

# Mind

## 48. Intentionality

Every mental phenomenon is characterized by what the scholastics of the Middle Ages called the intentional (and also) mental inexistence of an object, and what we would call, although not in entirely unambiguous terms, the reference to a content, a direction upon an object (by which we are not to understand a reality...), or an immanent objectivity. Each one includes something as an object within itself....

This intentional inexistence is exclusively characteristic of mental phenomena. No physical phenomenon manifests anything similar. Consequently, we can define mental phenomena by saying that they are such phenomena as include an object intentionally within themselves. (Franz Brentano, *Psychologie vom empirischen Standpunkt*)

An intentional phenomenon, according to Franz Brentano, is one that makes reference to, is directed upon, or is about other objects, perhaps even objects that do not exist. Intentional phenomena can in this sense be said to 'include an object intentionally within themselves'. Intentionality, then, is that special property of being directed upon something (*Gerichtetheit*). Brentano used this concept of intentionality to formulate a two-part thesis that has come to be known as *the thesis of intentionality*:

- (1) All and only mental phenomena are intentional.
- (2) No purely physical phenomenon is intentional.

This is to say, (1) intentionality, i.e., the special property of directedness or aboutness, is the mark of the mental, and (2) it sunders the mental from the purely physical. I will return to Brentano's thesis in a while, but at the moment my concern is with the concept of intentionality itself. For though Brentano is to be credited with the modern rediscovery of intentionality, his analysis of it is inadequate. My immediate goal is to define intentionality without appealing to the metaphors of directedness and inexistence.

I will begin by giving a schematic summary of Brentano's theory

of judgement as it is reported by Roderick Chisholm.<sup>1</sup> Brentano's theory differs sharply from the propositional/relational theory of judgement that I have been espousing in this work. On Brentano's theory, when one judges that  $(\exists x)Ax$ , one does not stand in relation to the proposition that  $(\exists x)Ax$ ; nor does one stand in a relation to the concept of being an  $A$ . Instead, one affirms or accepts  $As$ . Likewise, when one judges that  $\neg(\exists x)Ax$ , one does not stand in a relation to the proposition that  $\neg(\exists x)Ax$ . Rather, one denies or rejects  $As$ . In the same vein, to judge that  $(\exists x)(Ax \& Bx)$  is to accept  $As$  that are  $Bs$ . To judge that  $\neg(\exists x)(Ax \& Bx)$  is to reject  $As$  that are  $Bs$ . To judge that  $(\exists x)(Ax \& \neg Bx)$  is to accept  $As$  that are non- $Bs$ , and to judge that  $(\forall x)(Ax \supset Bx)$  is to reject  $As$  that are non- $Bs$ . In increasingly awkward steps Brentano thus attempts to extend his theory to complex judgements.

Non-propositional/non-relational theories of judgement are not rare. Evidence of them is found in works ranging from Plato's *Sophist* (240D, 260C–263D) and *Theaetetus* (188E–189A) to Russell's *The Problems of Philosophy* (chapter 12) and his introduction to the first edition of *Principia Mathematica* (pp. 43–4). Yet all such theories share a flaw, indeed, the very flaw that spells defeat for adverbial and multiple-operator approaches to intensional logic. (See §6.) Even if by various awkward maneuvers these theories can handle statements concerning particular judgements, they cannot handle general statements concerning judgements. To handle general statements, one must be able to bring the theory within the scope of quantifier logic, and this is precisely what non-relational/non-propositional theories are unable to do in a credible way. Consider, e.g., the following intuitively valid arguments:

Whatever  $x$  believes is necessary.  
 Whatever is necessary is true.  
 —————  
 $\therefore$  Whatever  $x$  believes is true.

Whatever  $x$  believes is true.  
 $x$  believes that  $A$ .  
 —————  
 $\therefore$  It is true that  $A$ .

$x$  believes that  $A$ .  
 —————  
 $\therefore$   $x$  believes something.

I argued in §6 that on the canonical syntactic treatment of such

arguments 'believes' is represented as a 2-place predicate and 'that *A*', as a singular term. Only then do the arguments submit to a plausible treatment within quantifier logic:

$$\begin{array}{l}
 (\forall y)(B(x, y) \supset Ny) \\
 (\forall y)(Ny \supset Ty) \\
 \hline
 \therefore (\forall y)(B(x, y) \supset Ty) \\
 \\
 (\forall y)(B(x, y) \supset Ty) \\
 B(x, [A]) \\
 \hline
 \therefore T[A] \\
 \\
 B(x, [A]) \\
 \hline
 \therefore (\exists y)B(x, y).
 \end{array}$$

Once this conclusion is reached, however, one is obliged to determine what special values are included in the range of the variable 'y' and to what the singular term 'that *A*' is semantically correlated. I argued in §8 that on the canonical semantical treatment propositions are the entities that fill the bill. I then concluded that, since 'believes' is a 2-place predicate, it expresses a binary relation whose range is made up of propositions. And this conclusion leads directly to the relational/propositional theory of judgement. However, on the non-relational/non-propositional theory of judgement there is no credible way even to express the above intuitively valid arguments, for this theory denies from the start that belief consists in standing in a relation to truth bearers. The non-relational/non-propositional theory thus falters at the earliest possible stage: it collides with logic itself.

Not only does the relational/propositional theory of judgement mesh easily with logical theory, it also makes possible the first step toward clarifying the phenomenon of intentionality, a step that is out of the question for Brentano's theory and the other non-relational/non-propositional theories. Let me explain. It is commonplace to say that thoughts, beliefs, judgements, etc. are about or directed toward other objects. This aboutness or directedness is what Brentano means by the intentionality of thought, belief, judgement, etc. Now each expression in the family 'thought', 'belief', 'judgement', etc. has at least three related uses. Each can be used to mean (1) a kind of intentional act, (2) the propositional object of the intentional act, or (3) a relation holding between

persons performing the intentional act and the propositional object of the act. The non-relational/non-propositional theory acknowledges only the first of these three uses, the one for intentional acts. This forces the theory to give its account of intentionality in the inevitably opaque terms of intentional acts, making metaphor and circularity unavoidable. By contrast, the relational/propositional theory acknowledges all three uses and, thus, is free to analyse intentional acts in terms of the associated relations and their propositional objects. The following is an illustration of how easy these analyses can be: *x* performs the intentional act of thinking that *A* if and only if *x* stands in the thinking relation to the thought that *A*.

However, I have said nothing yet concerning the intentionality of intentional acts, i.e., their directedness or aboutness. How does that arise? The answer is that it arises from the propositional objects, i.e., from the thoughts to which the person stands in the relation thought, belief, judgement, etc. After all, thoughts in the propositional sense are themselves things that are characteristically said to be about other objects; indeed they are often said to be about objects that do not exist. And the same thing holds not just for thoughts but for all complex ideas.<sup>2</sup> (Unlike complex ideas, simple ideas are just qualities, connections, and conditions. As such, they are not said to be about anything. E.g., the color red is not said to be about anything. Qualities simply qualify things; connections simply connect things; and conditions simply obtain. For example, recall the little triangle figure displayed in §40. One might say, e.g., that a set containing that triangle "involves" or "pertains to" the triangle but not that it is "about" the triangle; analogously, though the condition that the little figure is triangular involves or pertains to the triangle, it is not about the triangle. The condition is just there on the page. An intensional entity can be said to be about other things only if it has a complex logical form.<sup>3</sup>) Now a complex idea—whether it be a thought or a complex concept—can be said to be about other things even when the idea is not the object of any intentional act. (There are thoughts about Hamlet that no one has as yet had.) An intentional act, on the other hand, can be said to be about those and only those things that the associated complex idea can be said to be about. These facts strongly suggest that an intentional act can be said to be about other things for one reason only; namely, the intentional act

consists in standing in a certain relation to a complex idea that can be said to be about things. This is to say, an intentional act can be said to be about other things only secondarily through the complex idea that is the object of the act. Complex ideas are the objects that can in a primary way be said to be about other things.

But what features do complex ideas have that allow us to say that they are about other things, including even things that do not exist? Using the theory of qualities and concepts, we can answer this question. The answer goes in pragmatic stages, however, since from a linguistic point of view this idiom of aboutness is an extremely context-dependent affair.

The tightest way in which a complex idea might be said to be about an object occurs when the idea is the result of predicating something of the object itself. For example, the thought that  $Fx$  might be said to be about the object  $x$  since this thought is the result of predicating  $[Fy]$ , of the object  $x$ . A second and somewhat weaker way in which a complex idea might be said to be about an object occurs when the idea is descriptive in character, i.e., when the idea can be denoted by an intensional abstract of the form  $[A(1x)(Fx)]_x$ . Here the idea might be said to be about the unique object to which the associated descriptive concept applies, i.e., about the unique object that stands in the  $\Delta$ -relation to the concept  $[Fx]_x$ . In this case the object need not even exist: for example, in an appropriate conversational context I might truly say that the thought that the golden mountain does not exist is about the golden mountain.<sup>4</sup> A third and rather weak way in which a complex concept might be said to be about an object occurs when the concept applies to the object (i.e., when the object stands in the  $\Delta$ -relation to the concept). In this vein, a thought might be said to be about an object if the analysis of the thought contains a concept that applies to the object. Fourthly, there is a very weak way in which a complex concept might be said to be about objects. For any complex concept  $[A(\alpha)]_x$ , we might in an appropriate conversational context say that  $[A(\alpha)]_x$  is about  $As$  regardless of whether any  $As$  actually exist. Analogously, we might say that a thought is about  $As$  simply if a concept  $[A(\alpha)]_x$  occurs in the analysis of the thought. For example, we might say that the thought that no witches exist is about witches. Finally, in addition to the above ways in which a complex idea might be said to be about objects, there are myriad intermediate ways that depend on fine points of

logical form, antecedently determined interests, mutually held beliefs, etc.

Now every complex idea in some context or other could be said to be about something, even if that something does not exist. Which thing a complex idea is said to be about depends in large part upon the conversational context. In one conversational context, for example, it might be appropriate to assert that the thought that the witch blighted the sheep is about the witch, and yet in another, to deny this on the grounds that there are no witches. It all depends on what is deemed relevant in the context.

According to this analysis, then, the word 'about' has no semantically fixed extension; its extension is pragmatically determined. But we should not expect this to be otherwise, any more than we should expect the word 'relevant' to have a semantically fixed extension. Indeed, 'about' and 'relevant' have much in common: 'What are you talking about?' is very close to 'What is relevant to your conversation?'. The important point about the above analysis is that it successfully identifies all the formal features of complex ideas that determine what in a given context the thought or concept can be said to be about. And it accomplishes this using only purely logical terms provided by the theory of qualities and concepts.

So far I have reached the following conclusions. First, there are independent logical grounds supporting the relational/propositional theory of judgement. Secondly, using the relational/propositional theory, one is able to analyse intentional acts in terms of the associated relations and propositional or conceptual objects. Thirdly, the intentionality (i.e., the directedness or aboutness) of an intentional act can be accounted for by the fact that the intentional act consists in standing in an appropriate relation to a complex idea—either a thought or a complex concept—which given the right context can always be said to be about other objects, even objects that do not exist. Fourthly, using the theory of qualities and concepts, one can identify all the formal features that are at work in determining what in a given context a complex idea can be said to be about.

Yet the story is not complete, for there is an unsolved problem. Standing in just any relation to a complex idea does not constitute an intentional act. Only certain very distinctive relations will do—relations such as thinking, believing, judging, remembering, perceiving, desiring, deciding, intending, etc. These relations, naturally

enough, are called intentional relations. The problem is to give a non-circular definition of what an intentional relation is. If this problem can be solved, then the analysis of intentionality will be complete.

In the English-speaking world interest in intentionality and Brentano's thesis has been generated to a great extent through the efforts of Roderick Chisholm. Many of Chisholm's ideas on intentionality are expressed in his well-known published correspondence with Wilfrid Sellars 'Intentionality and the Mental'. In this correspondence Sellars asserts that every intentional sentence can be analysed into some sentence that uses only non-intentional vocabulary. In opposition, Chisholm echoes the second half of Brentano's thesis of intentionality by asserting that no intentional sentence can be analysed without appeal to further intentional vocabulary. Chisholm, however, also maintains that all intentional sentences have certain purely logical properties that are not shared by any non-intentional sentences. A bit later I will return to the Chisholm/Sellars dispute. At present, it is Chisholm's purely logical criterion for intentional sentences that interests me.

Over the years Chisholm has offered a variety of criteria. His best one appears in his article 'Intentionality' in *The Encyclopedia of Philosophy*. There he attempts to define the narrower concept of a *simple intentional sentence prefix*; he does this in the hope that the wider concept of an intentional sentence can then be defined in terms of this narrower concept. A sentence prefix is any expression that, when prefixed to a sentence, yields a new sentence. A simple sentence prefix is one that contains no meaningful proper part that is a sentence function, where a sentence function is a special kind of expression that, when supplied with a further expression, yields a sentence. Chisholm's definition is this:

A simple sentence prefix  $M$  is intentional if and only if for every sentence  $A$ ,  $M(A)$  is a contingent sentence.

Consider some examples. The simple sentence prefix 'John believes that' qualifies as intentional according to the definition since every sentence (e.g., ' $1 + 1 = 3$ ') is such that the result of prefixing 'John believes that' (e.g., 'John believes that  $1 + 1 = 3$ ') is contingent. By contrast, 'necessarily' does not count as intentional because there are sentences (e.g., ' $1 + 1 = 3$ ') such that the result of prefixing 'necessarily' (e.g., 'necessarily,  $1 + 1 = 3$ ') is not contingent.

Chisholm's ingenious definition seems to be founded on two insights into the nature of intentionality. The first harks back to Brentano's original characterization of intentional phenomena as those that are about objects and perhaps even about objects that do not exist. I have argued that intentional acts consist in standing in intentional relations to thoughts or to complex concepts and that thoughts and complex concepts are the sort of thing that can be said to be about other things, perhaps even objects that do not exist. Now if a relation can hold contingently between an individual and a thought or complex concept, then typically the relation will hold between the individual and the thought or concept independently of whether the thought is true and independently of whether the concept applies to anything. (Factive intentional relations, e.g., knowing, are an exception to this; but since they are definable ultimately in terms of non-factive intentional connections, e.g., believing, Chisholm's insight stands with this qualification.) However, if a thought is not true and if a complex concept does not apply to anything, then in an appropriate context we could say that they are about objects that do not exist. This, then, is the link up between Chisholm's definition and Brentano's original characterization. The second insight upon which Chisholm's definition seems to be founded concerns the concept of a phenomenon. If the condition that a certain relation holds between a certain pair of objects is either necessary or impossible, then the condition cannot be considered to be a phenomenon. The condition would be a phenomenon only if it were contingent. The insight is that fundamental intentional relations are ones that give rise to genuine phenomena and, therefore, must be able to hold contingently at least between certain relata.

Though Chisholm's definition helps us to uncover these two important insights into the nature of intentionality, it encounters a number of difficulties. To begin with, since a sentence prefix is a kind of operator, Chisholm's treatment of intentional language must be classified as a special case of the multiple-operator approach to intentionality. In this, his definition contains vestiges of the non-relational theory of judgement. Not surprisingly, then, general intentional sentences, which spelled defeat for the non-relational theory, also appear to cause trouble for Chisholm's treatment of intentional language. For example, since 'Whatever *x* believes is true' contains the predicate 'believes', which expresses an



intentional relation, it would seem to qualify as an intentional sentence. However, it is not clear that Chisholm's treatment helps to show this, for predicates are not sentence prefixes. It would seem better, therefore, simply to bypass Chisholm's linguistic superstructure and instead to define the concept of intentional relation straightaway. This is what I advocate.

To see the other difficulties with Chisholm's definition, one should realize why Chisholm limits the definition to just those sentence prefixes that are simple (i.e., to those sentence prefixes that contain no meaningful proper parts that are sentence functions). For if he did not limit it this way, then numerous easy counterexamples (generated by "meaningful proper parts") would arise. Now if the definition does any more than isolate some accidental feature of, say, English, it should apply to any language that might come to be spoken. But in that case, the very same sort of counterexamples beset the definition in spite of its limitation to simple sentence prefixes. For example, let English\* be the language that results when English is supplemented with certain new vocabulary items. In English\* let  $M_1$  be a syntactically primitive sentence prefix that is synonymous to the complex sentence prefix:

Socrates is Greek & (\_\_\_\_\_ or not \_\_\_\_\_).

Now observe that, for any sentence  $A$ ,  $M_1(A)$  is a contingent sentence. Take an example:

Socrates is Greek & ( $1 + 1 = 3$  or not  $1 + 1 = 3$ ).

is a contingent sentence. Thus,  $M_1(1 + 1 = 3)$  is contingent. In this way,  $M_1$  qualifies as intentional according to the definition. However,  $M_1$  is clearly not an intentional sentence prefix. Thus, the definition does not provide a sufficient condition. The next example shows that the definition does not provide a necessary condition. Let  $M_2$  be a syntactically primitive sentence prefix that is synonymous to the complex sentence prefix:

(\_\_\_\_\_) or John believes that (\_\_\_\_\_).

There are numerous sentences  $A$  such that  $M_2(A)$  is not contingent. For example, since ' $(2 + 2 = 4)$  or John believes that  $(2 + 2 = 4)$ ' is necessary,  $M_2(2 + 2 = 4)$  is necessary. Yet there is good reason for thinking that  $M_2$  is an intentional prefix: ' $(1 + 1 = 3)$  or John believes that  $(1 + 1 = 3)$ ' is necessarily equivalent to 'John believes

that  $(1 + 1 = 3)$ ', which is undeniably intentional.<sup>5</sup> For another kind of example that shows that the definition does not provide a necessary condition, let  $M_3$  be synonymous to the odd prefix:

9 believes that.

This prefix fails to satisfy the definition. E.g.,  $M_3(2 + 2 = 4)$  is not contingent since '9 believes that  $2 + 2 = 4$ ' is not contingent. At the same time,  $M_3$  would seem to qualify as an intentional prefix. To see why, notice that  $M_3(A)$  expresses the proposition that 9 believes that  $A$ . Using only the law of existential generalization (EG), we can derive an undeniably intentional proposition, namely, the proposition that something believes that  $2 + 2 = 4$ . Consider a related example. Let  $M_4$  be synonymous to the prefix:

This positron believes that.

Perhaps  $M_4$  is a counterexample since, for all we know, 'This positron believes that  $2 + 2 = 4$ ' is not contingent. For a final example, let  $M_5$  be synonymous to the prefix:

John introspects that.

This is a counterexample since, e.g., 'John introspects that someone other than John is in pain' is not contingent.<sup>6</sup> However,  $M_5$  is undeniably intentional.

What do these counterexamples show us about intentional relations?  $M_1$  shows that there are certain non-intentional Cambridge relations whose logical behavior resembles that of genuine intentional connections. In a similar vein,  $M_2$  shows that there are certain intentional Cambridge relations whose logical behavior fails to resemble that of genuine intentional connections. This suggests that the analysis of intentional relations should proceed in two steps. The first step is to define the concept of an intentional connection. And the second step is to define the concept of an intentional Cambridge relation in terms of this concept of intentional connection.<sup>7</sup> Once the first step is completed, the second step will be straightforward; for a Cambridge relation is intentional if its definition involves an intentional connection in a logically essential way. Thus, the crucial step is to define what it takes for a connection to be intentional.<sup>8</sup> But what do the remaining three counterexamples  $M_3$ ,  $M_4$ , and  $M_5$  show us? The lesson of  $M_3$  is

that universals cannot be intentionally connected to complex ideas; only individual particulars can. Next,  $M_4$  indicates that not just any individual particular can be intentionally connected to complex ideas; only certain ones can. And finally,  $M_5$  shows us that an individual particular cannot be intentionally connected to just any complex idea; there are epistemic limits. Nevertheless, for each intentional connection it is possible that at least some individual particulars are connected by it to at least some complex ideas.

Assembling the insights isolated over the last few pages, I offer the following analysis:

A connection is intentional if and only if it can contingently connect an individual to a complex idea independently of the veracity of the idea.<sup>9</sup>

(An idea has veracity *iff* it is a true thought or a concept that applies to something.) The claim here is that all and only intentional connections have this special logical character. We intentional beings are distinctive in that we can stand in contingent connections to complex ideas independently of the veracity of the ideas. There are complex ideas with regard to which we can (but need not) believe, contemplate, decide, doubt, remember, want, etc., and we can (but need not) do this quite independently of whether these ideas are true or whether they apply to anything. In thought we can (but need not) do all sorts of things that are about, or at least purport to be about, objects in the world, even though these things we do need not in any relevant way correspond to the actual conditions of objects in the world. Thus, to think is to engage the dual possibilities of truth and falsehood, possibilities born for us through the weaving together of forms.

There are a number of candidate counterexamples to this analysis of the concept of intentional connection. Though many appear promising at first, none hits its mark. Still, some of them are of philosophical interest in their own right. For that reason, as well as for the reason of imparting a better feel for the analysis, I will explain why the best of these counterexamples fail.

(1) Take first the *predication relation*. True, this relation is a genuine connection, and it can connect individuals to complex ideas. However, it can connect an individual to a complex idea only when that idea has instances (e.g., the individual itself). Since the

analysis requires that intentional connections hold independently of whether the idea has instances, the predication relation is not a counterexample.<sup>10</sup> The *predication operation* fails to be a counterexample for a related reason. Though this operation is a genuine connection, when it holds it holds necessarily. But on the analysis it must hold contingently in at least some possible cases

(2) Certain *causally grounded dispositional relations* constitute another kind of potential counterexample. (Dispositions that are not causally grounded are not connections and so are not a source of counterexamples. For more on dispositional counterexamples, see note 24.) Consider the dispositional relation holding between particulars *x* and conditions *y* such that *x* is a hunk of salt having the disposition to dissolve whenever condition *y* obtains. This relation holds, for example, between the hunk of salt *v* in my hand and the condition *c* that *v* is submerged in the water in my glass. This relation is certainly not intentional. Yet it can hold contingently; it might not have held between *v* and *c*, for the water in my glass might have already been saturated with salt. And it holds between *v* and *c* whether or not *c* ever obtains, i.e., whether or not *v* is ever actually submerged in my water. (In this, its logical behavior resembles that of a subjunctive conditional.) Still, this relation is no genuine counterexample since it fails to meet the analysis on two counts. First, it is not a genuine connection; it is not one of the relations that fix the primary logical, causal, or phenomenal order of the world. To be sure, genuine causal connections do underlie this disposition, but these are connections that hold between particular salt molecules and particular water molecules (or between events involving particular salt molecules and events involving particular water molecules or between conditions involving particular salt molecules and conditions involving particular water molecules). And it is only in virtue of such prior causal connections that *ad hoc* dispositional relations between hunks of salt and conditions hold at all. The second count against this possible counterexample is that the relation holds only between particulars and conditions. But the analysis requires that the relation be able to hold between a particular and a *complex* idea (such as a thought). This reveals a weakness in all causally oriented counterexamples: qualities, connections, and conditions—the determinants of the causal order—are *simple* ideas, but it is *complex* ideas that typically are in the range of intentional relations.<sup>11</sup>

(3) The relation of *speaker meaning*—i.e., the relation holding between speaker  $x$  and complex idea  $y$  such that  $x$  means  $y$  by uttering something—might well be a connection. If it is, then all the requirements of the analysis of intentional connection are met. This is no counterexample, though, since the relation of speaker meaning is an intentional relation.

(4) Consider, finally, the relation of *utterance-token meaning*—i.e., the relation holding between utterance tokens and what they mean. This relation can hold contingently between a particular (namely, an utterance token) and a complex idea (namely, the meaning of the utterance token) independently of the veracity of the idea. Is utterance-token meaning truly a connection? It hardly seems so. An utterance token and the relevant complex idea are not related to one another just on their own; the intervention of a third element is required, namely, the intentional activity of thinking creatures. Not unless these creatures make utterances with appropriate intentions and beliefs do utterance tokens become related to the relevant complex ideas. Intending and believing are the genuine intentional connections; the relation between the utterance token and the complex idea that comes to be its meaning is entirely derivative. Unlike intending and believing, it plays no role in the primary causal and phenomenal order of the world.<sup>12</sup>

I have been unable to find better candidate counterexamples than these, and I have thus grown to be convinced that the analysis is free of all serious counterexamples. The thesis to which I wish to be committed, then, is that every candidate counterexample can be disqualified; at most minor alterations in the analysis might be called for. (E.g., minor adjustments might be called for in order to deal with someone's special doctrine about modality, existence, and time.<sup>13</sup>) Rather than attempting to tie up every loose end, I will stop here in hopes that the reader will be able to see how these adjustments would be made.

Perhaps what is most distinctive about this analysis of intentionality is that it is given entirely in terms of logic, specifically, the logic for qualities and concepts. Now someone might worry that the analysis is at bottom circular, for the theory of concepts is none other than the theory whose purpose it is to treat the logic for intentional matters. But this worry would be unfounded. Logic is logic regardless of its field of application. And the theory of qualities and concepts definitely counts as logic. Its primitive

constants intuitively qualify as logical constants; they certainly are not smuggled in from psychology. (See also note 5 page 251.)

Having completed my analysis of intentionality, I am finally in a position to discuss Brentano's thesis of intentionality, which concerns the nature of the mental in general. Brentano's thesis will be the underlying theme of the next section.

#### 49. Experience and the Mental

According to Brentano's thesis of intentionality, all and only mental phenomena are intentional, and moreover, no purely physical phenomenon is intentional. In a while I will take up the second half of this thesis; for the present my concern will be with the first half. Is it really true that intentionality is the mark of the mental? The counterexample that springs to mind is that of pure, uninterpreted experience—pure sensation or pure inner (emotional) feeling—as posited by traditional empiricists. Any such experience would certainly be a mental phenomenon, but it would not be *about* anything. Brentano of course wants to deny that there is any such thing as pure experience. However, Brentano puts forward the first half of his thesis as analytic. Therefore, this half of the thesis would be undermined if pure experience were merely possible for some beings or other, not necessarily human beings. In the face of this threat, we would be wise to have an analysis of the mental that is neutral with respect to the possibility of pure experience.<sup>14</sup> Since this is the goal, we must first get clearer about the sort of thing pure experience is supposed to be. I will begin by speculating about sensation and later on will take up inner feeling. (Throughout I will use 'sensation' to mean pure sensation and 'inner feeling' to mean pure inner feeling.) I should stress, though, that if pure experience is not possible, then all experience is intentional and no analysis of the mental is required beyond the previous analysis of the intentional.

Just as in the case of judgement, so in the case of sensation there are both relational and non-relational theories. Relational theories of sensation most often assert that sensation consists in standing in the relation *sensing* to special mental particulars, e.g., appearances, sense impressions, phantasms, sense data, *sensa*, etc. So, for example, if I were to gaze at a bright red tomato in normal well-lit conditions, then according to this theory I would come to stand in the sensing relation to a mental particular that is itself a red object. And similarly, if I were to hallucinate a bright red tomato, I would

stand in the same sensing relation to the same sort of mental red particular.

Realists find it incredible that there are any such special mental particulars. They find it incredible, e.g., that, when I hallucinate something red, there is a real object sensed by me that is actually colored red. This reaction of the realists is certainly reasonable. Yet many realists carry it to an overreaction. They go on to conclude that sensation cannot in any way consist in standing in relations to objects and, therefore, that the non-relational theory of sensation must be right. The adverbial theory of sensation is one instance of a non-relational theory that is arrived at by this route. Versions of the adverbial theory have been advocated by Ducasse, Chisholm, Ayer, and, at times, Russell. On Chisholm's version of the adverbial theory, for example, my experience in the above two cases should be described by means of an explicitly adverbial construction, e.g., 'I am appeared to *redly*' or 'I sense *redly*'. All suggestion of the relational theory of sensation is thereby avoided.

Just as there is an analogy between non-relational theories of sensation and non-relational theories of judgement, there is also an analogy between adverbial theories of the language for sense experience and operator (and prefix) theories of the language for intentionality. In the last section I showed that the matter of generality created logical difficulties for non-relational theories of judgement and operator (and prefix) theories of the language for intentionality. It is predictable, therefore, that non-relational theories of sensation and adverbial theories of the language for sense experience should also run into logical difficulties on the issue of generality. For example, on the adverbial theory how are we to express the thought that the sense experience of one creature *u* is exactly like that of another creature *v*? Perhaps one could express this thought by means of some baroque higher-order adverbial theory. E.g., one might attempt to treat adverbs as a special new category of singular terms whose semantical correlates are "ways of experiencing". Accordingly, one might attempt to represent the above thought about creatures *u* and *v* in something like the following manner:

$$(\forall F)(F(\text{Senses}(u)) \equiv F(\text{Senses}(v))).$$

But notice how close this comes to being just a higher-order version of the relational theory of sensation, the theory that the adverbial

theory is designed to avoid. How much more direct it is to express the thought about creatures  $u$  and  $v$  by means of an explicitly relational theory:

$$(\forall x)(u \text{ senses } x \equiv v \text{ senses } x)$$

i.e.,  $u$  senses whatever  $v$  senses, and conversely.<sup>15</sup> For this and related reasons I am inclined to conclude that, if there is such a thing as pure sensation, then the best theory of it is the relational theory.

Once one adopts the relational theory, however, one is obliged to identify the sort of objects that are in the range of the sensing relation. What are they? In order to help find the answer to this question, let us consider again the thought that the sense experience of creature  $u$  is exactly like that of creature  $v$ . Intuitively, this thought could be true. So for the purpose of discussion, let us suppose that it is. Then, given the above relational analysis of the thought, it follows that, in this example at least, the objects of the sensing relation are not private objects.<sup>16</sup> Next, let us suppose that creature  $u$  or creature  $v$  or both  $u$  and  $v$  are hallucinating some or all of the time. Even in this case, it still seems that the thought that the sense experience of  $u$  is exactly like that of  $v$  could be true. So for the purpose of discussion, let us again suppose that it is. Then, given the above relational analysis of the thought, it follows that at least in this example the objects of the sensing relation are not ordinary physical particulars. Therefore, if in this example the objects of the sensing relation were particulars at all, they would have to be some kind of public particular that has no actual location, no actual causal efficacy, etc. This, however, sounds rather like a kind of particular that simply could not exist. It offends virtually every realistic intuition we have.

Barring such unacceptable particulars, one has no alternative but to conclude that at least in this example some objects of the sensing relation are universals, namely, certain appropriate qualities and conditions.<sup>17</sup> Such qualities and conditions are called *sensible qualities* and *sensible conditions*. By allowing at least some of the objects of the sensing relation to be sensible qualities and sensible conditions, one is still able, as desired, to hold a version of the relational theory of sensation. This version of the theory permits the objects of sensation to be public objects, as desired. For qualities and conditions are public objects. And at the



same time, this version of the relational theory avoids the realist objections that are so damaging to the traditional sense-data theory. According to the present version of the relational theory, one can, for example, sense the color red quite independently of whether any particular is actually colored red (as in hallucinations or dreams).<sup>18</sup> Likewise, given that one can sense conditions, one can do so quite independently of whether they actually obtain. For example, given that one can sense the condition that something colored red is surrounded by something colored blue, one can do so quite independently of whether there actually is something colored red surrounded by something colored blue. In general, on this version of the relational theory, the sensing relation is such that it can hold between an individual and a quality or condition independently of whether the quality or condition is *concretized*, i.e., independently of whether it either has an instance or obtains.

Now whenever the sensing relation holds in the way just indicated between some individual and some quality or condition, it of course does so only contingently. In addition, if there is such a thing as pure sensation, it would seem that the sensing relation is a genuine connection, rather than a mere Cambridge relation. But if so, look how close this comes to the analysis of intentional connections. An intentional connection is one that can contingently connect an individual to a complex concept or thought independently of the veracity of the concept or thought. On analogy, the sensing relation is a connection that can contingently connect an individual to a quality or condition independently of whether the quality or condition is concretized. Just as in thinking one engages the possibility of falsehood, so in sensing one engages the possibility of illusion.

With these remarks about sensing in mind, let us now turn to the topic of inner feeling, the kind of feeling traditionally thought to be associated with emotion. If the bodily-sensation theory of inner feeling were correct, then our job would be done, for inner feeling in that case would be just a species of sensing. If, on the other hand, this theory is not correct, then feeling must be taken up separately. From a logical point of view feeling does seem unlike sensing; whereas an individual can sense a sensible quality independently of whether anything has the quality, an individual can feel, say, a pure emotional quality only if he himself has the quality. For example, whereas one can sense red independently of whether anything is

red, one can feel sad only if something—oneself—is sad at least momentarily.<sup>19</sup> This might suggest that the relations of feeling and sensing are unrelated. Perhaps this is so. However, there is an attractive alternative which I will suggest.<sup>20</sup>

According to this alternative, sensing and feeling are not distinct basic modes of pure experience. There is in fact only one basic mode of pure experience, namely, *pure experiencing* itself. Sensing and feeling differ only in their objects. Sensing is the relation that results from restricting the range of the experiencing relation to sensible qualities and sensible conditions, i.e., qualities and conditions that can be experienced independently of whether they are concretized. And feeling is the relation that results from restricting the range of the experiencing relation to reflective qualities, i.e., qualities that an individual can experience only if they are qualities of that individual at least momentarily.<sup>21</sup>

On the view of the mind that is emerging there are two basic types of mental phenomena: pure experience and thinking (where the latter is taken to include all that is intentional).<sup>22</sup> The difference between them is that in pure experience we are typically connected to qualities and conditions and in thinking we are typically connected to complex concepts and thoughts. From a purely logical point of view, the difference between qualities, connections, and conditions, on the one hand, and thoughts and complex concepts, on the other, comes down to one of logical form. The former are simple with regards to logical form; the latter are complex. However, all of these entities, whether simple or complex, are intensional entities, and indeed they are the only intensional entities. That is, all and only these entities are ideas, as intensional entities are called in the theory of qualities and concepts. This suggests a unified analysis of what a mental connection is:

A connection is mental if and only if it is—or is necessarily included in—a connection that can contingently connect an individual to an idea independently of whether that idea is realized.<sup>23</sup>

(An idea is realized *iff* it is a true thought, a concept that applies to something, a quality that has instances, a connection that connects something, or a condition that obtains.)

Intuitively, what is distinctive about a mental being is this. He can stand in contingent connections to ideas, and he can do so in

such a way that it is not crucial how these ideas actually correspond to, or show up in, the world unless these ideas reflect special aspects of his own mental conditions, in which case they are realized in him. In so connecting an individual to ideas, mental connections thus provide him with a highly adaptive kind of "window on the world" whose reliability typically is variable, one exception being when it reflects certain special aspects of the individual's own mental conditions. Although non-mental individuals are connected to ideas in various ways, they are never connected to ideas in these unique ways. And this makes all the difference.

The candidate counterexamples that come to mind are all variants of those facing the analysis of intentionality, and they can be disqualified for corresponding reasons. One kind of candidate counterexample deserves special comment, however: namely, causally grounded dispositional relations. Consider the relation holding between  $x$  and  $y$  such that  $x$  is disposed to be activated by  $y$ —i.e., the relation holding between  $x$  and  $y$  such that  $y$  (dispositionally) activates  $x$ . Let us suppose what might well be false, that this relation is a connection. A critic of the analysis might claim that this relation can contingently connect a particular to a quality independently of whether the quality has any instances. He might claim, for example, that it can contingently connect a particular photoelectric cell to the color red independently of whether anything is actually colored red. If the critic is right, then the relation would be a true counterexample. I would dispute what the critic claims, however. In order for one thing to *dispositionally* activate another, it must be the kind of thing that can *actually* activate that other thing. But what actually activates a photoelectric cell, for example, is not a color itself but rather particular electromagnetic waves, which might be instances of the color. So the color red is not the kind of thing that can dispositionally activate a photoelectric cell, contrary to what the critic claims. True, a photoelectric cell can bear derivative Cambridge relations to a quality by being connected to instances of the quality. But there is no way for a photoelectric cell and a quality to stand in a genuine connection of the sort the critic imagines; the mediation of an instance of the quality is required. What is special about pure experience (if it truly exists) is that in it we are connected to sensible qualities without the need for the mediation of instances of the quality; we can sense red in

hallucinations, illusions, and dreams even if nothing is actually colored red.<sup>24</sup>

This analysis of the concept of a mental connection is stated entirely within the logic for qualities and concepts. Nothing but fundamental logical relations are appealed to: the predication relation, the thought-building operations (conjunction, negation, existential generalization, etc.), and the associated condition-building operations. Furthermore, the analysis, unlike the one envisaged by Brentano, does not rule out the possibility of pure experience. Thus the analysis holds independently of the first half of Brentano's thesis of intentionality, i.e., independently of the conjecture that all and only mental phenomena are intentional. This is fortunate, for there is a barrier blocking a full defense of the first half of Brentano's thesis: how could one ever demonstrate that pure experience is not at least possible for some being or other?

The second half of Brentano's thesis—i.e., the conjecture that no purely physical phenomenon is intentional—fares better than the first half, for there is no comparable barrier to its defense. Let us now look at this half of his thesis in relation to the issue of materialism. In order to state the doctrine of materialism, one must pay attention to the question of what it takes for an object to be physical in the materialists' sense.

Suppose that a particular is connected by some connection to some object. Then for brevity I will say that the particular *has* the connection. For example, every particular stands in the predication relation to the qualities that qualify it and to the concepts that apply to it. (In these cases the predication relation holds sometimes necessarily, sometimes contingently.) Thus, since the predication relation is a connection, every particular has the predication relation as one of its connections. There are also certain other logical connections that every particular has. Some of these the particular has necessarily. The following example involves one such necessary logical connection. Consider the particular  $x$  and an arbitrary property  $|Fy|_y$ , and the associated condition  $|Fx|^x$ . The condition-building operation of predication connects the three objects  $|Fy|_y$ ,  $x$ ,  $|Fx|^x$ , and it does so necessarily. Thus, the particular  $x$  has this necessary logical connection as one of its connections. In much the same vein, a particular also has necessary logical qualities, i.e., logical qualities that necessarily qualify the particular. Thus, every particular has various necessary logical qualities and connections;

in addition, every particular also has the predication relation as a purely logical connection.

Now materialism requires that, besides the purely logical qualities and connections of the sorts just described, the only qualities and connections that any particular has are physical qualities and connections. Particulars that are like this may be called *purely physical*. And those phenomena involving qualities and connections each of which is physical likewise may be called *purely physical*.<sup>25</sup> Materialism then is the doctrine that all particulars are purely physical and, in turn, all phenomena are purely physical.

Materialists historically have been unclear about the meanings of their basic terms; in particular, a fog surrounds the key term 'physical'. I believe that the only way one can remedy this situation is to analyse the basic concepts of materialism (and physicalism) within a purely logical theory; i.e., one must employ a strategy akin to the one I employed in analysing intentionality and the mental. Such an analysis can make use of the following necessary condition. Genuine physical connections can hold only between particulars and particulars, and perhaps between particulars and locations, particulars and times, particulars and stuffs, locations and locations, times and times, stuffs and stuffs, etc.;<sup>26</sup> they cannot hold between particulars and complex ideas. If a genuine connection holds between particulars and complex ideas, then it is not a physical connection. Whoever maintains otherwise would appear to have forgotten a category difference between physical and mental connections.

According to my analysis of intentionality, an intentional connection is one that can contingently connect a particular to a complex idea independently of the idea's veracity. Two conclusions follow. First, no intentional connection is a physical connection. Secondly, no intentional connection is the sort of special logical connection (characterized above) that the materialist permits purely physical particulars to have. From these two conclusions a third follows. For any intentional connection *x*, any particular *y*, and any complex idea *z*, the intentional phenomenon that *x* connects *y* to *z* is not a purely physical phenomenon. And this conclusion can be generalized, yielding the conclusion that no intentional phenomenon is a purely physical phenomenon. This is none other than the second half of Brentano's thesis of intentionality.

The same argument also works for mental connections and

mental phenomena generally. On my analysis, mental connections are necessarily included in connections that can contingently connect a particular to an idea (i.e., to an intensional entity) independently of that idea's being realized. No physical connection is like this. Thus, no mental connection is a physical connection. At the same time, no mental connection is the sort of special logical connection that the materialist permits purely physical particulars to have. It follows that, for any mental connection *x*, any particular *y*, and any idea *z*, the mental phenomenon that *x* connects *y* to *z* is not a purely physical phenomenon. Generalizing, one may conclude that no mental phenomenon is a purely physical phenomenon. This last conclusion is the main consequence that Brentano wanted to derive from his full thesis of intentionality. Fortunately, it is obtained here without appeal to the first half of Brentano's thesis.

From the conclusions reached in the preceding paragraph it follows that, if there are beings who in fact have mental connections to things, then there are beings who are not purely physical and there are phenomena that are not purely physical. Now, *à la* Descartes, I know directly that I am thinking thoughts. I also know directly that I am sensing smells, sounds, etc. and that I am feeling emotions. However, I can be thinking thoughts and having sensations and feelings only if I am standing in mental connections to things. And since I have mental connections to things, I am not purely physical, and phenomena involving my mental connections are not purely physical. Hence, if I am identical to my body, then my body is not purely physical, and phenomena involving my body's mental connections are not purely physical. And if I am not identical to my body, then phenomena involving nothing but my body are not identical to phenomena involving me, phenomena such as my thinking, my sensing, and my feeling. Either way, I am not identical to a purely physical object, and phenomena involving my mental connections are not identical to purely physical phenomena. Finally, all the foregoing conclusions about mental phenomena hold *mutatis mutandis* for mental events. Whether or not I am identical to my body, mental events involving me are not purely physical.

Let us consider now the Chisholm/Sellars dispute on intentionality and the mental which I described briefly in the previous section. In this dispute Chisholm, echoing Brentano, maintained

that no intentional sentence could be analysed without making further appeal to intentional vocabulary, and Sellars, espousing his own special version of materialism, maintained that every intentional sentence could be analysed without appealing to any intentional vocabulary. I will argue that, given my analysis of the concept of an intentional connection, both Chisholm and Sellars were in error. Consider Chisholm's position first. Given my analysis, there are at least some intentional sentences, e.g., 'Someone stands in an intentional connection to some thought', that can be analysed without intentional vocabulary; purely logical vocabulary suffices here. So, given my analysis, Chisholm's position is too strong. In fact, Chisholm's position that no intentional sentence can be analysed without appeal to further intentional vocabulary is nearly inconsistent with his view that the concept of an intentional sentence has a purely logical analysis. Next consider Sellars' doctrine that all intentional sentences can be analysed without appeal to intentional vocabulary. To see what the problem is with Sellars' doctrine, consider an intentional connection that connects a given individual (e.g., me) to a thought (e.g., the thought that I think). Let the 2-place predicate '*T*' express this intentional connection; let the name '*a*' name the individual, and let the name '*b*' name the thought. Now perhaps materialists could *describe* in non-intentional vocabulary the truth expressed by the sentence '*T(a, b)*'. Nevertheless, they could never produce a non-intentional sentence that *expresses* this truth. The reason goes as follows. Sentences, I have shown, express thoughts. However, thoughts conform to conception 2. Therefore, by the laws of the logic for conception 2, it follows that the only sentences that could ever express the same thing as the sentence '*T(a, b)*' are sentences having the logical form '*F(c, d)*' or '*G(c)*', where '*F*' and '*G*' are primitive predicates. But in this case '*F*' would have to express just the intentional relation expressed by the intentional predicate '*T*', and '*G*' would have to express just the complex intentional concept expressed by the intentional open sentence '*T(x, b)*'. So '*F*' and '*G*' would themselves have to be intentional predicates. And so '*F(c, d)*' and '*G(c)*' would have to be intentional sentences. Thus, it is impossible to express in non-intentional vocabulary the truth expressed by the original intentional sentence. It follows, therefore, that the materialist inevitably leaves something out.

I will close my discussion of the mental with a brief comment on minds and machines. Suppose that we should one day design and build a machine that performs physically as we do, both behaviorally and mechanically. A natural question to ask is whether the machine has a mind. According to behaviorism and materialistic versions of functionalism, this is just the question of whether the machine behaves or functions physically as we do. But *ex hypothesi* we already know that it does; *that* is not our question. We want to know something else, namely, whether the machine actually functions mentally. But this is to say, we want to know whether it stands in genuine mental connections to things. For intuitively, a thing functions mentally if and only if it stands in mental connections to things. It is not enough that it should behave or function physically as if it were mentally connected to things. Research on minds and machines that disregards this difference is likely to reach misleading conclusions about the basic nature of the mind.<sup>27</sup>

A closely related question—and one with the greatest moral significance—is whether a machine that behaves and functions physically as we do is conscious, i.e., whether it is aware of anything. Before one tries to settle this question, one should try to get clearer about what consciousness is; i.e., one should try to say what it is for a being to be aware of something. I will attempt to do this in my final, rather speculative section.

#### 50. Consciousness

Suppose for a moment that all mental connections are *conscious connections*; i.e., suppose that, necessarily, whenever an individual is mentally connected to something, he is also conscious of that thing. In this case, since I have already given a purely logical definition of the concept of a mental connection, it would be a straightforward affair to obtain a purely logical definition of consciousness also. Perhaps this is all there is to it. However, if certain commonly held psychological theories are correct, there are mental connections that are not conscious. (Examples might be standing belief, standing desire, unconscious desire, unconscious decision, etc. Let us call such mental connections *non-conscious*.) If there truly are non-conscious mental connections, then some other strategy for defining consciousness is in order. Fortunately, one can sidestep the complicated theoretical issue of whether there are non-conscious mental connections, for there is another way to define



consciousness that is neutral with regard to how that issue is settled. The key to the definition is the *unity of consciousness*.

At a given moment I might be sensing one thing, feeling another, thinking a third, and desiring and deciding still others. Yet despite the fact that several specialized conscious connections are all operating at once, my consciousness is not fragmented into a corresponding variety of exclusive programs on competing channels. I have a unified awareness of all the objects of these specialized conscious connections. What accounts for the unity?

The answer is that the consciousness relation—the relation being conscious of or being aware of—is itself a conscious connection whose unique global operation produces this unity. Thus, if an individual has a pure experience of something, then he will also stand in the consciousness relation to that thing; if an individual believes, wants, or decides something consciously, then again he will stand in the consciousness relation to it, and so on for any specialized conscious connection.<sup>28</sup> The consciousness relation is that mental connection whose operation must have this global character. The problem is to get at this character of the consciousness relation without circularity, i.e., without explicitly alluding (as I just did) to the distinction between conscious and non-conscious mental connections.

To be sure, if every mental connection were a conscious connection, the problem would be easy to solve: consciousness would simply be the maximal mental connection, i.e., the mental connection that is the union of all mental connections.<sup>29</sup> But I am looking for a definition that is compatible with psychological theories that entertain non-conscious mental connections, so another approach is needed.

Notice that the non-conscious mental connections have something in common. In each case, they mimic functionally the operation of a conscious counterpart, and moreover, in each case this conscious counterpart is one of the specialized conscious connections such as conscious belief, conscious want, conscious decision, etc. The consciousness relation itself is not the conscious counterpart of a non-conscious mental connection. What would that non-conscious connection be, unconscious consciousness, unaware awareness? The theoretical purpose served by non-conscious mental connections is to constitute a non-conscious functional analogue of conscious mental processes. But the consciousness

relation has no function that could possibly show up when mental processes are looked at mechanically as if there were no consciousness involved. All it does is to produce a unified consciousness. There cannot be a non-conscious connection with that function since non-conscious mental connections are, as we say, non-conscious. (This is why the consciousness relation never appears in functional psychologies based on the information-processing model.)

The conscious connections with non-conscious counterparts are those that do have mechanically recognizable functions. Conscious belief, conscious want, conscious decision, etc. are like this.<sup>30</sup> There are, however, certain other specialized conscious connections—such as attending, concentrating, meditating, contemplating<sup>31</sup>—that are not like this. Akin to the consciousness relation, they lack any apparent function when mental processes are viewed mechanically as though no consciousness were involved. Non-conscious attending, non-conscious concentrating, non-conscious meditating, non-conscious contemplating—these are useless Cambridge relations if any are. The only immediate function of relations like attending and meditating is to alter the quality of consciousness. (Thus I will call relations of this kind *qualitative conscious relations*.) For example, if an individual attends to or concentrates on something, then he will thereby acquire a *keen* awareness of that thing; if he meditates on or contemplates something, then he will therein have a *heightened* consciousness of it.<sup>32</sup>

Consider now the conscious connections that do have non-conscious counterparts. Notice that these conscious connections have special category limitations on their ranges. For example, it is impossible to believe consciously a color, a taste, or a smell; to want consciously a number; to decide consciously oneself, etc.<sup>33</sup> Naturally, the special category limitations on these conscious connections are inherited by their non-conscious counterparts. So, for example, just as it is impossible to believe consciously a color, a taste, or a smell, it is also impossible to believe non-consciously a color, a taste, or a smell; and so on.

Let us recall, however, that the consciousness relation is a global mental connection in the sense that, if an individual stands in any conscious connection to an object, then he will also be conscious of that object. The consciousness relation thus lacks the special category limitations of the sort in force for the mechani-

cally significant conscious connections and their non-conscious counterparts; in the sense that its range is free of these limitations the consciousness relation is *transcendental*.<sup>34</sup> I can be conscious of colors, sounds, smells, numbers, and I think that I am at the present conscious of myself and the present.<sup>35</sup> The qualitative conscious relations (attending, etc.) are like the consciousness relation in this. They too lack the special category limitations on their ranges, and in this sense they too are transcendental mental relations. So in general, if it is possible for an individual to stand in any conscious or non-conscious mental connection to an object, then it is also possible for an individual to be conscious of that object, and likewise, it is possible for some individual to attend to (concentrate on, contemplate, meditate on) that object. Now what is unique about the consciousness relation is that it is the maximal mental connection whose range has this transcendental character. That is, it is the mental connection that is the union of all transcendental mental connections, including itself. This is what is global about the consciousness relation.

Notice that this captures the unique global character of consciousness without circularity, i.e., without mention of the distinction between conscious and non-conscious mental connections. Thus, I am finally in a position to state my definition:

consciousness =<sub>df</sub> relating by means of transcendental mental connections.

That is,

$|x \text{ is conscious of } y|_{xy} =_{df}$   
 $|(\exists z)(z \text{ is a transcendental mental connection} \ \& \ \langle x, y \rangle \Delta z)|_{xy}.$

Inasmuch as the consciousness relation is the mental connection having this unique global character, its operation is what produces the unity of consciousness.<sup>36</sup>

From the definition of consciousness, definitions of conscious mental quality and connection follow immediately. Now, each new category of qualities and connections generates the possibility of a new category of beings and, thus, the possibility of a new category of associated phenomena. So it is with the category of conscious mental qualities and connections; it too generates the possibility of a new category of beings and associated phenomena. Indeed, that possibility is actual, for we are among the new beings.



