

THE NORMATIVITY OF CONTENT¹

Paul A. Boghossian
New York University

It is very common these days to come across the claim that the notions of mental content and linguistic meaning are normative notions. In the work of many philosophers, it plays a pivotal role. Saul Kripke made it the centerpiece of his influential discussion of Wittgenstein's treatment of rule-following and private language; he used it to argue that the notions of meaning and content cannot be understood in naturalistic terms. Kripke's formulations tend to be in terms of the notion of linguistic meaning, but his argumentative strategy makes it clear that he holds both versions of the claim. He wrote:

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*.²

Now, in fact I believe that the linguistic meaning and mental content versions of the normativity thesis raise somewhat different issues, and later on I will say something about why. In the present paper I shall concentrate mostly on the mental content case, addressing the version in terms of linguistic meaning only briefly later on.

Many writers have followed Kripke in defending a thesis of the normativity of content. Among others, John McDowell, Crispin Wright, Simon Blackburn and I have endorsed Kripke's thesis, with varying degrees of hesitation, in our respective discussions of his arguments. More recently, Robert Brandom has published an extended treatise devoted to developing this theme.³

In spite of all this discussion, the claim of normativity still seems to me to be shrouded in unclarity. In these remarks, I want to try to straighten the matter out: Is there any interesting sense in which the notion of thought content is normative? And, if so, what is it?⁴

I will orient my discussion around some recent work of Allan Gibbard's in which he, too, tentatively endorses the normativity thesis. A close examination of Gibbard's arguments provides, I think, an excellent opportunity for clarifying the relevant issues.⁵

Modest Normativity, Ambitious Normativity and Naturalism

Gibbard says:

The thesis that “mental content is normative” is this: that when I attribute mental content—when I say, for instance, that Ebenezer is thinking that he has lost his keys—I'm somehow speaking oughts. (p. 85)

Three important questions need answering. First, is the claim merely that content attributions *involve* oughts, or is it the more ambitious thesis that they are *exhausted* by oughts? Second, what notion of “involvement” is at issue? And, finally, what notion of “ought?” I will take these questions up in turn.

Gibbard doesn't address the first question directly, but he does say something—something surprising—that carries implications for it.

It seems that physics describes the entire world, but that it pictures a world without any thinking in it. Yet we are learning more and more about how a physical system might act just as if it had thoughts; we are learning how natural selection and ecological dynamics might produce organisms to which we'd ascribe thoughts without hesitation. We, presumably, are such organisms. Still there seems to be a big gap between describing an organism as a complex physical system and as a product of natural selection and patterns of environmental impact, and saying that it is actually having thoughts.

The last couple of decades have suggested a way out of this puzzle. Saying what a person is thinking, the proposal is, is saying something *normative*. (p. 83)

Gibbard's surprising claim here is that the normativity thesis would, if true, actually support *Naturalism*, rather than, as everybody else has been inclined to believe, pose an obstacle to it. If there is any way to make sense of his claim, it can only be on the thesis that content attributions are exhausted by oughts.

For suppose the claim were merely that content attributions involve oughts. Then it would seem that this would make life harder for Naturalism, rather than easier. Naturalism had a hard enough time accounting for

thought content when mere intentionality was at stake; how much harder to account for it if, to boot, content facts turn out to involve normative claims as well. This thought, indeed, is the principal reason why Kripke places so much emphasis on the normativity thesis.

Interestingly, there is a naturalistic line of thought that might seek solace in the more ambitious thesis. For if content attributions are *exhausted* by oughts, then they are merely normative judgments in disguise. If, furthermore, one is antecedently attracted to a non-factualist treatment of normative judgments in general, as Gibbard is, then—bingo—content facts disappear in a puff of non-factualist smoke. The Naturalist's picture depicts a world without thought, so the Gibbardian non-factualist might argue, because thoughts, by virtue of being entirely normative, are no part of the natural causal order of the world.

If this is the reasoning that underlies Gibbard's surprising claim, then it seems to me misguided. Even if content attributions can correctly be said to *involve* normative judgments, there is no reasonable prospect that they will *exhaust* them, for there is no reasonable prospect that we will be able to do away with the *causal* role of content properties. We don't simply attribute thoughts to people; we attribute thoughts to them that, we claim, cause and explain their behavior. As many philosophers have pointed out, much of ordinary psychology depends on this assumption. But we couldn't attribute a causal role to content properties if we were thinking of them as *exhausted* by oughts and as subject to a non-factualist treatment. If that's right, then the most we should expect is a modest normativity thesis.

Modest Normativity and Rules

What, however, about this modest thesis? Is it true that content attributions involve oughts? And what, anyway, does it mean for a content attribution to involve an ought?

Gibbard is clear that what's at issue is a *constitutive* thesis:

If meaning is "fraught with ought", the point is not just supposed to be the obvious one that *oughts* can depend on meanings. *Oughts* after all, can depend on all sorts of things that are clearly natural. Whether we ought to take a walk can depend on the weather; that doesn't make the weather normative in any philosophically special sense. The claim of normativity is that *oughts* are somehow constitutive of meaning, that what something means *is a matter of* certain *oughts*' obtaining. (pp. 84–85)

So, it's *constitutive* of content facts' obtaining, that oughts obtain. But what does this in turn mean? Gibbard follows many in understanding this notion of constitutivity as follows: *q* is constitutive of *p* just in case *p* analytically implies *q*—i.e., just in case facts about the meaning of 'p' and the meaning

of 'q', along with logic, suffice to derive q from p. I shall come back to the question whether that is the best way to understand the thesis that oughts are *constitutive* of content.

How might we argue that content attributions constitutively involve oughts? The literature suggests two strategies. One is based on the connection between the notions of content and correctness, the other on the connection between content and rules. Gibbard favors the latter strategy and rejects the former. I shall argue that this gets things exactly the wrong way round.

Gibbard describes the rules-based considerations as follows:

What I'm thinking is a matter of the rules I am following in my thinking....This brings me to the best candidate I can find for a pattern of phenomena that might indicate that mental content is normative, that thought attributions are normative attributions. Consider a couple of simple stylized examples of a broad kind that permeates our thinking. First,

One ought not both to accept that it was snowing and to accept that it wasn't, for the same time and place.

This, in sample application, is a normative principle that seems somehow *constitutive* of negation. Consider too this example:

One ought not to accept that snow is white, to entertain the thought that something is white, but not to accept that something is white.

These patterns seem to characterize having thoughts with *something* content. The 'ought' here is an 'ought' in the subjective sense, not the fishy objective sense of correctness. (pp. 86–87)

Gibbard is here appealing to a specifically *normative* version of a conceptual-role semantics. The idea is that we come to grasp concepts by adopting rules, of permission or of obligation, for the conditional acceptance of contents involving them. It is a familiar idea, even if it is often underestimated how difficult it is to make it work. But there are three problems in using such a theory to justify a thesis of the normativity of content.

The first problem has to do with the controversial status of a specifically normative version of a conceptual role semantics. A number of philosophers are attracted by the idea that the meaning of the expressions in one language of thought are determined by their conceptual role. But it is not compulsory to think of this role in normative terms, as having to do with the permission or obligation to infer certain thoughts from other thoughts and/or perceptions. As a matter of fact, most of the proponents of such a view favor a directly naturalistic account, in terms of the *dispositions* to engage in certain patterns of reasoning and inference, rather than in terms of the obligation or permission to engage in those patterns.

A second difficulty has to do with the plausibility of such a view of meaning, understood as a *general* theory of meaning. Putting aside the question of whether the theory is best understood normatively or dispositionally, a conceptual role semantics is clearly the best idea that anyone has had about how to explain our grasp of the concepts of the logical constants, the only example that Gibbard actually offers. But there is a difficulty seeing how this sort of view can be made to work across the board, for all contents. Indeed, isn't it precisely the moral of the extensive discussion that *natural kind concepts* have received over the past twenty years or so that there are very few concept-constituting truths involving them, so that one couldn't hope to explain our grasp of those sorts of concepts by appealing to our adoption of concept-constituting truths involving them? If that is right, then the normativity of *content* cannot be grounded in the sort of rule-based conceptual role semantics that Gibbard alludes to because that sort of semantics cannot be assumed to apply to all contents.

Finally, this particular strategy for motivating a thesis of content normativity strikes me as peculiar, since the normativity thesis is supposed to be intuitive and is not supposed to depend for its appeal on a controversial theory of how content is determined. Indeed, in Kripke's original presentation of it, it was supposed to be so pre-theoretic that it could be used as a *constraint* on theories of content, rather than resting on a controversial instance of them. Is there such a thesis that could be said to flow from relatively uncontroversial and intuitive features of content?

Modest Normativity and Correctness

The second strategy I mentioned above would appear to provide just what's needed here, for it proceeds from the totally uncontentious observation that whether one is thinking correctly depends on what one is thinking, on the content of one's thought. Thus, if one is thinking that 68 plus 57 = 125, then one is thinking correctly, whereas if one is thinking that 68 quus 57 = 125, then one is thinking incorrectly, since 68 quus 57 equals 5, not 125. (Quus is defined as follows: 'x quus y' yields the same as 'x plus y' for all x, y, under 57, but yields 5 otherwise.) Furthermore, it seems right to say both that correctness is a normative matter, a matter of whether one ought to do what one is doing, and that the correctness conditions of one's thought are *constitutive* of what one is thinking.

Consider the following example discussed by Gibbard. 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.

Since, however, “correctness is a normative matter,” this would appear to yield:

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

Don’t we have here, then, a straightforward argument for the claim that content is normative?

Gibbard raises the following objection to this line of thought. He argues, first, that the sense of “ought” implicated in these considerations is an *objective* ought, not a subjective one. He next argues that to suppose that such an objective ought is both genuinely constitutive of thought content and genuinely normative yields the absurd result that *all* facts are normative.

For consider some thought concerning, for instance, the exact number of people who were wholly within the Parma city limits at 6 a.m. today. One such thought is correct, and so, in a sense, there is such a thought that I ought to have. But that’s not to say that I ought to have it given the evidence currently available to me, but only that I ought to have it in light of everything that is the case, whether or not I have any way of knowing it. So the ought implicated in (1) is objective.

Next, suppose that this objective ought is both genuinely normative and genuinely constitutive of our thought that Mallory reached the summit. Then we get the conclusion that all facts are normative. For if it is genuinely constitutive, then (1) gives, as Gibbard puts it, an “analytic equivalence.” And if it 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.”

Gibbard’s argument here rests on two assumptions that require further examination. The first is that the relation between a thought’s correctness conditions and the corresponding ought claims is biconditional in form. The second is that constitutivity is a matter of analytic implication. I shall look at these assumptions in order.

- (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.

Are both of these claims true? Let us put 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 is familiar as the standard 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. What's going on here? Since the belief that Mallory reached the summit is correct if and only if he did, why is there an asymmetry between (4) and (5)? Why don't both ought statements hold?

The reason is that the relation between correctness and normativity isn't completely straightforward: not every correctness fact generates a corresponding ought fact. In particular, oughts—even objective ones—are constrained by the requirement that they should imply cans; whereas correctness facts aren't. Thus, although it is true that, for any p, if p, then it is correct to believe that p, it doesn't follow that if p, one ought to believe that p, for it's clearly impossible to believe *everything* that's true.

I'm inclined to hold, therefore, that one can infer p from 'One ought to believe that p,' but not the other way round. Unfortunately, though, this doesn't by itself disarm the worry that every fact will turn out to be normative, for even given the inferences that I'm prepared to allow, it is still the case that every fact analytically implies a normative claim. For isn't it right to say that every fact gives rise to a *permission* to believe in it, and so that for every p,

If p, then it is permissible to believe that p.

And permissions are just as normative as oughts.

This fear that a correctness-based normativity thesis will over-generalize has worried a number of writers. But the problem arises, I think, only given the inadequate idea that a fact's being constitutively normative is a matter of its analytically implying an ought or permission. Let me suggest, however, a different way of looking at the matter.

When philosophers use the idea of some fact B being *constitutive* of some other fact A, they typically mean, I think, not that A analytically implies B, but rather a stricter condition along the following lines: that it's a condition on *understanding* what it is for A to obtain that one understands what it is for B to obtain. In other words, B is constitutive of A means: grasping the *concept* of an A-fact requires grasp of the concept of a B-fact. (Strictly speaking, of course, it doesn't matter what people mean by "constitutive" for it is a technical notion. What matters is how to characterize the normativity thesis so that it doesn't—by implying that all facts are normative—come out to be obviously false.)

Now, I hope it is clear that this is a stricter condition than mere analytic implication. A might analytically imply B, in the sense that the concepts ingredient in A and B are sufficient for seeing that B follows from A. However, it need not follow that B is constitutive of A because it might not be a condition on understanding A that one have all the concepts ingredient in B.

For example, the concepts ingredient in *Snow is white or grass is green* are sufficient for seeing that *Snow is white or grass is green* follows from *Snow is white*. But *Snow is white or grass is green* needn't be constitutive of *Snow is white* because it needn't be a condition on understanding *Snow is white* that one have the concept *or* or that one have the concept *grass* or, for all I know, that one have the concept *green*.

Applying this conception of constitutivity to the issue of normativity, I want to say that content attributions are normative just in case oughts are constitutive of content attributions. And I want to say that oughts are constitutive of content attributions just in case it's a condition on understanding a content attribution that one understand that if it is true, then the corresponding ought claim is also true.

Does Mallory's having reached the summit still come out as a normative fact, on this view? Well, it seems clear that one could perfectly well understand what it is for Mallory to have reached the summit without so much as having the concept of an ought or permission. By contrast, Marco could not be said to understand what it is for Ebenezer to *believe* that Mallory reached the summit unless he understands that Ebenezer's belief is correct only if Mallory did reach the summit, unless he understands, that is, that Ebenezer ought to believe that Mallory reached the summit only if he did.

Objective and Subjective Norms

Having made this adjustment in how a correctness-based normativity is to be understood, is there any further significance to the fact that the implicated ought is merely objective? Does that fact, by itself, lessen the interest of the corresponding normativity thesis?

What is true is that it will not always be transparent how one is to obey the norm encoded in (4). Subjectively speaking, one might well be required by the evidence at one's disposal to believe p, even if (unbeknownst to one) it is the case that not p.

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.⁶ On the contrary, I would maintain that the holding of this norm is one of the defining features of the notion of belief: it's what captures the idea that it is constitutive of belief to aim at the truth. The truth is what you ought to

believe, whether or not you know how to go about it, and whether or not you know if you have attained it. That, in my view, is what makes it the state that it is.

Furthermore, I would argue, it is because belief is governed by this objective ought that the less controversial *subjective* oughts hold of it as well. 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. All of these familiar 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.

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 now in a position 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 Mallory reached the summit only if he did. 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 *Ebenezer 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 what it is for S to believe that p that one understand that S 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* Gore to 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 be 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 Gore is the next President. But, in and of itself, this doesn't seem to 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 Gore 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? 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.

Let me review the dialectic up to this point. I have said that a judgment type is normative just in case you can't understand judgments of that type without understanding 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, just as Kripke, Sellars and others have claimed. If, however, this is not true, and contents may be understood through their 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, in all likelihood, 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. So there is no corresponding difficulty in imagining a judge who thinks of someone merely as a believer and not as a desirer, and so no difficulty in claiming that someone could have the concept of belief but not the concept of desire. Perhaps a deep transcendental argument could nullify this conceptual appearance, but I don't offhand see how it would go.

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?

Can I coherently think of a creature as wanting various things to be true, without thinking of it as having any views whatsoever about how they actually are? This does seem bizarre. Don't I have to think of Creature as having some beliefs about how things are, in order to coherently think of it 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 Creature could 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 Creature as wanting that all sorts of propositions be true without thinking of it as having any beliefs whatever?

What is it about this that strikes one as absurd? I think that the answer stems from the following fact: 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.

One might object: If Lars says, "I live in Aarhus because I want to live in Aarhus," we surely wouldn't want to say that he believes that he currently doesn't live in Aarhus, or that he is unsure whether he does. On the contrary, it is built into the example that he knows himself to be living in Aarhus.

It is easy, however, to say why this particular objection doesn't work. Someone's saying, "I want to live in Aarhus," while clearly believing himself already to be living in Aarhus, can mean only that he is expressing satisfaction at his prior choice of habitat, and indicating his inclination to continue in that choice. In ways that confirm the claim about the relations between desire and belief that I've been making, it isn't plausible to take him to be expressing a desire for a state of affairs that he already knows to obtain.

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 of the normativity of content 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 Clinton won the election:

Being glad that Clinton 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 state that doesn't seem to me to obviously 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 it is 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.

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 idea that I know of concerning how to understand content naturalistically is through an informational semantics, the sort of theory that Fred Dretske and Jerry Fodor have done so much to develop and defend. But an informational semantics has 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. First, they take propositional attitudes to be relations to mental representations, say, though this is not essential, to sentences in a language of thought. For example, they 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*. Then, they attempt to specify naturalistically what it is for the sentence *S* to mean that *p*.

If you look at how these theories attempt to elaborate the second part of this program, you will see that they inevitably go through the notion of belief, or at least through its computational counterpart. That is, 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 *C*, *S* would be placed in the belief box iff *p*. In a word, 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 held true:

C: PA(*p*) > *p*.

That there should be such a set of conditions for belief is already incredible enough.⁷ 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.

If that is right, then as I've said, the philosopher with the most reason to believe in the primacy of belief and hence in the normativity of content is, ironically enough, the naturalist. Whether this particular instance of the primacy of belief in our understanding of content is a mere artifact of naturalism or whether, as I'm inclined to believe, it reflects something deeper about content is, as I've already stressed, a question that will have to await a fuller discussion on another occasion.

Notes

1. This paper started out as a commentary on Allan Gibbard's "Thoughts and Norms," at a symposium of the Central Division of the American Philosophical Association in Pittsburgh in April 1997. In effect, it is an extended reflection on its opening pages. Thanks to Stephen Schiffer, Peter Kung, Masahiro Yamada and Derek Parfit, and to audiences in Pittsburgh, Parma and Budapest for comments on earlier drafts. For many of the central ideas of this paper, I am indebted to conversations with David Velleman. (I am not sure how many of them he would actually endorse.)
2. Saul Kripke: *Wittgenstein on Rules and Private Language* (Cambridge: Harvard University Press, 1982), p. 37.
3. See John McDowell: "Wittgenstein on Following a Rule," *Synthese*, 1984; Crispin Wright: "Kripke's Account of the Argument Against Private Language," *Journal of Philosophy*, 1984; Simon Blackburn: "The Individual Strikes Back," *Synthese*, 1984; Paul Boghossian: "The Rule-Following Considerations," *Mind*, 1989.
4. In my own earlier paper, I fell in with the normativity thesis without really clarifying what it amounts to, relying mostly on the existence of an a priori connection—which there is—between the meaning of a word and its *correct* use. However, I did not take it as evident that any such connection refuted naturalistic accounts of meaning.
5. Allan Gibbard: "Thoughts and Norms," this volume. All references to Gibbard are to this paper and will be incorporated into the text.
6. Compare: One ought to buy low and sell high.
7. For discussion see my "Naturalizing Content," in B. Loewer and G. Rey (eds.): *Meaning in Mind* (Oxford: Blackwell's, 1990).