"I'M SORRY, DAVE, I'M AFRAID I CAN'T DO THAT": NON-NOMOLOGICAL USES FOR BELIEFS

Jane Duran

One's tendencies to want to discuss human functioning in terms of folk psychological constructs seem to remain unabated despite the current contretemps over the use and accuracy of such constructs.[1] In previous pieces I have argued that contemporary epistemic justification theory – specifically, certain sorts of coherentism – can be ameliorated by the introduction of "naturalized" material, i.e., material which actually makes use of notions relevant to the agent's cognitive functioning.

In a recent piece, I argued more specifically that the justificatory set for an agent seen as a coherentist might be deemed to be comprised of elements specified as the result of a series of intentional acts on the part of the epistemic agent in response to a skeptical challenger's queries.[2] But if epistemology can and should be naturalized, as many have argued, [3] it remains to be seen precisely what ramifications, if any, are to be had for the naturalization of epistemology by taking note of certain problems in philosophy of psychology, problems which are certainly relevant to the naturalization of epistemology. If, as some have argued, cognitive science tends to bottom out at the neurological level, [4] then any account of an agent's functioning which relies on intentionality is inherently problematic. Even those who have supported the notion of a place for folk psychology or a crudely-delineated intentionality have acknowledged that such constructs are probably not nomological. If cognitive science does bottom out, it at least has the virtue that whatever scientific work comes along to develop or replace it will probably be more susceptible to nomological status.

On the other hand, if (as another line of argument has tried to convince us) folk psychological constructs such as desires and beliefs can serve us well enough – providing that we do not ask too much of them – then another set of problems arises. They are not, of course, problems similar to those posed by the bottoming out of cognitive science on the neurological end.

Rather, they are the problems created by trying to be specific about what it is that folk psychological constructs can actually do for us.

Any attempt at the naturalization of epistemology which tries to employ an intentionality then has two distinct areas of difficulty, the first being the obvious difficulty of remaining explicit about how naturalized concepts can do work in, say, epistemic justification theory, and the second is the level of difficulty imposed by the use of problem-ridden constructs themselves.

In pieces antecedent to this piece, I have focused on what intentionality can do for epistemic justification theory.[5] I now want to address some of the problems posed by folk psychological constructs themselves. I want to argue that accounts which try to save beliefs, desires, etc., are essentially on the right track, but that they are still faced with gargantuan difficulties, since to assert that such concepts can be employed non-nomologically is to have achieved only a Pyrrhic victory. I want simultaneously to be specific about what the difficulties are, and how such accounts might address them. Finally, I will indicate what the foregoing might mean for a certain sort of view of naturalized epistemic justification theory.

I

In "Thoughts Without Laws; Cognitive Science Without Content," Millikan argues that the original desire to produce a psychological science with intentional components seemed to yield a view where "... it [was] implied that both the intentionality and the specific content of these intentional items must correspond to the kinds of lawful interactions these had with one another, with the environment, and with behavior."[6] One conclusion of her paper is that a science which deals (even implicitly) with the intentional "... need not be ... (nomological) to play a crucial role in the development of cognitive science." But the question remains how far one can go in terms of precision with the intentional before our intuitive desire to talk about beliefs and other intentional states "bottoms out." It may very well be the case that, as Millikan argues, non-nomological intentional constructs may be helpful in the formulation of hypotheses, and so forth.[7] But the difficulty remains what sorts of questions could be asked (or hypotheses formulated) without a more precise and nomologically susceptible set of criteria with which to work.

Even portions of the computational model of mind which are less folk-construct laden than homey folk psychology itself have

been criticized in much the same manner as have the original folk psychological constructs themselves.

The computational models have been constructed so that a crude version of any one of them usually adverts to a neuronal structure that is the equivalent of the computer's 1's and 0's. But when one wants to employ such a model one runs into immediate difficulties if the model is employed simultaneously for some sorts of intentional states. The literature on visual images, for example, bears this out. Block has written recently in *Philosophical Review* of the dispute between those who would like to limit the sorts of representations employed in the computational model essentially to descriptions, and those who disagree with this view.[8] Block wants to try to make out the case that if the arguments of the anti–pictorialists can be refuted (and if there is some room for visual imaging to be counted as a "representation" within the cognitive model) then cognitive sciences tend to bottom out. This sort of assertion is important for our argument if we remember that visual images are quintessential intentional objects.

> To the extent that our knowledge is to be explained by appeal to the nature of primitive operations, to that extent the kind of computational explanations at the heart of cognitive science explain less ...
> Then "How do we determine whether our current input is a word?" would not have a cognitive science answer ... but rather it would only be explainable in terms of physiology. That is, cognitive science explanations would "bottom out" sooner than expected. The primitive operations of cognitive science are not themselves explainable by cognitive science, so to the extent that the explanation of psychological phenomena requires appeal to primitive processes – to that extent, the psychological phenomena do not have cognitive science explanations.[9]

The difficulties here might be incomprehensible without a clear understanding of what "representation" is taken to mean within the framework of the computational model. The term is used in a different sense by some who have at least addressed certain of the broader concerns of cognitive science; such usage does not concern us here, however.[10] Within the scope of standard usage, the notion of *representation* is crucial to an overall understanding of the model itself. The general import of the term is "that which is semantically encoded." The representation is not a piece of the brains's hardware or (as Pylyshyn has it)

"functional architecture"; it is the semantic correlate of the syntactic neuronal flip–flopping to which we have earlier alluded.

One is tempted to say that a system's representational aspects amount to what philosophers have traditionally labeled intentionality. Propositional attitudes or intentional states may be deemed to be a set of representations; the postulation of such entities is necessary in order, according to the standard account, to fill in some gaps in everyday causal chains. Our contemporary desire to be able to talk about action in the world in accordance with brain states – rather than, say, homunculi, or Berkeleyan spirit-percept interaction – makes the most sense when we can posit some component which helps us along the way from the firing of synapses to an activity like eating ice cream. One theorist of the underpinnings of the computational model has written:

> Plainly, what is going on is, my behavior is being caused by certain states of my brain. Yet – and this is the crux of the problem – the only way to explain why those states caused me to type the specific sentences *about* walking, writing, the mountains and so on is to say that these states themselves are in some way related to the things referred to (writing, walking, mountains) ... My brain states are not, as we have noted, causally connected in appropriate ways to walking and to mountains. The relationship must be one of content: a semantic, not a causal, relation.[11]

The syntactic "wiring" of the brain, then (the neuronal structure), correlated with its semantic interpretations "causes" me (eventually) to eat the ice cream. This model provides an account which enables one to move by increments from the sort of hardware or physiological level to the world of mundane activities which one would like to be able to say are caused by the brain.

But the difficulties persist and are by no means alleviated when one would like to be more precise about the wiring, the semantic correlates, and how all of this fits into our everyday (folk psychological) notions of intentionality. The visual imagery debate has achieved the importance which it currently has precisely because it brings home to us, in an effective way, the split between the terminology/constructs which we would like to employ and that which the computational model – the going model – allows us to employ.

Block has already been cited as holding out hope that the

arguments of the anti-pictorialists can be refuted, and thus as pouring cold water on our efforts to keep cognitive science from bottoming out. Pylyshyn has himself written effectively on the visual imagery debate, but his arguments run in the opposite direction. It is important to get clear on his points, since they would seem to hold out more hope for a less than bottomed-out model.

In a piece directed specifically to the visual imagery problem, [12] Pylyshyn begins by reminding us of Kant's schemata. They are not (says Kant in a small excision from *The Critique of Pure Reason*) "images of objects." Pylyshyn then goes on to present both philosophical argument and some empirical evidence for the notion that even those sorts of representations which one would naturally want to think of as visual images should not be construed in that fashion. Several sentences are worth quoting in this context because they seem to sum up the anti-pictorial argument (the very argument which Block thinks is refutable), an argument which itself, as we have seen, is associated with reducing the intentional standpoint to the representations of the computational model.

Pylyshyn notes that there are strong reasons for not taking any representations to be visual images *simpliciter*:

> Although we often appear to go through a process of recalling an image of a scene... The fact that we can recall a scene or part of a scene, by addressing aspects off the perceptually interpreted content of the scene argues that what we have stored is *already interpreted* and not in need of reperception as we supposed.[13]

And again (with regard to experiments showing that small children have difficulty reproducing phenomena they have just witnessed),

> children have not mastered adult conventions and various physical principles, *so their reproductions and imitations are a source of more dramatic illustration of the abstract conceptual or descriptive nature of mental representation. The principle is not by any means, however, confined to children*.[14]

What is it, according to Pylyshyn, which is responsible for the descriptive representations (i.e., semantic, interpreted and largely non-visual representations) which we are wont, in our folk psychological way, to associate with mental imagery?

Pylyshyn (and he is not alone here) is disconcertingly vague on this. The tendency in the literature is to speak of "analogue representations" and "transducers." It is admitted that these semantically-loaded representations must be analogues of – must bear some relationship to – the properties of whatever it is that one is (folk psychologically) trying to "image." But it is precisely here, many have argued, that the computational model leaves so much to be desired that one is thrown back to the humdrum folk psychology. The upshot of this debate then, is this: on certain sorts of arguments, any computational model bottoms out and an account of the wiring or neurological processes would be necessary to give us the true picture. Other sorts of arguments – like Pylyshyn's against visual images – leaves us with the sensation that we have not bottomed out, but that it is very difficult to get beyond the level of talk about transducers to account for intentionality. Either account seems unsatisfactory.

II

Millikan argues that a virtue of the non–nomological folk constructs is that they "...enable us to do a great deal of explaining and a certain amount of fallible predicting of human behavior."[15] But it is important to try to be precise about what level of explanation and prediction is employable here. I take it that Millikan's overall argument is correct; we are drawn to folk psychological constructs (and they recur in the computational model, albeit in another guise) because on an everyday basis they are helpful. Millikan has several examples of this kind of folk psychological aid and succor in her piece. Consider John, who wants to meet the girl next door:

> The intentional characterization of John, 'He wants to meet the...' where the blank space is filled in and read transparently, *does* give us a handle on what John might well do, and certainly a handle by which we might later be able to explain why John did what he did do, though not a handle that fits into a deductive or nomological scheme.[16]

Or consider the intentional status of my desire to visit China:

> There is nothing wrong with me or my desire if I desire to visit China yet circumstances prevent my ever getting there or I decide I want to do something else more. But

> there is something wrong with me or with my belief if I
> believe that China is west of Europe.[17]

Finally, part of Millikan's conclusion is that "...folk psy-
chology, embellished with something akin to the traditional no-
tion that thoughts are like impressions or pictures, might be
placed in the context of modern physiology to yield a theory of
content for beliefs and desires – a naturalistic theory of inten-
tionality."
But the difficulty, surely, is just this. We understand, intui-
tively, that folk psychological constructs *are* useful and that
they do give us a handle on, at least, human and primate
behavior. It is precisely for these reasons that the literature on
intentionality continues to burgeon, [19] and that any view which
wants to discuss the wiring in detail – any view which shows
signs of bottoming out – has to be strategically and energet-
ically defended. Folk psychology is useful, but it is crude. The
constructs employed are so crude that, as Millikan readily ad-
mits, they do not fit into a deductive or nomological scheme.
When more closely examined, according to, for example, Stich,
they are next to useless.[20]
What, then, can be gleaned by the use of these constructs?
Certain cues to behavior and certain (fallible) predictions about
the future – but presumably we already knew that. The real
question is whether the constructs can do something for us with
which we are not already, on the basis of the mundane tasks of
living, familiar. The question is whether the use of the con-
structs holds any epistemological surprises for us, when em-
ployed in scientific or explanatory contexts.

III

What the folk psychological constructs can do for us is, in many
ways, outside of the purview of the computational model itself.
That is, any sort of critical view of the computational model of
mind encounters the very difficulties to which we have alluded
in our citations from Block and Pylyshyn; it is the purpose of
the model, ultimately, to aid us with the precise formulation (or
dissolution) of such constructs, not the other way around. But
cognitive science, broadly construed, comprises work in many
categories and disciplines, and some of the work so labeled does
employ these constructs in new and interesting ways, without
the necessity to investigate the status of the constructs them-
selves.

John Searle, in "The Intentionality of Intention and Action,"[21] reminds us of the intentional patterns inherent in speech acts of which we may be largely unaware. Viewed from the standpoint of the speech act itself, the pattern of intentionality is a useful construct, however crude, which aids us in the prediction of behavior at a broad-stroke level in much the same way that Millikan's example about John's beliefs about the girl next door does. Again, when one tries to be more precise about the intentionality (or the beliefs, *simpliciter*) one run smack up against the very theoretical difficulties we have just delineated. But, stopping short of that level, there seems to be a place for the intentional stance.

A useful exemplar here is Searle's description of what one intends to accomplish by certain sorts of speech acts: his example is of someone ordering someone else to leave the room. If the second person is gone after an interval of a few minutes, but later reports that she left to open a window, the first person has not accomplished her speech act goal.

> But what this illustrates is that the content of my order is not simply that you leave the room, but that you leave the room by way of obeying *this order*, that is; the logical form of the order is not simply that I order you (that you leave the room) but rather it is self-referential in the form I order you (that you leave the room by way of obeying this order).[22]

Here Searlian intentionality and Millikan's beliefs hook up in a nice way to yield something which does give us a handle on behavior. It is then to return to Millikan's example, not only the case that John's desire to meet the girl next door helps us guess what John might well do, but, at least in some situations, it helps us guess more specifically how John might go about it. To be precise, it helps us make the inference that not only will John engage in actions which might put him into proximity with the girl next door, but that in his interactions (if any) with the girl next door he will intend that *she recognize his intentions* (at some point).

The previous example has been adduced merely to help us along with the contention made by Millikan that beliefs (and other intentional states) "give us a handle" by which we may be able to predict and/or explain action, "though not a handle that fits into a deductive or nomological scheme."[23] Intentional states may give us handles which are, to some extent, rather precise. It is simply that they will not be precise in a nomological way.

In previous pieces I have argued that part of the "naturalization" of a cohering justificatory set involves paying attention to the construction of such a set by an epistemic agent responding to a challenger on a speech act level, and noting the extent to which the construction of such a set is a socially–bound practice which reflects the norms and standards of the participants.[24] In two pieces I have tried to be precise about how the amelioration of some of the problems posed by classical coherentism proceeds along the lines of naturalizing the coherentist's justificatory set (that is, limiting the size of the set by context) and paying attention to those elements of cognitive science which might aid us in this naturalization and contextualization.

Specifically, in "A Contextualist Modification of Cornman," I noted that "...Cornman, and coherence theorists in general, have a problem with underdetermination of justificatory sets because the size, scope and rigor of such sets is well beyond that which, on a model of descriptive adequacy, would be produced by a functioning epistemic agent."[25] I then went on to be still more precise about such a model, utilizing lines of intentionality similar to those addressed by Millikan and Searle and similar to those employed as exemplary in this paper.

But the truth is that employment of such lines helps us only so far, that is, it helps us in approximately the same manner and to the same degree that our knowing that John desires to meet the girl next door helps us to predict John's behavior. At a certain level, intentional constructs either bottom out (as we have seen that Block claims), or, failing that, they are analyzed along the lines of constructs (such as transducers) which are themselves unanalyzable and not helpful primitives. Goldman, Kornblith and others have shown us that epistemology and cognitive science have a natural intersection. What is required is still further elucidation of what the intersection amounts to if work in epistemology is to be interestingly aided by cognitive science.[26]

University of California at Santa Barbara
and
Mount St. Mary's College

NOTES

1. See, for example, Stich, Stephen, *From Folk Psychology to Cognitive Science*, Cambridge: Bradford of MIT Press, 1983; Block, Ned, in various publications, but especially "Mental

Pictures and Cognitive Science," in *Philosophical Review*, XCII, No. 4 (October 1983); and Millikan, Ruth Garrett, in several pieces, but particularly "Thoughts Without Laws; Cognitive Science Without Content," in *The Philosophical Review*, XCV, No. 1 (January 1986), p. 48.

2. See my "Intentionality and Epistemology," in *The Monist*, vol. 69, No. 4, (Winter 1987).

3. There is a plethora of journal articles, but among books I cite Alvin Goldman's *Epistemology and Cognition*, Cambridge: Harvard University Press, 1986, and Hilary Kornblith's *Naturalizing Epistemology*, Cambridge: Bradford of MIT Press, 1985.

4. Block, in *op. cit.*, *passim*.

5. See *The Monist*, *op. cit.*; see also "Descriptive Epistemology," in *Metaphilosophy*, vol. 15, Nos. 3&4; "A Contextualist Modification of Cornham," in *Philosophia*, Vol. 16, Nos. 3&4.

6. Millikan, *op. cit.*, p. 48.

7. *Ibid.*, p. 56.

8. Block, in *op. cit.*

9. *Ibid.*, p. 533.

10. The usage is different, for example, in some of Millikan's work.

11. Pylyshyn, Zenon, *Computation and Cognition*, Cambridge: Bradford of MIT Press, 1984, p. 130.

12. Pylyshyn, Zenon, "Imagery and Artificial Intelligence," in *Readings in Philosophy of Psychology*, Vol. 2, Ned Block, ed., Cambridge: Harvard University Press, 1981, pp. 170-194.

13. *Ibid.*, pp. 172-173.

14. *Ibid.*, pp. 175-176.

15. Millikan, *op.cit.*, p. 50.

16. *Ibid.*, p. 65.

17. *Ibid.*, p. 68.

18. *Ibid*, p. 77. It is worth noting that Millikan does not address the literature on pictorial representation. The traditional notion to which she refers is not a favored piece of theory within the scope of the computational model.

19. See, for example, the entire issue of *The Monist*, Vol. 69, No. 4.

20. Stich Stephen, in *op. cit.*

21. In *Perspectives on Cognitive Science,* Donald Norman, ed., Ablex Publishing Co., Norwood, NJ, 1981.

22. *Ibid.*, pp. 213-214.

23. Millikan, in *op.cit.*, p. 65.

24. See my pieces, *op.cit.*, fn. 1 and fn. 4.
25. In *Philosophia*, *op.cit.*, p. 386.
26. See the work cited in fn. 2.