

Is Meaning Normative?¹

Paul A. Boghossian

1. Introduction

The claim that meaning is a normative notion has become very influential in recent philosophy: in the work of many philosophers it plays a pivotal role.² Although one can trace the idea of the normativity of meaning at least as far back as Kant, much of the credit for its recent influence must go to Saul Kripke who made the thesis a centerpiece of his much-admired treatment of Wittgenstein's discussion of rule-following and private language.

Kripke, as you may recall, attributes to Wittgenstein the exciting and potentially paradoxical thesis that there can be no facts to the effect that people mean things by their words. To establish this startling claim, he proceeds by elimination: all the facts that are potentially relevant to fixing the meaning of an expression are examined and rejected.

Among the most promising candidates for being the facts in virtue of which I mean, for example, addition by '+' are facts concerning my dispositions to use that symbol in a certain way. Although this may not appear obvious at first sight, in fact the two leading naturalistic theories for the fixation of content – informational semantics and inferential role semantics – are both versions of a dispositional theory in the relevant sense.

Against this popular idea about naturalizing meaning, Kripke deploys a number of considerations: that our dispositions are finite; that one cannot read off our dispositions what we mean because our dispositions may include dispositions to make errors; and so forth. However, even if it were possible to overcome these objections, Kripke argues, one could still not identify meaning facts with dispositional facts because at bottom the relation between meaning and future use is *normative*, whereas the dispositionalist construes it descriptively.

In a moment, we shall look at how this observation is supposed to work in some detail. But the important point to note is that, if the alleged normativity

¹ This paper overlaps with, and tries to improve upon, the discussion of this issue in Boghossian (2003).

² See John McDowell (1984), Crispin Wright (1984); Simon Blackburn (1984); Allan Gibbard (2003).

of meaning is to be used in this anti-naturalist way, to knock out proposed theories of meaning, it must be established intuitively and pre-theoretically, as something that every theory of meaning would have to respect, and not on the basis of assumptions about the nature of meaning that a naturalist could regard as optional.

In what, sense, then, is meaning intuitively normative and how does that help defeat naturalist conceptions of it?

2. Kripke on the Normativity of Meaning

For expository purposes, I will work with the simple version of the dispositional theory that Kripke operates with, for nothing essential will depend on the simplification. As Kripke tells it, then, the dispositionalist attempts to explain what it is for me to mean addition by ‘+’ by saying that it consists in my being disposed to respond to the question

$$x + y = ?$$

with a number that is the *sum* of x and y , as opposed to my being disposed to respond with a number that is their *quum* (where the quum of $x + y$ is their sum if $x, y < 57$, and is 5 otherwise).

There are, of course, serious problems getting a dispositionalist account to work, even putting aside issues about normativity. The main difficulty is to get a dispositional theory to be extensionally adequate so that it yields the correct verdicts about what people mean by their expressions. And the main problem here is to naturalistically specify a set of optimality conditions which will be such that, under those conditions a thinker will be disposed to apply an expression to something if and only if that item is in its extension.

According to Kripke, though, we needn’t bother trying to see if we can ever find such a set of optimality conditions because there is a problem of principle in seeking to reduce meaning facts to facts about dispositions. Kripke writes:

Suppose I do mean addition by ‘+’. What is the relation of that supposition to the question of how I will respond to the problem ‘68+57’? The dispositionalist gives a descriptive account of this relation: if ‘+’ meant addition, then I will answer ‘125’. But this is not the proper account of the relation, which is normative, not descriptive. The point is not that, if I meant addition by ‘+’ I will answer ‘125’, but that I should answer ‘125’....The relation of meaning and intention to future action is normative and not descriptive. (Saul Kripke 1982, 37)

Notice that Kripke’s formulations – both of problem and of solution – tend to be in terms of the notion of *linguistic meaning* rather than in terms of *mental content*, although his argumentative strategy makes it clear that he holds that both notions are normative. In fact, it is important for present

purposes to distinguish between them. I will begin with the thesis as applied to linguistic meaning.

Is it really true that, if I mean addition by '+', then, if am asked what the sum of 58 and 67 is, I *should* answer '125'? What if I feel like lying or misleading my audience? Is it still true then that I should answer '125'? If I want to mislead, it looks as though I should not say '125' but rather some other number.

Of course, we can say that, if you mean addition by '+' and have a desire to tell the truth, then, if you are asked what the sum of those two numbers is, you should say '125.' But that is mere hypothetical normativity, and that is uninteresting: every fact is normative in that sense. (Compare: if it's raining, and you don't want to get wet, you should take your umbrella.)

If there is to be an interesting thesis of the normativity of meaning, we ought to be able to derive a *should* or an *ought* from the mere attribution of meaning to someone and without having to rely on any auxiliary desires that that person may or may not have. But can we do that?

Does it follow from the mere fact that I mean addition by '+' that I should not lie or mislead? There may be, for all I know, a *moral* prohibition against lying or misleading; but are there such prohibitions flowing from the nature of meaning itself?

I suppose there have been philosophers who have held complicated views to the effect that the very possibility of meaning requires that lies and deception not be very common.

But, first, it is not at all clear that these philosophers are right. And, second, and as I emphasized at the beginning, what we are after is some intuitive thesis to the effect that meaning is normative, not something that would lie at the end of some complicated 'transcendental' argument. What, however, is the intuitive normative truth that falls directly out of the attribution of meaning, so that such normativity can then be used as a constraint on theories of meaning?

One thought that might seem to be in the right neighborhood is this:

If I mean addition by '+' then, although I may not be disposed to say '125', in response to the question $68+57=?$, it is *correct* for me to say '125'.

The trouble is that it is not clear that, at least as it is being used here, "correct" expresses a normative notion, for it may just mean "true." Of course, if I mean addition by '+' then I will only have said something *true* if I say '125.' But there is no obvious sense in which truth is a normative notion.

In my own earlier work on what are referred to as the "rule-following considerations," I underestimated the force of this point. In that paper, I wrote:

Suppose the expression 'green' means *green*. It follows immediately that the expression 'green' applies *correctly* only to these things (the green ones) and not to those (the non-greens)...The normativity of meaning turns out to be, in other words, simply a new name for the familiar fact that...meaningful expressions possess conditions of *correct use*. Kripke's insight was to realize that this observation may be converted into a condition of adequacy on theories of the determination of meaning: any proposed candidate for the property in virtue of which an expression has meaning, must be such as to ground the 'normativity' of meaning...it ought to be possible to read off from any alleged meaning-constituting property of a word what is the correct use of that word. (Boghossian 1989, 513)

As I now see it, there is nothing wrong with the substance of this passage: there *is* a problem capturing an expression's satisfaction conditions in dispositional terms. But it is misleading at best to use the label "the normativity of meaning" in connection with this problem since, as I remarked above, there is nothing obviously normative about the notion of a truth condition or a satisfaction condition. (Putting the word 'normativity' in shudder quotes, as I did, helps, but in retrospect, more skepticism about its use was probably in order.)

Well, what conditions must an expression satisfy if it is to express a genuinely normative notion? This is, of course, a difficult question, but a first stab might well look like this:

An expression E expresses a normative notion only if it is *constitutive* of our *understanding* of E that its application implies an *ought* or a *may*.

Later on, I will come back to what we should take "constitutive" to mean in this context, and what "ought." But however liberal we may be with these notions, it seems to me, we have been shown no clear reason to think that "true" is a normative notion (I shall come back to this).

Consider, however, the corresponding observation in the case of *content and belief* rather than in that of *meaning and assertion*:

If it is addition that I am thinking in terms of, rather than quaddition, then although I may not be disposed to believe that $68+57=125$, because I have, we may suppose, a tendency to make certain kinds of arithmetical error, still it is only *correct for me to believe* that $68+57=125$.

In this case, it looks as though there may be a real chance that the notion of correctness in question doesn't just mean "true" but expresses something genuinely normative, something that would ground an *ought*. ...although I may not be disposed to believe that $68+57=125$, still that is what I *ought* to believe, given that I mean addition by '+'. This formulation does not offend in the way in which the claim about linguistic meaning did.

Or consider the following claims:

If you mean *negation* by ‘not’, then you ought not to accept both ‘p’ and ‘not-p.’

If you mean *if* by ‘if’ then if you accept that ‘It is raining’, and you accept that ‘If it is raining, then the streets are wet,’ then you ought to accept that ‘The streets are wet.’

Each of these claims seems plausible and yet neither of them seems to depend for its truth on any auxiliary desire by the thinker.

Not so fast, you might think: don’t these oughts depend on your wanting to believe the truth, so that they are at best hypothetical imperatives, just like the imperatives in the case of linguistic meaning and assertion?

I think that there is a good basis for saying that the answer to this question is ‘No,’ that the aim of truth is built into the nature of belief in the way that it is not built into the nature of assertion. If that’s right, then we don’t need to invoke any auxiliary desires in order to explain why these ought statements come out true.

At least if we are looking at belief and mental content, then, rather than at assertion and linguistic meaning, there seems to be a chance that the thesis of the normativity of meaning might actually be true.

Gibbard’s Objections

Allan Gibbard has objected to the idea that we can ground the normativity of content in this way, through content’s link to what it is correct to believe. Let us switch to an empirical example. Concerning Mallory’s last day on Everest we can say:

- (0) It is correct to believe that Mallory reached the summit iff Mallory did reach the summit.

Now, suppose we understand “correct” in a normative sense, so that it is taken to imply an ought. Then (0) would appear to yield:

- (1) One ought to believe that Mallory reached the summit iff Mallory did reach the summit.

Now Gibbard raises two objections.

First, the sense of ought involved in (1) is clearly (what is known as) an “objective” ought, an ought that kicks in in light of what is true rather than in light of the evidence currently available to one. To appreciate this distinction, suppose you have a choice of two flights, one operated by British Airways and the other by Lufthansa, leaving at more or less the same time, from the same airport, the British Airways flight costing significantly less than the Lufthansa flight. Naturally, you opt for the British Airways flight, seeing no reason to prefer the Lufthansa flight. Unfortunately, the BA flight ends up

being canceled because of a late-breaking labor strike and you are stranded. You think: I should have taken the Lufthansa flight. There is a sense in which that thought is true – in an “objective” sense, in light of what has turned out to be true. Subjectively speaking, in terms of the evidence available to you at the time of decision, it wasn’t true that you should take the Lufthansa flight: by hypothesis, the evidence available to you favored the BA flight.

Similarly in the case of our belief about Mallory, perhaps Mallory was a clever illusionist who only made it seem as though he had scaled the summit of Everest. Then, in a subjective sense, I ought to believe that he scaled the summit even if he didn’t. So the sense of “ought” in which (1) is guaranteed to be true must be an “objective” sense and not a subjective one.

Now, Gibbard continues, if we nevertheless insist that this objective ought is genuinely normative, then we will have to say that all facts are normative; and that would clearly be absurd.

For suppose that this objective ought is both genuinely normative and genuinely constitutive of our thought that Mallory reached the summit. Well, if it is genuinely constitutive, then (1) gives, as Gibbard puts it, an “analytic equivalence”: Meaning facts alone enable us to see that the biconditional in (1) is true. And if the implicated ought is genuinely normative, then, given the analytic equivalence, the right hand side, that Mallory reached the summit, describes a normative fact as well. But, surely, that Mallory reached the summit, is a non-normative claim if anything is.

I think there are a number of ways of resisting Gibbard’s argument here, the most direct one being this. It is not true, I think, that the relation between a thought’s correctness conditions and the corresponding ought claims is biconditional in form. (1) may be broken up into two conditional statements:

- (2) One ought (objectively) to believe that Mallory reached the summit, only if Mallory reached the summit.

And

- (3) If Mallory reached the summit, then one ought (objectively) to believe that he did.

Putting them in their general form:

- (4) For any p: One ought to believe that p only if p.
 (5) For any p: If p, then one ought to believe that p.

It’s clear, I think, that these two claims are not on a par. The first, I think, may be taken to be the correct expression of the norm for belief – that one ought to believe only what’s true (more on this below). But no one thinks that it’s a norm on belief that one believe *everything* that’s true. How could it be, given the metaphysical impossibility of believing everything that’s true?

I'm inclined to hold, therefore, that one can infer p from 'One ought to believe that p ,' but not 'One ought to believe that p ' from p . This by itself disarms the fear that a correctness-based normativity thesis will lead to a normative explosion.

Objective and Subjective Norms

But even if that is so, isn't there still something irredeemably fishy about this "objective" sense of ought: something that one ought to believe just because it's true rather than because it's compelled by the evidence available to one. It can objectively be the case that one ought to believe p , even though, subjectively, one might well be required by the evidence at one's disposal to believe not p .

I think it is relatively easy to defend the importance of the idea of objective oughts.

The main problem with the norm encoded in (4) is that it is not a norm that is directly followable. One may well have to follow *other* norms as a means towards obeying it.

But the mere fact that (4) is a norm whose satisfaction isn't transparent doesn't mean that it isn't important, or that it's not a real norm.

We are often in the position of attempting to comply with some non-transparent norms by following other more transparent ones. Traders on the stock markets are attempting to comply with the rule: Buy low, sell high. But there is no direct way to recognize when a stock's price is low relative to the price for which one will be able to sell it. So traders follow certain other rules as a means of attempting to comply with the non-transparent rule that truly captures the aim of their trading activity. Some will use rules based on technical indicators, others will use rules based on fundamentals. These are rules that may be followed directly, by doing what the rules call for when their input conditions are recognized to obtain. However, we would give a seriously misleading account of their activity if we left out the fact that the following of these rules isn't an end in itself but only a means of complying with the non-transparent rule of buying low and selling high.

Just so, I think, with the "objective" norm that one ought to believe only what's true. Once again, this is not a rule that can be followed directly, but that can only be followed by following certain other rules, the so-called norms of rational belief. For example: that we ought to believe that which is supported by the evidence and not believe that which has no support; that we ought not to believe p if some alternative proposition incompatible with p has a higher degree of support; that we ought to believe p only if its degree of support is high enough, given the sort of proposition that it is. And so on.

But, just as before, our story would be incomplete if we left out the fact that our following of these rules is a means of following the norm that we ought to believe only what is true. All of these epistemic norms are grounded in the objective norm of truth. It is that ought that supplies their rationale,

even if it has proven extremely difficult to say – in the theory of knowledge – exactly how.

It is this objective norm that captures the idea that it is constitutive of belief to aim at the truth, and so that something goes wrong if a belief is false. That, in my view, is what makes belief the state that it is.

Normativity of Content vs. Normativity of Belief

Early on in the paper, I said that there was a difference between the normativity thesis as it applies to mental content and as it applies to linguistic meaning. We are in a better position now to say what that difference is and why it holds. To put the matter concisely, the linguistic version of the normativity thesis, in contrast with its mentalistic version, has no plausibility whatever; and the reason is that it is not a norm on assertion that it should aim at the truth, in the way in which it is a norm on belief that it do so. Thus, the only imperatives that flow from attributions of linguistic meaning are hypothetical imperatives.

Kripke says: If I mean addition by ‘+’ then it doesn’t follow that I *will* say that ‘ $68+57=125$ ’, but only that I *ought* to say that it does. But it seems to me that neither claim follows. In particular, the ought claim doesn’t follow because, even though I mean addition by ‘+’ and know therefore that it would only be correct to say that ‘ $68+57=125$ ’, I might still not choose to say it because I might deliberately not choose to say what I know to be correct. Deciding knowingly to assert what is false is not to undermine the very possibility of assertion.

By contrast, no desire or decision is needed for it to be true that I ought to believe that $68+57 = 125$. Indeed, the very fact that the imperative here is not hypothetical is, as I’ve just been arguing, a defining feature of belief. It is what makes it the state that it is.

The difficult question in this vicinity, I think, is not about the truth of the claim that attributions of the form *Wolfgang believes that p* are constitutively normative, nor is it about its importance; rather, it is about its *source*: Does the fact that such attributions are normative reveal something about our notion of *content*, or does it reveal something, rather, about our notion of *belief*? Do we have here a thesis of the normativity of content, or a thesis of the normativity of belief?

We have said that belief attributions are normative because it’s a condition on understanding them that one understand that one ought to believe that *p* only if *p*. If we look at things this way, then it does seem as if what’s responsible for the normativity is the concept of belief and not that of content. After all, contents can figure in *other* attitudes about some of which there aren’t norms. If it’s *content* as such that’s normative, why aren’t there norms governing these other attitudes? If it’s genuinely constitutive of content that it be normative, shouldn’t it carry this normativity with it wherever it goes?

Take a concrete example. Suppose I say of Ebenezer that he *wants* that Howard Dean be the next President. In making this attribution, am I in any way speaking oughts?

There are views, of course, according to which there are facts about what is objectively desirable. On such views, one could say that desires are correct only in so far as they line up with those objective facts. If such a species of evaluative realism were true, that would ground oughts about desires. But the source of these oughts would lie squarely in the evaluative realism, whereas what I'm asking is whether there are oughts about desires in virtue of the mere fact that they are contentful states. To be sure, Ebenezer's desire has conditions of satisfaction – it will be satisfied if and only if Dean is the next president. But, in and of itself, this doesn't translate either into a correctness fact or into an ought of any kind. Of course, Ebenezer may have this particular desire because he believes it to be a way of securing the satisfaction of another of his desires, and so his desire may be said to be correct to the extent that his belief is true. But that would be entirely a matter of the correctness of the underlying belief; it wouldn't introduce a sense in which the desire itself may be subject to normative evaluation.

It's not clear to me, then, that there are norms on desire merely qua contentful state. The matter is perhaps even clearer in the case of pure thinking, the pure entertaining of a proposition. Suppose I say that Ebenezer is merely entertaining the thought that Dean will be the next president. He doesn't believe it, he doesn't desire it – he's merely thinking about it, turning it over in his mind. In attributing this content to him, am I in any way attributing oughts? It seems not. As far as entertainings are concerned, you can do what you want with them.

Doesn't all this imply that the notion of thought content is not normative as such, that the answer to the question that constitutes the title of this paper is 'No'?

I think that it does, *unless* the following is true: that we understand the role that contents play in propositional attitudes generally only *through* our understanding of their role in belief. If our grasp of the notion of content were somehow to depend in a privileged and asymmetric way on our grasp of the concept of belief, then our only route to the notion of a contentful state would be through our grasp of a constitutively normative notion, and – although we would have arrived at this result in a way not envisioned by its proponents, still – that would be enough to substantiate the claim that content itself is normative, in spirit if not in letter.

Let me review the dialectic up to this point. I have said that a judgment type is normative just in case it's constitutive of our understanding of judgments of that type that they imply oughts. I have also argued that attributions of belief are normative judgments in this sense. If, then, we could be said to understand content only through our understanding of belief, then the notion of content would turn out to be a constitutively normative notion.

We would understand content only through belief, and belief only through normative notions. This may not be quite what Kripke, Sellars and others had in mind, but it would still count as an interesting thesis. If, however, it is not true that content depends on belief, that content may be understood through its role in other non-normative attitudes, such as desire or the pure entertaining of a proposition, then we would not have a thesis of the normativity of content but only the rather different thesis of the normativity of belief.

Is Belief Special?

Now, I take it that the concept of a proposition, or content, just is the concept of whatever it is that is the object of the attitudes. And, of course, the notion of content could, in principle, be introduced in connection with the notion of, say, desire: *prima facie*, at least, it doesn't seem to have a *privileged* connection to belief. So our question isn't so much whether the notion of content *can* be understood in connection with attitudinal concepts other than that of belief, but whether any non-belief based understanding would covertly presuppose an understanding of its role in belief. In other words, is belief, in some appropriate sense, conceptually primary?

Let us ask this question not in full generality, but in connection with the notions of desire and belief, extending the question to the other attitudes only later. Is there any asymmetry in our understanding of belief and desire? Do we understand the one notion through understanding the other? Or are they on a par, either both depending, or neither depending, on the other?

A functionalist about the concepts of the attitudes would, of course, deny that there is any asymmetry. He would view the two concepts as graspable only jointly. However, just as I earlier didn't want to assume a controversial theory of how content is determined, so I don't now want to assume a controversial theory of the concepts of the propositional attitudes. To whatever extent it is possible, I want to ask about our understanding of the attitude concepts in a pre-theoretic and intuitive manner.

Let us begin, then, with the following question: Could someone have the concept of belief without having the concept of desire? *Prima facie*, this would appear to be so: it does seem possible for someone to have the idea of *accepting a content as true* without having any idea of what it would be to *desire a content to be true*.

One way in which this conceptual appearance could be falsified is if it turned out that I couldn't coherently think of someone as believing something without also thinking of them as desiring something, if it were conceptually impossible to think of someone as a believer without also thinking of them as a desirer. But this doesn't seem impossible. At least at the intuitive level, there appears to be no difficulty in thinking of someone as a pure believer: that is, as a creature who only has views about how things are, but no conception of how she would want them to be.

Let us now ask the converse question: Could someone have the concept of desire, but not yet the concept of belief? Could someone understand the idea of wanting the world to be a certain way, but have no idea at all of what it would be to take it to be a certain way, to accept its being a certain way?

This does seem bizarre. Don't I have to think of someone as having some beliefs about how things are, in order to coherently think of them as having wants about how things should be.

Where the desires in question are conceived of as perceptually mediated *de re* desires, the purported possibility does seem incoherent. It's hard to understand how someone could be said to want that perceptually presented apple without (in some appropriate sense) believing that there is an apple there.

What about *de dicto* desires? Could I think of a person as wanting that all sorts of propositions be true without thinking of it as having any beliefs whatever? It's hard to imagine. The reason is that we think that someone can want *p* at some time *t* only if he either believes it to be not *p* at *t*, or if he is unsure whether it is *p* at *t*. You cannot want *p* at a given time, if at that time you already believe that *p* has occurred. You can be glad at *t* that *p* has occurred, if you already believe that *p* has occurred; but you cannot want it to occur. If I now want it to snow, that can only be because I currently believe it not to be snowing.

If that's so, then understanding desire involves understanding the idea of wanting things to be different than they are actually believed to be, and so presupposes the concept of belief.

These considerations are admittedly sketchy. But they suffice, I think, to make a plausible *prima facie* case for an asymmetry in our grasp of belief and desire: grasp of the concept of desire seems to asymmetrically depend on our grasp of the concept of belief in just the way that, I have argued, the normativity of content thesis requires. (If this is right, then we would have here a significant objection to the functionalist analysis of the concepts of belief and desire, for those analyses treat grasp of these concepts as symmetric: either you grasp both or you grasp neither).

A full treatment would require conducting similar investigations into the relations between all the other non-normative attitudes with reference to which the notion of content can be understood and seeing whether for each of them it is true that an understanding of that attitude depends asymmetrically on an understanding of belief. Until such an investigation is carried out, one which I think has an interest quite apart from the focus of our present concerns, the question whether content is normative will have to remain unresolved.

Just looking at the matter in a cursory way, it's impressive, I think, how many concepts of the propositional attitudes depend asymmetrically on the concept of belief: for example, all of the following seem to me to conceptually presuppose the belief that Bush won the election:

Being glad that Bush won

depressed that he won
 sad that he won
 angry that he won
 conflicted about the fact that he won

In fact, the only sort of attitude concept that doesn't seem to presuppose that of belief is that of the mere entertaining of a proposition. What isn't clear to me, however, is whether this appearance isn't after all illusory. For is it clear that the peculiar notion of entertaining a proposition is not just a negative notion, the notion of

Thinking about a proposition without taking up any doxastic or conative attitude towards it.

Naturalistic Theories of Content

However this more extensive investigation may ultimately turn out, what is clear, ironically enough, is that the philosopher with perhaps the most reason to believe the asymmetric dependency thesis, and hence the normativity thesis, is the *naturalist* about mental content. For the most promising ideas that I know of concerning how to understand content naturalistically come in either one of two forms: either through an informational semantics or, especially in the important case of the logical constants, through an inferential role semantics. Both of these sorts of account, however, have to understand the fixation of content through its role in the fixation of *belief*.

To see why, reflect on how such theories propose to naturalize content. Let us begin with the case of an informational role semantics and, to ease exposition, let us assume that propositional attitudes are relations to mental representations, say, though this is not essential, to sentences in a language of thought. For example, let us take the state of desiring that *p* to be, or to be realized by, the state of desiring a sentence *S* of mentalese that means that *p*. Informational semantics theories attempt to specify naturalistically what it is for a mentalese sentence *S* to mean that *p*.

If you look at how these theories attempt to achieve this, you will see that they inevitably go through the notion of belief, or at least through its computational counterpart. All such theories attempt to understand what it is for an arbitrary mentalese sentence *S* to mean that *p* by specifying the conditions under which *S* would be placed in the belief box (to use Stephen Schiffer's useful metaphor):

S means that *p* iff under optimal conditions *O*, *S* would be placed in the belief box iff *p*.

In other words, these theories depend on the idea that there is a set of conditions under which one will believe something when and only when it's true. As such, they seek to understand the notion of content through its role in the fixation of belief.

Not only is this the way things are typically done; it's very hard to see how they are to be done otherwise: for no attitudinal state other than belief has anything like the hope of covarying naturalistically with the conditions under which its content is realized.

To put the point abstractly, the idea would be to come up with the specification of a set of conditions *C* and a propositional attitude *PA*, such that, for all the atomic propositions in an organism's repertoire, the following holds true:

C: $PA(p) \leftrightarrow p$.

That there should be such a set of conditions for belief is already incredible enough.³ But it doesn't have any chance, it seems to me, where the propositional attitude in question is anything other than belief. For example, and in part for reasons reviewed earlier, there is no chance that there is a set of conditions under which one will desire something only if that content is already realized.

All of this is even more true in the case of an inferential role semantics, as its name already implies. In the case of an inferential role semantics, the idea is that *S*'s meaning is fixed by the set of beliefs that lead one to accept it and by the set of beliefs to which accepting it gives rise.

Interestingly enough, then, the philosopher with the most reason to believe in the primacy of belief and, hence, given the way I have set things up, in the normativity of content is, ironically enough, the naturalist about content. But if this is the only way in which the normativity of content can turn out to be true, it shows what an uninteresting thesis it is, especially in the context of the dispute with the naturalist. For on this setup, the naturalist will only have reason to believe in the normativity of content if his naturalistic theories of content fixation are true. But that means that there is no longer any way in which to use the putative normativity of content to argue that those naturalistic theories are not true.

I don't want to leave you with the misleading impression that I have suddenly become a naturalist about content. I have not. For reasons that I have given elsewhere, I still don't see how to pull off the trick of specifying a naturalistically adequate set of optimality conditions that will allow naturalistic theories to specify the meanings of expressions correctly. But this is not because content attributions are constitutively normative but because intentional facts seem not to be reducible to naturalistic facts.

³ For discussion see Boghossian (1990).

Bibliography

- Blackburn, Simon (1984). The Individual Strikes Back. *Synthese* 58:281-302.
- Boghossian, Paul A. (1989). The The Rule-Following Considerations. *Mind* 98:507-549.
- Boghossian, Paul A. (1990). Naturalizing Content, in B. Loewer and G. Rey (eds.). *Meaning in Mind*. Oxford: Blackwell.
- Boghossian, Paul A. (2003). The Normativity of Content. *Philosophical Issues: A Supplement to NOUS* 13:31-45.
- Kripke, Saul (1982). *Wittgenstein on Rules and Private Language*. Cambridge: Harvard University Press.
- McDowell, J. (1984). Wittgenstein on Following a Rule. *Synthese* 58:325-364.
- Wright, C. (1984). Kripke's Account of the Argument Against Private Language. *Journal of Philosophy* 81:759-777.