

## COMMENTS AND CRITICISM

## MULTIPLE MEANINGS AND THE STABILITY OF CONTENT

**E**ric Lormand<sup>1</sup> provides meaning holists with a bold new response to what many people, including Lormand, take to be the main objection to meaning holism: that meaning holism leads to meaning instability. His proposal is particularly interesting because he claims not only to meet the challenge that meaning stability poses, but to transform this challenge into an argument in favor of meaning holism. "My positive aim is to show how meaning holism, even at its most extreme, can accommodate and also increase meaning stability" (51). The innovation responsible for this feat is, according to Lormand, the idea that concepts have multiple meanings. If he is right about this, and he actually succeeds in making a virtue out of what has been taken to be the main difficulty for meaning holism, then a good deal of current work in the philosophy of mind and language is deeply misguided. We shall argue, however, that Lormand does not in fact succeed because his account is inadequate in dealing with the threat that meaning stability poses for meaning holism. In the end, the intriguing hypothesis that symbols have multiple meanings has little effect on the issue of whether meanings are stable.

Meaning holism, like many -isms, is a family-resemblance notion; it covers a variety of related views about the nature of representation. A core idea in this family is that the content of a given type of symbol in a system of representation depends upon its relation to most other types of symbols in the system. This means, for example, that the English word 'bird' has the content it does not just because of its relation to some small number of prescribed words (say, 'animal', 'flies', 'wings') but also because of its relation to most other words of English ('pencil', 'tunnels', 'parka', 'galaxy', and so on). In spite of the counterintuitiveness of this result, meaning holism has many adherents in philosophy and cognitive science. Indeed, in contemporary philosophy of language, it has become something of a characteristic doctrine. The reason for this, as Lormand points out, is that it has proved exceedingly difficult to draw a line between those links among symbols which contribute to their meaning and those which do not. Assuming that some links among symbols contribute to their mean-

<sup>1</sup> "How to Be a Meaning Holist," this JOURNAL, XCIII, 2 (February 1996): 51-73.

ing, the only principled view seems to be that *all* of them do; hence the pressure toward meaning holism.

Meaning holism is connected in a fairly direct way to worries about the stability of content. By the *stability of content*, we mean the fact that two symbols can express the same content even though they fail to bear the same relations to other symbols; or, to put it differently, that the content of a symbol remains the same in spite of changes in its relations to other symbols.<sup>2</sup>

The point is easiest to see in reference to mental content. Suppose for the moment that thinking takes place in an internal system of representation, where beliefs and other propositional attitudes are understood as relations to structured mental representations. On this view, thoughts are much like sentences in that they are composed of simpler elements that themselves have semantic content. Call these constituents *concepts*.<sup>3</sup> Now, when meaning holism is applied to mental content, this means that the content of a concept depends upon its relation to most other concepts in the internal system of representation. Usually, the relations that are thought to matter are a concept's inferential relations, that is, the causal relations it bears to other representations in cognitive and perceptual processes.<sup>4</sup> But then, to the extent that people have different beliefs or different inferential tendencies, the inferential relations in which their concepts participate will diverge. As Lormand points out, these divergences make it extremely difficult to characterize interpersonal—or even intrapersonal—agreement and disagreement. The situation is especially graphic on the strong version of meaning holism, which Lormand adopts for expository purposes. On this version, *any* change in belief involving a given concept changes the meaning of that concept, and since the conceptual system is taken to be thoroughly interconnected, any change in the meaning of one concept changes the meaning of every other concept (55). The result would seem to be that no two people can agree or disagree with

<sup>2</sup>Notice that the first formulation concerns two token symbols and whether they are of the same content type, whereas the second concerns a symbol type (where the type is determined by the content) and conditions on instantiating that type. We shall generally let context make clear whether types or tokens are at issue.

<sup>3</sup>Following Lormand, we shall use bracket notation to refer to mental representations, where the content of the representation is given by the bracketed expression. For example, [bird] is the concept bird, and [some birds fly] is a representation with the content that some birds fly.

<sup>4</sup>See, for example, Ned Block, "An Advertisement for a Semantics for Psychology," in Peter A. French, Theodore E. Uehling, and Howard K. Wettstein, eds., *Midwest Studies in Philosophy, Volume X: Studies in the Philosophy of Mind* (Minneapolis: Minnesota UP, 1986), pp. 615-78.

one another. Similarly, since any individual's beliefs are constantly changing, her concepts must be changing. So, it looks like, a single person could not even agree with her former or future selves. Taking this reasoning to its limit, Lormand points out that a single person could not even execute a simple deduction without equivocation, for in accepting the premises she would thereby change the meaning of the concepts involved (56).

These are certainly bracing consequences, and meaning holism would clearly be a far more palatable view if only these consequences could be avoided. In this context, Lormand's defense of meaning holism is especially welcome. Rather than ignore the problem of the stability of content (like most meaning holists), he argues that, contrary to appearances, meaning holism is actually perfectly compatible with the stability of content. On his account, the innovation that is responsible for the compatibility is the idea that a token symbol has multiple meanings.

Lormand introduces his characterization of multiple meanings by asking us to imagine all of the potential causal/inferential links in which a given concept might participate. In general, these are to be divided into (possibly overlapping) classes, each of which Lormand calls a *unit*. He suggests that, to a first approximation, we think of a unit "as a separable rough test for the acceptable use of that representation" (57). In Lormand's example, one unit associated with [bird] might be [feathered flying animal], another [thing called 'a bird' by mommy] and a third [thing that is similar enough to  $B_1, B_2, \dots, B_n$ ] (where  $B_1, B_2, \dots, B_n$  are things taken to be birds). Lormand's proposal about multiple meanings amounts to the claim that each unit associated with a representation constitutes one of its meanings and that each representation may have as many meanings as it has units. We are supposed to think of *the meaning* of a concept as the collection of all of its individual meanings.

It is this innovation that leads Lormand to claim that meaning holism allows for meaning stability. "Two representations may have different total sets of units but still share meaning, if the sets overlap, or share at least one unit. Even after a change in one unit, a representation still has many of the same meanings it did before" (59). The suggestion, in other words, is to allow that the meaning of every concept changes with any change in belief, but to insist that this requires only that at least one meaning of each concept changes.

Lormand illustrates this proposal with an example involving two children who have rather different cognitive lives. Larry, we are told,

has a mental representation [bird],<sup>5</sup> which is associated with the units [feathered flying animal], [thing called 'a bird' by mommy], and [thing that is similar enough to B<sub>1</sub>, B<sub>2</sub>,...,B<sub>n</sub>]. His brother, Moe, who perhaps is blind, also has a mental representation [bird] which is associated with [thing called 'a bird' by mommy] but not with [feathered flying animal] and not with [thing that is similar enough to B<sub>1</sub>, B<sub>2</sub>,...,B<sub>n</sub>]. Lormand's suggestion is that, despite their differences, the two children are capable of genuine agreement and disagreement precisely because they associate at least one of the same units with [bird]—namely, [thing called 'a bird' by mommy]. This really is the crucial point for Lormand: shared meanings, or units, are the basis for a certain amount of stability.<sup>6</sup>

It is important to see, however, that the proposal does not really establish stability in *the* meaning of concept. At best, it establishes stability in some of the meanings of a concept. Moreover, there is no simple relation between two concepts sharing meanings and their having the same meaning. It might turn out, for example, that two people share only one meaning for [bird] (as in the case above), but that they share several meanings for concepts that, intuitively, are not the same. It might happen, for example, that Larry erroneously believes that birds sometimes make croaking sounds, and that Susie is afraid of them, while Moe correctly believes that Susie is afraid of frogs and also believes that frogs sometimes make croaking sounds. In this case, Larry's [bird] shares more meanings with Moe's [frog] than it does with Moe's [bird].<sup>7</sup> Having "the same meaning," then, is

<sup>5</sup>Here, and below, we speak loosely of his "mental representation [bird]." Strictly speaking, we should use some more neutral locution, since the issue of meaning stability just is whether or not Larry and Moe both possess the same mental representation: [bird]. Perhaps the best strategy would be to construe claims about their "[bird] mental representations" as references to mental representations that are individuated nonsemantically.

<sup>6</sup>The nature of the agreement in such a case is rendered more explicit following Lormand's related suggestion that beliefs involving a concept be understood as determining sets of more specific beliefs, one for each unit associated with the concept. So if Larry thinks that Tweetie is a bird, this really amounts to his having a set of beliefs which includes the belief that Tweetie is a feathered flying animal, the belief that Tweetie is a thing called 'a bird' by mommy, and the belief that Tweetie is similar enough to B<sub>1</sub>, B<sub>2</sub>,...,B<sub>n</sub>. Since Moe also associates [thing called 'a bird' by mommy] with [bird], if he thinks that Tweetie is a bird, he too believes that Tweetie is a thing called 'a bird' by mommy. Hence their agreement about Tweetie consists in their having identical beliefs that Tweetie is a thing called 'a bird' by mommy.

<sup>7</sup>It is particularly important to remember here that we are using the symbolism for referring to concepts in a loose way, as noted earlier in footnote 5. In particular, one cannot assume on the theory under consideration that Moe's [frog] and Larry's [bird] are less similar in meaning than Moe's [bird] and Larry's [bird].

not a simple function of sharing meaning units.<sup>8</sup> Nonetheless, Lormand might be prepared to live without a richer notion of sameness of meaning. He might claim that he only needs meaning units to be stable, since, with the stability of these in place, he can make sense of agreement and disagreement across individuals and over time. Larry and Moe agree that Tweetie is a thing called ‘a bird’ by mommy, and they disagree about whether Tweetie is a feathered flying animal. Allegedly, standard versions of meaning holism cannot even provide this, so Lormand’s proposal may look like a big step forward.

What is more, Lormand does not stop at the claim that multiple meanings secure some amount of stability. He adds that the multiple-meaning view secures more stability than purely referential theories of content. His argument has two stages. First, he claims that so long as “a representation has a long list of meanings anyway (that is, those yielded by its units), we can simply add referents to this list of meanings” (59). “Such a theory,” he says, “would at least provide all the meaning stability of a referential theory, since every meaning posited by the referential theory would also be posited by the holistic [theory]...” (59). Second, he adds that the holistic theory that incorporates referents is superior to a purely referential theory because the holistic theory “provide[s] all the psychologically-relevant meanings that purely referential theories omit” (59-60). For Lormand, psychologically-relevant meanings are more fine-grained than referents. This is because psychology is supposed to require a taxonomy of mental representations which is responsive to interesting psychological generalizations, yet coreferential representations often have different psychological consequences. Lormand’s proposal is to treat the meaning units (for example [is called ‘a bird’ by Mommy]) as providing fine-grained meanings for concepts. According to Lormand, the resulting holistic theory is superior to a purely referential theory because shared meaning units establish points of stability over and above shared referential meanings.

In Lormand’s hands, then, holism looks promising. The multiple-meanings version of holism appears to allow for a considerable amount of content stability—indeed, more stability than its nonholistic competitors. But appearances are deceiving. Lormand’s proposal does not really give us what it promises. What it does, in effect, is trade the issue of the stability of the meaning of a concept for the issue of the stability of the units which constitute the various meanings

<sup>8</sup>These issues are complicated considerably if, as Lormand suggests, we allow “weak” or “nondiagnostic” tests (for example, [birds are smaller than the Milky Way]) to count as individual meanings (58).

of the concept. The problem with this strategy, however, is that, given a strong form of holism, *the contents of the units for a given concept are no more stable than the contents of the concepts themselves.*

Consider the following simple case. Suppose that Larry acquires a single new belief that changes the meaning of one of the units in his concept [bird]. Perhaps he acquires a new belief about feathers, which changes the meaning of his concept [feather], and so changes the meaning of the unit [feathered flying animal] in his concept [bird]. Can this change be isolated? It looks like it cannot. Since one of the units in [bird] has changed, the meaning of [bird] changes, so, if Larry also believes that birds are animals, then the meaning of [animal] changes as well. Since [animal] changes, [human] changes. Since [human] changes, [mommy] changes. Since [mommy] changes, [thing called 'a bird' by mommy] changes. If Larry believes that [mommy's hair is similar to mine], or has some such belief, then changing the meaning of [mommy] also changes the meaning of [similar] and thereby changes the meaning of [things similar enough to B1, B2,...,Bn]. (See figure 1.) As this simple chain of changes illustrates, the holist cannot insulate changes in meaning. Since the units themselves are representations composed of other concepts, these "unit concepts" will themselves be linked to concepts that are more immediately affected by the initial change, and thus they, too, will change meanings.

The upshot is that, on a thoroughgoing holism, changing the meaning of even a single unit will involve a change in the meaning of just about every other concept in the representational system, including all of the other concepts that are labeled as "units."<sup>9</sup> The multiple-meanings hypothesis has done nothing at all to diminish the threat of meaning instability. Before, when concepts were understood to have just a single meaning each, we faced the problem that a single change in belief would result in every concept changing its content. Now, with the multiple-meaning hypothesis, the problem is that with a single change in belief, every concept changes *all of its*

<sup>9</sup>The specific beliefs used in this example are purely illustrative. All that is needed is that all concepts are linked to each other somehow or other. The same point could perhaps be made more simply by noting that any given concept is likely to show up in the tests for nearly every other concept. As Lormand says: "A complete specification of the tests for a typical conceptual representation is likely to include virtually all of one's other representations" (58). So changing the meaning of a single concept will fairly immediately change the meanings of all the concepts to which it is connected. Since these concepts include those used in the units of a representational system, the change will lead to changes in the meanings of all the units associated with all of the concepts of the representational system.

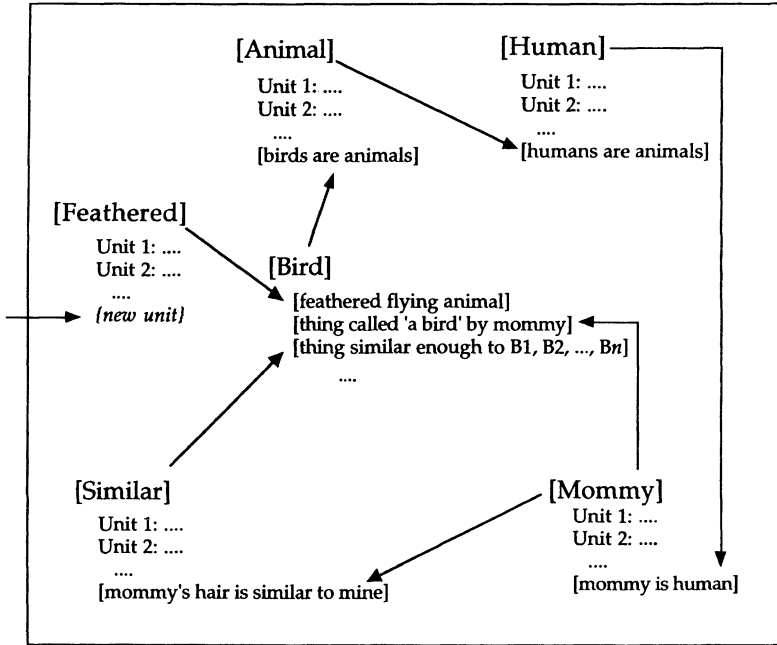


Figure 1. An illustration of how a change in the meaning of a single unit of a typical concept—[bird]—leads to changes in the meanings of all the units of that concept. Arrows indicate linked changes in meaning.

meanings. Either way, instability appears to follow from the commitment to holism.

As far as we can see, the only way that a multiple-meaning approach could contribute to stability would be if the semantic theory had two parts, one of which was not holistic. Lormand makes a suggestion apparently along these lines, in a footnote. He remarks that

...the meanings of all the representations in a system can be specified completely in terms of a relatively small set of semantically atomic (or simple, primitive, basic) representations in the system. A representation can have a primitive meaning—and so be a semantic atom while also having other nonprimitive meanings—and so satisfy inferential-role semantics and meaning holism. Units consisting entirely of atomic representations, then, can yield meanings which are independent of other representations in the system (58).

Insofar as we understand this suggestion, we do not think it helps.

Note that in adopting a two-component theory, Lormand abandons his advertised project of defending the strong form of holism

against the objection that it leads to meaning instability. What has become of his “positive aim” of showing “how holism, even at its most extreme, can accommodate and also increase meaning stability”? But, moreover, all of the stability that is achieved is due to the nonholistic component of the theory. If this is not obvious, consider that, apart from those few concepts in the “small set of semantically atomic (or simple, primitive, basic) representations,” we need to assume, in accordance with meaning holism, that the concepts in the representational system are thoroughly interconnected. Then, as before, any change in belief leads to changes in the meanings of all of these concepts, and consequently to the meanings of any unit in which one or more of these concepts figure. The only unit meanings that remain stable are those which do not contain any concept whose meaning depends upon the network. So meaning stability is still a problem exactly to the extent that the meanings of a representational system are determined holistically.

Much the same point applies to Lormand’s suggestion that we might “simply add referents” to the list of meanings a concept has, thereby gaining all of the stability that a nonholist reference-based theory has. Although this move would lend some stability to the conceptual system, the stability achieved is in direct proportion to the extent to which the account is nonholistic. In any case, standard two-factor versions of inferential-role semantics already have this much stability. The interesting question is whether Lormand’s proposal provides the meaning holist with a new and stronger response to the threat of meaning instability. Lormand suggests that his theory provides more stability than either a nonholistic reference-based theory or a standard two-factor theory because, in addition to the stability of the referential component, his account includes more fine-grained stability from the various stable meanings that concepts have. As we have shown, however, Lormand’s theory fails to deliver these additional points of stability, since, to the extent that the account is holistic, not only will all concepts change meanings with any change in belief, but all units will change meaning as well.

In sum, Lormand does not show how meaning holism can meet the objection that it suffers from meaning instability. Multiple meanings do not secure stability, since, assuming meaning holism, the multiple meanings of a concept are themselves unstable. So Lormand’s original bold claim that even a thoroughgoing holism is compatible with meaning stability—and that stability is actually a point in favor of meaning holism—is simply false. Moreover, isolating some units whose meanings are determined nonholistically does



not help, since the conceptual system remains unstable exactly to the extent that it is holistic. Meaning holists need to look elsewhere for ways to meet the threat of meaning instability; Lormand's multiple-meaning hypothesis leaves the threat exactly as it was.\*

ERIC MARGOLIS

Rice University

STEPHEN LAURENCE

London School of Economics

\*The order of the authors' names is arbitrary.