speaker of Chinese who "understands" (with all the attendant murkiness of this notion), the contrast is instead drawn in relatively simple and clear terms: straight out, one is aware that one's minimal content is, say, sets or numbers. So we take several steps back from understanding to awareness of minimal content. Crucially, and unlike understanding, minimal content does not come in degrees. As a result, we do not have to rely on, or attempt to resolve, conflicting intuitions as to whether or not a system such as the

Chinese Room understands. In the cases I considered, each individual clearly is aware of their respective minimal contents—despite radically different levels of understanding—whether those contents are the same (e.g. theorists A and C compared) or different (e.g. theorists A and B compared).

Another advantage is that a move against the Chinese Room commonly made by functionalists and others is defused. They often wish to bring in causal interaction of the system with its environment. Since such interactions are relevant in our own case for language understanding, they argue that such causal interaction must also be extended to the room system, prior to its getting any serious attention. I think this move is fundamentally mistaken (see chapters 2 and 3); still, the importance of such interaction may initially seem plausible when it is a question of language understanding in general. In part, it may seem plausible because of the pervasiveness of terms referring to things that occur in our environment. Whatever the reasons, it is clear that such causal interaction with the environment (other than the trivial ones for input and output) is not even initially plausible when we speak of abstract contents such as numbers or sets.

### Explaining the Appearance

Given that it is at least initially plausible that we sometimes have some limited privileged access to the content of our intentional states, a special restriction is placed on any proposed account of intentionality: It must either include an account of this feature or explain why it appears to be the case but is not.

The force of this restriction might be made clearer by an analogy. George Berkeley argued that material substance did not exist (or rather that the notion itself was incoherent). But ordinary thinking and ordinary experience seem to present material things as uncontroversial data. In sections 34–81 of *A Treatise Concerning Principles of Human Knowledge*, Berkeley particularly addresses himself to objections based on this type of consideration—objections that his view denies various obvious truths, or that it obliterates various obvious distinctions, such as that between a real thing and a chimera.<sup>34</sup> In point of fact, while rejecting material substances, he took great care to show exactly why one cannot, say, simply walk through real walls, whereas one can “walk” through imagined walls, even though on his view both are collections of ideas. (Whether he was ultimately

successful is entirely immaterial to the little moral I wish to draw from my discussion of him.) In doing so, Berkeley offered an explanation of how certain obvious data that seemed to refute his view were not only consistent with but explained by his account.

That Berkeley did explain why such obvious data are in accord with his theory is precisely why Samuel Johnson's famous "refutation" of Berkeley,<sup>35</sup> though cute in its vivid and immediate appeal to what seems to be a conflicting datum, can never be taken seriously. Had Berkeley failed to provide these explanations, Johnson would have had a formidable point. Unlike Berkeley, many present-day philosophers prefer to reduce, eliminate, or give an objective externalist account of the mental rather than rejecting the material; nevertheless, if they are to avoid a Johnsonian refutation (in reverse), they, too, must explain certain appearances, though their task is quite different from that which Berkeley confronted.

It is the kind of privileged access to one's (minimal) contents brought out by the various examples discussed earlier that constitutes the sort of data that produce such severe obstacles for strictly third-person methodologies. The obstacles thus raised are raised as data. Such data do not require some elaborate theoretical analysis (internalist or otherwise) to shore them up so that they can play this role. Not only do we have mental states; we also sometimes directly know (some of the) content of those states, its minimal content. When I speak here of 'being aware of the content' or 'knowing what the content is', I repeat that I do *not* refer to its *ontological* status, but only to how it is described in virtue of the subject's non-inferential knowledge of it. Whatever its ultimate ontological status, it is data of the utmost importance.

Heil (1988, p. 247) cites with approval Davidson's claim that we are victims of a certain misleading picture of the mind—to wit, ". . . the content of one's mental states are taken to be based on inward glimpses of those states or on the grasping of particular entities." Davidson recommends abandoning such a picture, and Heil claims that once we do so we "remove at least one of the reasons for supposing that externalism undermines privileged access" (ibid.). I applaud this abandonment. However, this picture is not the only obstacle for externalism or for strictly objective accounts. Abandoning it does not clear the way for such accounts of privileged access.

When I argue (in the next chapter) that these accounts cannot provide for any privileged access to contents, I do not rely on any analysis of this privileged access that presupposes such "inward glimpses of those states or

on the grasping of particular entities." I am not committed to this model of introspection or to the grasping of *entities*. The metaphor used to characterize or single out the datum is inconsequential relative to the datum itself. I do not think that content is an entity, nor do I think we have "inward glimpses" of it, nor do my arguments depend on any such assumptions. *Still, we are sometimes aware of content.* We are not, however, aware of it in the sense of having special or detailed knowledge of its nature or ontological status. To hold the latter would be to invoke the thoroughly discredited view that our privileged access to our own mental states gives us infallible and incorrigible knowledge of the *very nature* of those states. I certainly do not hold that view.

Nor should my recent comments be construed as my harboring some view such as that there exist "pure data," that there is a "given," or that common-sense beliefs are somehow epistemically privileged. I hold no such views. Such views are, to my mind, also thoroughly discredited. Though there are no "pure data," it is far from evident that the indicated datum—that sometimes we have privileged access to (at least part of) our own content, our minimal contents—is such that it is essentially dependent on this particular (wrong-headed) picture of the mind. Rejecting the picture does not, *ipso facto*, eliminate the datum in question; puzzles remain regarding an account of this datum.

As I see it, any account that employs an exclusive third-person methodology must fail in accommodating privileged access to (minimal) content or even recognizing the latter. If that is correct, the only option for someone employing a strictly third-person methodology is to dissolve the datum by explaining it away as mere appearance—that is, by showing that privileged access to our own contents is just an illusion. If a theory fails in the latter and lacks the resources to account for the appearance, then the theory must be rejected.

If I am right, minimal content and awareness of it are fundamental to a theory of mind. Recognition of this provides a better foundation for the analysis of mind than can any theory based on a strictly third-person methodology. Since reliance on the latter methodology has dominated the philosophical study of the mind, and since minimal content is evident only from the first-person perspective, it is not surprising that minimal content has gone largely unnoticed.