D. C. DENNETT

Brain Writing and Mind Reading

What are we to make of the popular notion that our brains are somehow libraries of our thoughts and beliefs? Is it *in principle* possible that brain scientists might one day know enough about the workings of our brains to be able to "crack the cerebral code" and read our minds? Philosophers have often rather uncritically conceded that it *is* possible in principle, usually in the context of making some point about privacy or subjectivity.[1] I read Anscombe to deny the possibility. In *Intention*[2] she seems to be arguing that the *only* information about a person that can be brought to bear in a determination of his beliefs or intentions is information about his past and future actions and experiences; a person's beliefs and intentions are whatever they must be to render his behavioral biography coherent, and neurological data could not possibly shed light on this. This is often plausible. Suppose Jack Ruby had tried to defend himself in court by claiming he didn't know (or believe) the gun was loaded. Given even the little we know about his biography, could we even make sense of a neurologist who claimed that he had scientific evidence to confirm Ruby's disclaimer? But in other cases the view is implausible. Sometimes one's biography seems completely compatible with two different ascriptions of belief, so that the Anscombean test of biographical coherence yields no answer. Sam the reputable art critic extols, buys, and promotes mediocre paintings by his son. Two different hypotheses are advanced: (a) Sam does not believe the paintings are any good, but out of loyalty and love he does this to help his son, or (b) Sam's love for his son has blinded him to the faults of the paintings, and he actually believes they are good. Presum-

AUTHOR'S NOTE: Earlier drafts of this paper were read at the University of Maine, Tufts University, and the University of Cincinnati Colloquium on Brain and Mind, November 1971.

[1] See, in another context, A. I. Melden's use of the notion in *Free Action* (New York: Humanities, 1961), pp. 211–215.

[2] G. E. M. Anscombe, *Intention* (2nd ed.; Oxford: Blackwell, 1963).

ably if (a) were true Sam would deny it to his grave, so his future biography will look the same in either case, and his past history of big-heartedness, we can suppose, fits both hypotheses equally well. I think many of our intuitions support the view that Sam really and objectively had one belief and not the other, and it goes against the grain to accept the Anscombean position that in the absence of telltale behavioral biography there is simply nowhere else to look. Couldn't the brain scientist (in principle) work out the details of Sam's belief mechanisms, discover the *system* the brain uses to store beliefs, and then, using correlations between brain states and Sam's *manifest* beliefs as his Rosetta Stone, extrapolate to Sam's covert beliefs? Having deciphered the brain writing, he could read Sam's mind. (Of course, if we could establish this practice for Sam the art critic, we would have to reopen the case of Jack Ruby, but perhaps, just perhaps, we could then devise a scenario in which neurologists were able to confirm that Ruby was the victim of a series of unlikely but explainable beliefs — as revealed by his "cerebroscope.")

I admit to finding the brain-writing hypothesis tempting,[3] but suspect that it is not coherent at all. I have been so far unable to concoct a proof that it is incoherent, but will raise instead a series of difficulties that seem insuperable to me. First, though, it would be useful to ask just why the view is plausible at all. Why, for instance, is the brain-writing hypothesis more tempting than the hypothesis that on the lining of one's stomach there is a decipherable record of all the meals one has ever eaten? Gilbert Harman offers the first few steps of an answer:

> We know that people have beliefs and desires, that beliefs and desires influence action, that interaction with the environment can give rise to new beliefs, and that needs and drives can give rise to desires. Adequate psychological theories must reflect this knowledge and add to it. So adequate models must have states that correspond to beliefs, desires and thoughts such that these states function in the model as psychological states function in the person modeled, and such that they are representational in the way psychological states are representational. Where there is such representation, there is a system of representation; and that system may be identified with the inner language in which a person thinks.
>
> This reduces the claim that there is an inner language, which one

[3] I claimed it was a distinct possibility with regard to intentions in "Features of Intentional Actions," *Philosophy and Phenomenological Research*, 29(1968):232–244.

thinks in, to the trivial assertion that psychological states have a representational character.[4]

The first point, then, is that human behavior has proven to be of such a nature that the only satisfactory theories will be those in which inner representations play a role (though not necessarily a role that is not eliminable at another level of theory). Diehard peripheralist behaviorists may still wish to deny this, but that is of concern to historians of science, not us. It is Harman's next point that strikes me as controversial: where there is representation there is system, and this system may be *identified with* a person's inner language. Are all representations bound up in systems? Is any system of representations like a language? Enough like a language to make this identification more useful than misleading? Or is Harman's claim rather that whatever sorts of representations there may be, the sorts we need for human psychology must be organized in a system, and this system must be more like the system of a language than not? Assuming Harman's claim survives these questions, we still would not have an argument for the full-fledged brain-writing hypothesis; two more steps are needed. First, we need the claim that these psychological models with their language-style representations must be realized in brainware, not ectoplasm or other ghostly stuff.[5] This ought to be uncontroversial; though psychologists may ignore the details of realization while elaborating and even testing their models, the model-making is ultimately bound by the restriction that any function proposed in a model must be physiologically or mechanically realizable one way or another. Second, it must be claimed that it will be possible to determine the details of such realizations from an empirical examination of the brainware and its causal role in behavior. This second point raises some interesting questions. Could the functional organization of the brain be so inscrutable from the point of view of the neurophysiologist or other physical scientist that no fixing of the representational role of any part were possible? Could the brain use a system that no outsider could detect? In such a case what would it mean to say the brain used a system? I am not sure how one would go

[4] "Language Learning," *Nous*, 4(1970):35. See also his "Three Levels of Meaning," *Journal of Philosophy*, 65(1968):590–602, esp. p. 598.

[5] Cf. Wilfrid Sellars, "Notes on Intentionality," *Journal of Philosophy*, 61(1964): 663, where he discusses mental acts as tokens expressing propositions, and claims that all tokens must be sorts of tokens and must have a determinate factual character, and proposes identifying them with neurophysiological episodes.

about giving direct answers to these questions, but light can be shed on them, I think, by setting up a crude brain-writing theory and refining it as best we can to meet objections.

Again Harman gives us the first step:

> In a simple model, there might be two places in which representations are stored. Representations of things believed would be stored in one place; representations of things desired in the other. Interaction with the environment would produce new representations that would be stored as beliefs. Needs for food, love, etc., would produce representations to be stored as desires. Inferences could produce changes in both the set of beliefs and the set of desires. (*Ibid.*, p. 34)

No doubt we would also want to distinguish more or less permanent storage (belief and desire) from the more fleeting or occurrent *display* of representations (in perception, during problem solving, sudden thoughts, etc.). In any case we already have enough to set some conditions on the brain-writing hypothesis. Some formulations of it are forbidden us on pain of triviality. For instance, claiming that there is brain writing, but that each representation is written in a different language, is just an oblique way of asserting that there is no brain writing. I think the following six conditions will serve to distinguish genuine brain-writing hypotheses from masqueraders.

(1) The system of representations must have a generative grammar. That is, the system must be such that if you understand the system and know the finite vocabulary you can generate the representations — the sentences of brain writing — you haven't yet examined. Otherwise the language will be unlearnable.[6] Only if there were a generative grammar could the investigator get himself into a position to extrapolate from manifest beliefs and desires to covert beliefs and desires. There need not be a single generative grammar covering all representations, however. Just so long as there is a finite number of different "languages" and "multilingual" functional elements to serve as interpreters, the learnability condition will be met.

(2) Syntactical differences and similarities of the language must be reflected in physical differences and similarities in the brain. That is, the tokens of a syntactical type must be physically distinguishable by finite test from the tokens of other syntactical types. That does not

[6] See Donald Davidson, "Theories of Meaning and Learnable Languages," in Y. Bar-Hillel, ed., *Logic, Methodology, and Philosophy of Science* (Amsterdam: North-Holland, 1965), pp. 383–394.

mean that all tokens of a type must be physically similar. What physical feature is peculiar to spoken and written tokens of the word "cat"? There must simply be a finite number of physical sorts of token of each type. Tokens and "strings" of tokens may of course align themselves in physical dimensions other than those of natural language. For instance, lexical items might be individuated not by shape but by spatial location, and ordering in the strings might be accomplished not by a sequence in space or time but by degree of electric potential.

(3) Tokens must be physically salient. This is a "practical" point. Tokens might bear physical similarities, but similarities so complex, so diffuse and multidimensional that no general detection mechanism could be devised; no frequency filters, stereo-locators, litmus papers, or simple combination of these could be built into a token detector. If tokens turned out not to be physically salient — and this is rather plausible in the light of current research — the brain-writing hypothesis would fail for the relatively humdrum reason that brain writing was illegible. It is worth mentioning only to distinguish it from more important obstacles to the hypothesis.

(4) The representation store must meet Anscombe's condition of biographical coherence. The sentences yielded by our neurocryptographer's radical translation must match well with the subject's manifest beliefs and desires, and with common knowledge. If too many unlikely beliefs or obvious untruths appear in the belief store, we will decide that we have hit upon something strange and marvelous — like finding the Lord's Prayer written in freckles on a man's back — but not his belief store. To give a more plausible example, we might discover that certain features of brain activity could be interpreted as a code yielding detailed and accurate information about the relative tensions of the eye muscles, the orientation of the eyeball, the convexity of the lens, etc., and this might give us great insight into the way the brain controlled the perceptual process, but since a man does not ordinarily have any beliefs about these internal matters, this would not be, except indirectly, a key to his belief store.[7]

[7] Discovering such a code is not establishing that the information the code carries for the scientist is also carried for the person or even for his brain. D. H. Perkel and T. H. Bullock, in "Neural Coding" (in F. Schmitt, T. Melnechuk, et al., eds., *Neurosciences Research Symposium Summaries*, vol. 3 [Cambridge, Mass.: M.I.T. Press, 1969]), discuss the discovery of a code "carrying" phasic information about wing position in the locust; it is accurately coded, but the "insect apparently makes

(5) There must be a reader or playback mechanism. It must be demonstrated that the physical system in which the brain writing is accomplished is functionally connected in the appropriate ways to the causes of bodily action, and so forth. Of course, if we were to find the cortex written all over with sentences expressing the subject's manifest beliefs, we would be convinced this was no coincidence, but until the operation of the mechanisms that utilized the writing was discovered we would not have a theory. (A person who discovered such a marvel would be roughly in the same evidential position as a clairvoyant, who (we can imagine) might be able to predict with uncanny accuracy what a person would say, etc., and yet could not be supposed to have any authority — in a court of law, for instance — about a person's beliefs.)

(6) The belief store must be — in the main — consistent. If our translation manual yields sentences like "My brother is an only child" and pairs of sentences like "All dogs are vicious" and "My dog is sweet-tempered" one of several things must be wrong. If the subject declines to assert or assent to these anomalous sentences we will discredit the translation manual (cf. Quine on radical translation); if the man does issue forth with these sentences we will conclude that we have discovered a pathological condition, and our brain-writing system will be viewed as a sort of assent-inducing tumor.[8]

A more graphic way of looking at this point is to ask whether the neurocryptographer could do a bit of tinkering and thereby *insert* a belief in his subject. That is, if he can *read* brain writing he ought to be able to *write* brain writing. Let us suppose we are going to insert in Tom the false belief: "I have an older brother living in Cleveland." Now can the neurocryptographer translate this into brain writing and do a bit of rewiring? Let us suppose he can do any rewiring, as much and as delicate as you wish. This rewiring will either impair Tom's basic rationality or not. Consider the two outcomes. Tom is sitting in a bar and a friend asks "Do you have any brothers or sisters?" Tom says, "Yes, I have an older brother living in Cleveland." "What's his name?" Now what is

no use of this information." (Blocking this input and substituting random input produces no loss of flying rhythm, ability, etc.)

[8] This condition of rationality has some slack in it. We do permit some small level of inconsistency, but large-scale illogicality must be indicative of either a defect in the subject so serious as to disqualify him as a believer at all, or a defect in our translation hypotheses. See my "Intentional Systems," *Journal of Philosophy*, 68(1971):87–106.

going to happen? Tom may say "Name? Whose name? Oh, my gosh, what was I saying? I don't have an older brother!" Or he may say, "I don't know his name," and when pressed he will deny all knowledge of this brother, assert things like "I am an only child and have an older brother living in Cleveland." In neither case has our neurocryptographer succeeded in wiring in a new belief. This does not show that wiring in beliefs is impossible, or that brain writing is impossible, but just that one could only wire in one belief by wiring in (indefinitely?) many other cohering beliefs so that neither biographical nor logical coherence is lost.[9]

Now suppose we have a brain-writing theory that meets all of these conditions: we have a storage facility functionally tied to behavior that is somehow administered to preserve logical and biographical coherence, and the mode of storage involves elements having physically salient syntactical parts for which we have a generative grammar. This system is going to take up some room. How much room do we need? Marvin Minsky has an optimistic answer: "One can't find a hundred things that he knows a thousand things about. . . . I therefore feel that a machine will quite critically need to acquire the order of a hundred thousand elements of knowledge in order to behave with reasonable sensibility in ordinary situations. A million, if properly organized, should be enough for a very great intelligence."[10] If Minsky's estimate were realistic, the brain, with its ten billion neurons or trillions of molecules would be up to the task, no doubt. But surely his figure is much too low. For in addition to all the relatively *difficult* facts I have mastered, such as that New York is larger than Boston and salt is sodium chloride, there are all the easy ones we tend to overlook, like New York is not on the moon, or in Venezuela, salt is not sugar, or green, or oily, salt is good on potatoes, on eggs, tweed coats are not made of salt, a grain of salt is smaller than an elephant. . . . Surely I can think of more than a thousand things I know or believe about salt, and salt is not one of a hundred, but one

[9] I examine this case more fully in "Mechanism and Responsibility," in Ted Honderich, ed., *Essays on Freedom of Action* (London: Routledge and Kegan Paul, 1973), pp. 157–184. Joan Straumanis has pointed out to me that there is some experimental evidence that suggests that another outcome of the rewiring experiment could be that Tom spontaneously and unconsciously fabricates a web of cohering beliefs to "protect" the inserted belief and his others from each other (a sort of pearl-in-the-oyster effect).

[10] M. L. Minsky, ed., *Semantic Information Processing* (Cambridge, Mass.: M.I.T. Press, 1968), p. 26.

of thousands upon thousands of things I can do this with. Then there is my knowledge of arithmetic; two plus two is four, twenty plus twenty is forty. . . . My beliefs are apparently infinite, which means their storage, however miniaturized, will take up more room than there is in the brain. The objection, of course, seems to point to its own solution: it must be that I *potentially* believe indefinitely many things, but I generate all but, say, Minsky's hundred thousand by the activity of an extrapolator-deducer mechanism attached to the core library. So let us attach such a mechanism to our model and see what it looks like.

It has the capacity to extract axioms from the core when the situation demands it and deduce further consequences. If it is to do this, it will need to have an information store of its own, containing information about what items it would be appropriate at any time to retrieve from the core, and, for instance, the metalinguistic information it needs to analyze the contradiction in "all cats are black" and "my cat is brown." Now perhaps it does this by storing the information that what is black is not brown, or maybe *that* information is in the core storage, and the metalinguistic information stored in the extrapolator-deducer mechanism is to the effect that the core element, "what is black is not brown," is relevant to an analysis of the contradiction. Now how will the extrapolator-deducer mechanism store its information? In its own core library of brain-writing sentences? If it has a core library, it will also need an extrapolator-deducer mechanism to act as librarian, and what of *its* information store? Recalling Lewis Carroll's argument in "What the Tortoise Said to Achilles,"[11] we can see that the extrapolator-deducer will be hamstrung by a vicious regress if it must always rely on linguistically stored beliefs which it must retrieve and analyze about what can be deduced from what. This a priori point has been "empirically discovered" in the field by more than one frustrated model-builder. As one team sums it up: ". . . a memory that merely stores *propositions* leads to technological, or organic, monstrosities and frustrates, rather than facilitates inductive operations."[12]

The conclusion is that writing — for instance, brain writing — is a

[11] *Mind* (1895), reprinted in I. M. Copi and J. A. Gould, eds., *Readings on Logic* (New York: Macmillan, 1964). Harman offers a similar argument in "Psychological Aspects of the Theory of Syntax," *Journal of Philosophy*, 64(1967):75–87.

[12] H. von Foerster, A. Inselberg, and P. Weston, "Memory and Inductive Inference," in H. L. Oestreicher and D. R. Moore, eds., *Cybernetic Problems in Bionics* (New York: Gordon and Breach, 1968).

dependent form of information storage. The brain must store at least some of its information in a manner not capturable by a brain-writing model. Could it do without brain writing altogether? I think we can get closer to an answer to this by further refining our basic model of belief.

Representations apparently play roles at many different levels in the operation of the brain. I have already mentioned the possibility of codes representing information about the tension of eye muscles and so forth, and these representations do not fall into the class of our beliefs. At another level there is the information "we use" to accomplish depth perception. Psychophysicists ascribe to us such activities as *analyzing* depth cues and *arriving at conclusions* about distance based on *information we have* about texture gradients, binocular interaction, and so forth. Yet it is nothing conscious that I do in order to perceive depth, and if you ask me what beliefs I have about texture gradients I draw a blank. Closer to home, a child can demonstrate his understanding of addition by reeling off sums without being able to formulate or understand propositions about the commutativity of addition. His performance indicates that he has caught on to commutativity, but should we say that among his beliefs is the belief that addition is commutative? To give one more case, while driving down a familiar road I suddenly am struck by the thought that its aspect has changed — somebody has painted his shutters or a tree has blown down, or something. Do I have a *belief* about how it used to be that grounds my current judgment that it has changed? If so, it is a belief to which I can give no expression and about which I am quite in the dark. Somehow, though, the information is there to be used.

Suppose we partition our information store into the part that is verbally retrievable and the part that is not. I would not want to claim that this separates our beliefs from everything else. Far from it. Our preanalytical notion of belief would permit young children and dumb animals to have beliefs, which must be verbally irretrievable. Perhaps, though, a strong case can be made out that at least our verbally retrievable beliefs are stored in brain writing. The picture that emerges is not, I think, implausible: there are on the one hand those representations that are available for our conscious, personal use and apprehension, and on the other hand those that operate behind the scenes to keep us together. If any representations are stored in brain writing, the former

will be, for they are in intimate relation to our natural languages. Included in this group will be the bits of factual knowledge we pick up by asking questions and reading books, especially the facts that only language-users could apprehend, such as the fact that Thanksgiving is always on Thursday. With regard to this group of representations Minsky's figure of a hundred thousand looks more realistic, provided we have an extrapolator-deducer mechanism.[13]

If ever it seems that we are *storing sentences*, it is when we are picking up facts in this verbal manner, but are these things we pick up our beliefs? Sometimes we salt away a sentence because we like the sound of it, or because we will later be rewarded for producing it on demand, or just because it has a sort of staying power in our imagination. In Chekhov's *Three Sisters*, Tchebutykin, reading a journal, mutters: "Balzac was married in Berditchev," and then repeats it, saying he must make a note of it. Then Irina dreamily repeats it: "Balzac was married in Berditchev." Did they acquire a belief on that occasion? Whether *they* did or not, the sentence has stuck in my mind, and yet I wouldn't say it was one of my beliefs. I deliberately have not looked it up in the encyclopedia; probably it's true — why would Chekhov insert a distracting falsehood, for mischief? No doubt if someone offered me a thousand dollars if I could tell him where Balzac was married, I'd say Berditchev (wherever that is), but it would be wrong for him to conclude that this was a belief of mine.

If brain writing served only for such storage of words and sentences that we pick up for various reasons, at least we could all breathe a lot easier about the prospects of evil scientists reading our every seditious thought and turning us over to the authorities. Imagine the Loyalty Commissar asking the neurocryptographer if the man in the cerebroscope is a true patriot. "Let's see," says the scientist, "Here is the sentence we've been looking for: 'I pledge allegiance to the flag . . .' Would finding the sentence, "America's the greatest land of all," satisfy the Commissar? I think not.

The matter of verbally retrievable beliefs is in any case more complicated than the picture we've just been examining. Whereas if I am

[13] Minsky's anthology, *Semantic Information Processing*, is a collection of several brilliant attempts to provide working models for just such question-answering, extrapolating systems. It is a gold mine of philosophically tantalizing suggestions and problems.

asked who won the Super Bowl in 1969 it does seem a bit as if I am searching for a ready-made sentence to utter in response, in other cases this sort of account does not ring true at all. Suppose I am watching the shell game, intent on which shell the little pea is under. At any moment it seems to be true that I have a belief about where the pea is, and can tell you if you ask, but it does not seem plausible that this is accomplished by a rapid writing and erasing of successive sentences: "now it's left, now it's center, now right" and the flashing on and off of the negation sign in front of "it's under the center shell." For one thing, if asked to give you my perceptual beliefs of a moment I may have to work a bit to formulate them, yet the perceptual representation was what it was before I was asked. The representationality — or intentionality — of something (e.g., a belief or perception) is compatible with its being vague or indeterminate in some respects.[14] The effort of retrieval is often an effort to formulate a sentence that is an approximation of a belief, and we are often distressed at the hard edge of determinacy our verbal output substitutes for the fuzziness of our convictions.

The answer we formulate, the judgment we find an expression for when asked for our belief, is determinate and individuated, because it consists of a specific string of words in our natural language, whether we then speak it aloud or not. These representations, not the beliefs to which we have verbal access but the occurrent, datable judgments themselves, have the syntactic parts we have been looking for, and about these the brain-writing hypothesis looks much more workable. Not only are judgments determinate; they are, as Harman has pointed out, lexically and syntactically unambiguous.[15] If it occurs to me that our mothers bore us, I know for sure whether I am thinking of birth or ennui. So it is proper to view a judgment not as a sentence simpliciter but as a deep structure or sentence under an analysis. Judgments, unlike beliefs, occur one at a time; we have at any moment indefinitely many beliefs, but can be thinking just one thought. We saw that the brain-writing hypothesis with regard to storage of beliefs did not really effect any economies of design, because however systematic and efficient one's grammar is, one still needs infinite space to store infinitely many tokens, but with

[14] See G. E. M. Anscombe, "The Intentionality of Sensation: A Grammatical Feature," in R. J. Butler, ed., *Analytical Philosophy, Second Series* (Oxford: Blackwell, 1965).

[15] "Language Learning."

regard to representation of judgments the situation is different. A finite mechanism incorporating a generative grammar would be an efficient means of representing, one at a time, any of an infinite set of propositions.

The interesting thing about judgments is that although each of us is authoritative about the *content* of his judgments and although each of us is authoritative about his sincerity or lack of sincerity in giving outward verbal expression of a judgment, we are not in a privileged position when it comes to the question of whether our judgments are reliable indicators of our beliefs.[16] Normally there is harmony between our judgments and our behavior, and hence between our judgments and our beliefs, but when we are afflicted by Sartre's *mauvaise foi*, our sincerest judgments can be lies about our beliefs. I may judge to myself that a man is innocent, while believing him guilty.

This suggests that even if we were to discover a brain-writing system that represented our judgments, the mind reading that could be accomplished by exploiting the discovery would not uncover our beliefs. To return to the case of Sam the art critic, if our neurocryptographer were able to determine that Sam's last *judgment* on his deathbed was "My consolation is that I fathered a great artist," we could still hold that the issue between the warring hypotheses was undecided, for this judgment may be a self-deception. But at this point I think we are entitled to question the intuition that inspired the search for brain writing in the first place. If discovering a man's judgments still left the matter of belief ascription undecided, and if in fact either ascription of belief will account for, explain, predict Sam's behavior as well as the other, are we so sure that Sam determinately had one belief or the other? Are we sure there is a difference between his really and truly believing his son is a good artist, and his deceiving himself out of love while knowing the truth in his heart of hearts? If there were brain writing, of course, there would have to be a physical difference between these two cases, but now, what reasons do we have for supposing there is brain writing?

We are thrown back on our conviction that the brain must be an organ that *represents*, but I hope it is no longer obvious that the brain must represent in sentences. In fact we know that at least some of the

[16] See John Vickers, "Judgment and Belief," in K. Lambert, ed., *The Logical Way of Doing Things* (New Haven, Conn.: Yale University Press, 1969), and A. W. Collins, "Unconscious Belief," *Journal of Philosophy*, 66(1969):667–680.

representation must be accomplished in some more fundamental way. Must there be a *system* for such representation? I cannot yet see that there must. In particular, I cannot yet see that there must be a *learnable* system, in Davidson's sense, for it is not clear to me that the brain must — or can — learn (the way a child learns a language) its own ways of representing. Certainly information can be transmitted by means of unlearnable languages. Consider a string of nine light bulbs in a row; there are 512 different patterns of lit and unlit bulbs possible in this array, and so we can use the array to transmit the natural numbers from 0 to 511. There are all sorts of systems one might use to assign patterns to numbers. An obvious and efficient one would be binary notations: 000000001 is 1, 000000010 is 2, and so forth. Once a person knows the system he can generate the numbers he hasn't yet observed, but suppose instead of using this or any other system, the patterns are assigned numbers randomly; a random assignment will carry the information just as well; the outsider will simply not be in a position to predict the balance of the assignments, having learned some of them. Can the brain "use" information carried by unlearnable systems? At some levels of "coding" it obviously does — where the "codes" carry very specific and limited information. In general can the brain get along without learnable representation systems? Until we can say a lot more about what it is for a system to use representations, I for one cannot see how to answer this question. If the answer is no, then there must be brain writing, but how it overcomes the difficulties I have raised here is beyond me. If the answer is yes, then the only way translation of representation can be accomplished is "sentence by sentence," assigning meaning to representations by determining their functional role in the behavior of the whole system. But where competing translations are behaviorally indistinguishable, the content of the representation will be indeterminate.