CORE
Metadata, citation and similar papers at core.ac.uk
Provided by SAS-SPACE

Groundwork for a Nonconcessive Expressivism

Mark Eli Kalderon

June 18, 2006

1 The Frege–Geach Problem and Nonconcessive Expressivism

The Frege–Geach problem was first raised by **?** (1939: 33-34) and independently by **?** (1958, 1960, 1965) and Searle (1962, 1969) and was originally directed at expressivist proposals such as Ayer's (1946: 108) emotivism:

> It is worth mentioning that ethical terms do not serve only to express feeling. They are calculated also to arouse feeling, and so to stimulate action. . . . In fact we may define the meaning of the various ethical words in terms both of the different feelings they are ordinarily taken to express, and also the different responses which they are calculated to provoke.

Ayer's talk of 'defining' the meanings of ethical words suggests that he identified the meaning of an ethical word with the attitudes it expresses and provokes.

Ayer's emotivism is subject to the Frege–Geach problem because it is a form of atomistic reduction. It is atomistic in that it assumes a one-one correlation between the meanings of ethical words and the linguistic actions they perform (such as the expression and provocation of the relevant attitudes). It is a reduction in that it identifies the meaning of an ethical word with attitudes it expresses and provokes. Atomistic reductions face the following dilemma: Suppose that freestanding and embedded occurrences of 'wrong' mean the same. The problem is that words can occur in embedded contexts and fail to express the attitudes they do in freestanding contexts. This difficulty would be avoided if the atomistic reduction applied only to the meaning of freestanding occurrences. However, if the account applies only to freestanding occurrences, then it is incomplete, for the expressivist would lack an account of the meaning of ethical words in embedded contexts. Furthermore, some guarantee must be given that freestanding and embedded occurrences mean the same despite their meanings being differently determined. For, if freestanding and embedded occurrences differ in meaning, then the expressivist is apparently committed to the invalidity of recognized forms of valid argument (due to the fallacy of equivocation).

The proper target of the Frege–Geach problem is not expressivism per se, but those expressivist semantics, like Ayer's, that take the form of an atomistic reduction. A natural Fregean diagnosis for the failure of atomistic reduction is that it runs afoul of the attitude/content distinction. The meaning common to freestanding and embedded occurrences of a word despite the difference in attitude expressed is its *content*. Ayer's emotivism fails because it mistakes the attitude expressed by an ethical word for its content.

The challenge posed by the Frege–Geach problem is to explain the meanings of words in embedded contexts without assuming that they express the attitudes they do in freestanding contexts. In this light, the force of Horwich's proposal is brought into relief: For Horwich does not *identify* the meaning of a word with the attitudes expressed by freestanding occurrences; rather, the fact that freestanding occurrences express the attitudes they do *constitutes*, in part, the meaning the word has in embedded contexts. Horwich describes expressivist proposals, like Blackburn's, that endeavor to give separate explanations for the meanings of words in different contexts, *concessive*. Such concessive proposals must give some guarantee that the differently determined meanings in freestanding and embedded contexts are the same. However, Horwich observes that if the central positive claim of expressivism is understood as a claim about constitution, separate explanations need not be given. The meaning of 'ought' even in embedded contexts is explained, in part, by its use in freestanding contexts to express pro-attitudes. Horwich's proposal thus closely parallel's his (1998) minimalism about truth. Just as the success of minimalism depends on explaining the meaning of the truth predicate in contexts like 'Everything Edgar says is true' in terms of the truth predicate's role in the T-equivalence along with other linguistic facts; so the success of nonconcessive expressivism depends on explaining the meaning 'ought' in embedded contexts in terms of its freestanding use to express pro-attitudes along with other linguistic facts.

Horwich denies that the Frege–Geach problem is a good objection and thus spends a lot of time diagnosing why theorists have been taken in by it. Whether or not it is a good objection depends on what its proper object is. In contrast to Horwich, I think that it is good objection to a natural form that expressivist accounts may take, and its enduring challenge to the expressivist is to give accounts of meaning that do not take this natural if problematic form. This might seem like a mere difference in emphasis, but it is a difference in emphasis that affects the dialectical situation: From Horwich's perspective, it is natural to focus on establishing the possibility of a nonconcessive expressivism; from the perspective I recommend, it is natural instead to focus on actually giving these nonconcessive expressivist explanations. Indeed, subsequent discussion suggests that actually giving such explanations might be the only way of establishing their possibility.

2 'Hiyo!'

Dreier's (1996) 'hiyo' example purports to establish that not every expressive function that can be assigned to a word could be meaning-constituting. The challenge for Horwich is to explain why the expressive function of 'ought' could be meaning-constituting whereas the expressive function of 'hiyo' could not.

One may accost a person by uttering 'Hi' or 'Yo'. Perhaps the speech act of accosting could be used to constitute the meaning of a novel predicate, 'hiyo'. Uttering a sentence containing the predicate 'hiyo' is stipulated to perform the expressive function of accosting. Given this stipulation, a competent speaker accepts 'x is hiyo' just in case he intends to accost x. Does this expressive function along with other linguistic facts suffice to explain its meaning? Suppose the meaning of the material conditional is constituted by the rules of inference governing it. Would this along with the expressive function of 'hiyo' in freestanding contexts suffice to explain the meaning of 'hiyo' when it occurs embedded in the conditional? Could the meaning, so explained, render valid the inference from 'If a dingo is nearby, then Bob is hiyo' and 'A dingo is nearby' to 'Bob is hiyo'? Do you know what is meant by accepting the conclusion of this argument?

According to Dreier (1996: 43):

> No, you don't. It doesn't mean anything intelligible. And the reason is obvious—there is no sense to be made of the idea that someone might use that conditional to infer 'Bob is hiyo' from 'A dingo is near'. It is obvious that the idea of inferring is out of place when the conclusion is a speech act of accosting.

Horwich agrees that 'hiyo' manages to mean nothing intelligible but seems otherwise to have a diametrically opposed reaction to the thought experiment:

> Could it be that the trouble with 'hiyo' rule is simply that it would not put us in a position to deploy 'Bob is hiyo' in *inference*—to rely on it as a premise or to derive it as a conclusion? But if that is a problem, it remains to see why. For isn't it up to us to decide whether we are going to accept 'Bob is hiyo of the moon is blue' on the basis of 'Bob is hiyo'? What's to stop us deciding to apply the usual deduction rules to 'hiyo' sentences? (Horwich, ms)

Horwich is right, I think, at least to this extent: There could be a language game involving 'hiyo' with rules of acceptance that would license accepting 'Bob is hiyo or the moon is blue' if one accepted 'Bob is hiyo'. The question is whether this is anything other than an empty formalism. More specifically, a legitimate worry may be raised about whether the licensed linguistic transformation counts as genuine inference. Dreier claims it obviously doesn't, but perhaps it is worth spelling this out, at least partially and metaphorically. When a person infers 'GÃdel is a fraud' from 'Schmidt is the author of the incompleteness theorem' and 'If Schmidt is the author of the incompleteness theorem, then GÃdel is a fraud', they do so because the truth of the premises rules out the possibility of the conclusion being false, and so rules out accepting the premises while rejecting the conclusion. The fundamental problem is that ruling out seems out of place with accosting, and hence the linguistic transformation licensed by the 'hiyo' language game could not be an inference no matter how much it formally resembles one.

('Ruling out' is understood as a potentially nonliteral description of a genuine phenomenon just as Wollheim's, 1980, talk of 'seeing in' is understood as a nonliteral description of a genuine phenomenon. It is the conditions that make talk of 'ruling out' apt and not

necessarily the representationalist order of explanation, i.e. explaining inference in terms of ruling out possible truth-determining circumstances, that I intend to convey by 'ruling out'. For more on the representationalist order of explanation and its alternatives see my 2001.)

I said that I *think* that Horwich is right at least to the extent that there could be a language game whose rules of acceptance license linguistic transformations that formally resemble genuine inference. The reason for the hedge is this: Whether or not this is so depends on how much is built into the notion of acceptance. If acceptance is just the terminal state of a move in an arbitrary language game, then I see no obstacle to the possibility of a language game whose rules of acceptance license linguistic transformations that formally resemble genuine inference. For example, there could be a 'Jaberwocky'-inspired language game governed by rules of acceptance such that if a competent speaker accepts 'If 'tis brillig and the slithy toves gyre and gimble in the wabe, then all mimsy are the borogoves and the mome raths outgrabe' and ''Tis brillig and the slithy toves gyre and gimble in the wabe' then the speaker must also accept 'All mimsy are the borogoves and the mome raths outgrabe'. However, more substantive conceptions of acceptance would constitute an obstacle. Suppose acceptance must be understood to be coordinated with rejection. In rejecting a sentence a competent speaker rules out the possibility of its being true and so rules out the acceptance of that sentence. Unfortunately, it is precisely because ruling out seems incompatible with accosting that the linguistic transformation failed to count as genuine inference. And if ruling out is necessary for rejection, and acceptance must be understood along with rejection, then since ruling out seems incompatible with accosting, accosting could not be the basis for acceptance. We could not accept 'Bob is hiyo' on the basis of our intention to accost Bob, and so could not accept 'Bob is hiyo or the Moon is blue' on the basis of accepting 'Bob is hiyo'.

In introducing a novel language game, we are free to stipulate rules of acceptance that license linguistic transformations that formally resemble inference so long as acceptance is thin—so long as acceptance is conceived to be the terminal state of a move in an arbitrary language game. Horwich's remarks about inference are only coherently interpreted in terms of the this thin conception of acceptance. However, later on, Horwich invokes a more substantive conception of acceptance when he claims that if a competent speaker accepts a sentence then that speaker believes the content of the accepted sentence. Suppose belief must be understood to be coordinated with denial. Denial involves ruling out. If accepting a sentence is sufficient for believing the content of that sentence, then the possibility of ruling out is built into the very notion of acceptance. Given this more substantive conception of acceptance, in introducing a novel language game, we are not free to just stipulate rules of acceptance that license linguistic transformations. Some guarantee must be given that in conforming to these rules we have managed thereby to infer anything. (This is the lesson of Prior's, 1961, runabout inference ticket.) A nonconcessive explanation of the use of the novel word in embedded contexts would guarantee this, but it must be given to do so.

Horwich should not claim that there is no problem about inference, but rather, as I believe he intended, that the problem about inference is not the fundamental problem:

> The basic problem is that when someone genuinely accepts something and

> someone else accepts its negation there is a sense of substantive disagreement. There is a conflict, a clash, a feeling that the other person is somehow in bad shape. But the 'hiyo' rule would not give rise to these phenomena. (Horwich, ms)

The observations about the role that ruling out plays in inference substantiate Horwich's insight here. Substantive disagreement is only possible if ruling out is. The difference between the accosting expressed by 'hiyo' and the pro-attitudes expressed 'ought' is that competent speakers cannot disagree with an accosting but can disagree with the relevant pro-attitudes. The pro-attitudes expressed by 'ought' must display sufficient structural complexity to allow for the possibility of ruling out in a way that accosting does not. It is only in terms of these attitudes that the nonconcessive explanation of the meaning 'ought' in embedded contexts can consistently be given. (See Gibbard, 2003, chapter four for a similar suggestion.)

3 Disagreement

What is it about the nature of disagreement that can distinguish attitudes whose expression could be meaning-constituting from attitudes whose expression could not? What is it about the nature of belief that explains why contradictory beliefs result in substantive disagreement?

Horwich begins with the idea that contradictory beliefs practically conflict:

> The conflict associated with contradictory belief consists, I would suggest, in their tendency to engender contradictory *desires*. Suppose I want it to be the case that X occurs; this desire might be derived from fundamental (and thus fairly universal) desires in light of my belief *that p*; and so someone who denies *that p* is not unlikely to want X not to occur. So we can see how the divergent beliefs can easily correlate with a further tension. And we thereby see how we—for the sake of minimizing desires that clash with our own and thereby jeopardize our happiness—would like others to believe what we do. (Horwich, ms)

The further suggestion is that the pro-attitudes expressed by 'ought' could practically conflict in a way that accosting could not.

There is something to Horwich's idea that substantive disagreement fundamentally involves a practical conflict (I make a suggestion about what this might be in the final section); however, I doubt that he has succeeded in specifying the relevant kind of practical conflict.

First, I am unsure why accostings could not participate in practical conflicts of the kind Horwich describes. Suppose Edgar believes that a desire of his would be frustrated by Bernice's accosting Bob. This might prompt Edgar to desire that Bernice not accost Bob. Suppose that Bernice desires to accost Bob and intends to do so by uttering 'Bob is hiyo'. Isn't there a practical conflict here: Bernice desires that she accost Bob and Edgar desires that she not accost Bob. It is hard to see how conflicting desires could be the source of

substantive disagreement since they are present even in cases where, intuitively, there is no such disagreement

A related doubt can be raised about whether the pro-attitudes expressed by 'ought' must practically conflict in the right kind of way. As I have emphasized ruling out does not only operate interpersonally as it does in cases of substantive disagreement, it also operates intrapersonally as it does in cases of inference. This is manifest in the difference between a change of mind and a mere change of attitude. If the attitudes of a competent speaker change as the result of inference, the speaker has changed his mind. But not every change of attitude is a change of mind. Edgar might on a whim desire a cappuccino. If the whim passes and the desire ceases, Edgar has not changed his mind about the desirability of cappuccino, even though he has changed his attitudes towards a cappuccino. Edgar's earlier state of desiring a cappuccino on a whim does not rule out his later not desiring a cappuccino in the way that accepting the premises of an inference rules out accepting the negation of its conclusion. Suppose that Edgar comes to have a pro-attitude towards X. He might come to accept 'X ought to obtain' on this basis, and he might express this pro-attitude by uttering 'X ought to obtain'. Ordinarily, we would consider Edgar's coming to accept 'X ought to obtain' when he previously did not as a change of mind. But it is hard to see how, on Horwich's account, it must be.

The practical conflict, as Horwich conceives of it, has the structure of Stevenson's (1944) disagreement in attitude: A disagreement in attitude arises when (i) the disputants have incompatible attitudes towards the same object and (ii) at least one of them has a motive for changing the attitude of the other. There are intrapersonal conflicts that corresponds to interpersonal disagreements in attitude. Thus, for example, Edgar in his youth may approve of and take pride in a tattoo that later in life he may feel embarrassed about and regret. So a single person can have incompatible attitudes towards the same object at different times. That satisfies the first condition on a disagreement in attitude, but what about the second condition?

Of course, once Edgar regrets the tattoo, it is too late to change the earlier pro-attitude that led to his getting it in the first place. Is this incompatible with Edgar's later self having a motive to change the attitudes of his earlier self? We can have motives that are impossible to act on (as when, in a moment of despair, one wishes one were never born). While the impossibility of acting on the motive is incompatible with forming an intention on its basis, it might be compatible with having the motive nonetheless. Consider a duplicate of Edgar in a nearby possible world with backward causation. The duplicate's later regret may motivate him to prevent the earlier attitude. If Edgar's duplicate has the motive why doesn't Edgar? Of course Edgar's duplicate can form an intention that Edgar cannot but this is perfectly explicable (as is the duplicate's capacity to form different de re attitudes).

Suppose you remain unconvinced. Notice, all that the second condition requires is that at least one has the motive. Perhaps Edgar's earlier self can have a motive to change the attitude of a potential future self. Suppose when Edgar gets the tattoo, he does so with foresight: He recognizes that later he might become the kind of person for whom the tattoo would be an object of regret. This may fail to deter him from getting the tattoo if he lacks respect for the kind of person he might become. Indeed, Edgar might conceive of the

tattoo as a partial means of not becoming such a person, thinking that if he ever does regret the tattoo, he will deserve whatever distress he experiences for coming to have that kind of character. Edgar's earlier self could be motivated to change the attitude of his later self, even if Edgar's later self could not be motivated to change the attitude of his earlier self; and so the second condition is satisfied along with the first.

So disagreements in attitudes can be intrapersonal as well as interpersonal. The problem is that not all intrapersonal cases of disagreement in attitude can be described as cases of substantive disagreement: Some are mere changes of attitude rather than changes of mind. So a practical conflict having the structure of a disagreement in attitude is not sufficient to give rise to substantive disagreement, ruling out, and the like. Gibbard (2003: 70) cites instances of Nagel's (1970) practical solipsism as examples of intrapersonal disagreement in attitude that are not changes of mind, but, less exotic examples will do. Think of the way taste changes over time (in the sense associated with the *de gustibus* motto if not in the sense of 'taste' whose standard Hume sought to establish). In his youth, Edgar listens to MotÃrhead, in middle age, he listens to Schnittke. Edgar's later self need not disagree with his former self over the merits of MotÃrhead even if it is no longer to his taste and even if, in his youth, he vowed to keep on rocking, and never, ever stop. So the mere fact that the pro-attitudes expressed by 'ought' figure in a disagreement in attitude in a single person over time is insufficient for the change of attitude to be a change of mind. Intrapersonal disagreement in attitude is insufficient for inference.

4 Negation

Consider now a minor infelicity that might be a symptom of a deeper problem. Suppose Edgar's belief about Mexican winters prompts a desire to visit Mexico in winter. Suppose Bernice denies that Mexican winters are as Edgar believes them to be. Bernice would not be prompted to desire visiting Mexico in the way that Edgar is. But it doesn't seem true to claim that she would desire to not visit Mexico in winter. Perhaps, she would simply be indifferent to visiting Mexico. It is not true that someone who denies that p is not unlikely to want X not to occur; rather someone who denies that p is not unlikely to not want X to occur. This seems like a slip, but it does structurally parallel a problem that Horwich faces in distinguishing not accepting a sentence from accepting the negation of that sentence.

Edgar believes that human activity caused global warming. Bernice and Emma do not. Bernice suspends judgment about whether humans caused global warming or even if global warming is taking place: She believes that the recorded increase in temperature is within the normal variation to be expected within a global epoch. Whereas Bernice suspends judgment about whether human activity caused global warming, Emma explicitly denies that it has. Edgar and Emma disagree. But do Edgar and Bernice?

It might be claimed that Edgar and Bernice disagree but not about whether human activity caused global warming. Edgar and Bernice disagree about whether the relevant observations are sufficient evidence to believe that human activity caused global warming. Edgar accepts that they are, and Bernice rejects that they are. But this is a disagreement about standards of evidence, not global warming. However, a case can be made that Edgar and Bernice disagree

about global warming as well as standards of evidence. When Edgar sincerely utters 'Human activity caused global warming' he expresses his acceptance of that sentence. More than that, he implicitly demands that his audience accept that sentence as well. (Horwich sees this demand as grounded in the practical conflict that a difference in belief gives rise to: 'So we can see how the divergent beliefs can easily correlate with a further tension. And we thereby see how we—for the sake of minimizing desires that clash with our own and thereby jeopardize our happiness—would like others to believe what we do.') Of course, Edgar's audience need not accede to his demand—they may in turn demand that Edgar give some reason for them to accept the sentence, and if he fails to do so Edgar is normally under conversational pressure to withdraw his utterance. So when Edgar utters 'Human activity caused global warming' he implicitly demands that Bernice accept that sentence. Bernice rejects this demand. Do they not then disagree? Don't Edgar and Bernice disagree about whether to accept 'Human activity caused global warming'? Wouldn't this give rise to a practical conflict—concerning, say, what kind of public policy initiative to support?

I do not mean to settle this issue. I mean only to distinguish two potential contrasts involved in talk of disagreement. Disagreement might be understood as the contrast between a competent speaker accepting a sentence and a competent speaker accepting its negation (this was the contrast between Edgar and Emma); or disagreement might be understood as the contrast between a competent speaker accepting a sentence and a competent speaker not accepting that sentence where not accepting a sentence is potentially broader than accepting the negation of that sentence (this was the contrast between Edgar and Bernice). In speaking of 'substantive disagreement' Horwich has in mind the former and not the latter contrast, but is he entitled to?

Consider Horwich's explanation of negation. The meaning of negation is constituted, in part, by the attitude expressed by accepting the negation of a sentence, i.e. the rejection of that sentence. So rejection is not explained in terms of the acceptance of the negation of a sentence; rather the acceptance of the negation of a sentence is explained in terms of its rejection. But what does rejection mean? Horwich faces a dilemma here: On the one hand, suppose that rejection is understood in such a way that Bernice rejects the acceptance of 'Human activity caused global warming' just like Emma. However, only Emma disagrees with Edgar in the intended sense. So understood, Horwich's expressivist semantics for negation would fail (since Bernice would count as accepting the negation of the sentence which she manifestly does not). On the other hand, suppose that rejection is understood in such a way that Emma but not Bernice rejects 'Human activity caused global warming'. Then while Horwich would be able to explain what it is for a competent speaker to accept the negation of a sentence, his account would be incomplete, for he would so far lack the means to explain what it means to not accept that sentence. We would know the state of mind Edgar is in when he accepts the sentence; and we would know the state of mind that Emma is in when she rejects the sentence; what we would not know is Bernice's state of mind when she neither accepts nor rejects the sentence but suspends judgment concerning it. Suspending judgment involves more than a failure to accept or reject a sentence (as when an otherwise competent speaker lacks the concepts to do so), it also involves a judgment that rules out both acceptance and rejection. But what state of mind is that? (Compare

Unwin's, 1999, 2001 criticisms of, respectively, Blackburn and Gibbard; see also Gibbard, 2003: 72-4, for discussion of this.)

Notice this dilemma structurally parallels the dilemma posed by the original Frege–Geach problem. And like the Frege–Geach problem, there is a natural Fregean diagnosis for it: The dilemma arises by running afoul of the attitude/content distinction. If the rejection of S involves ruling out an *attitude* (the acceptance of S), then Bernice would count as rejecting S just like Emma even though only Emma disagrees with Edgar in the intended sense. If the rejection of S involves ruling out the *content* of S, then only Emma and not Bernice would count as rejecting S as required. However, so understood, the rejection of S would be the acceptance of the negation of S. After all, the semantic function of negation just is to rule out the content of the embedded sentence. But to explain rejection of S as the acceptance of the negation of S is to abandon the expressivist explanation of negation.

This is not a decisive objection to nonconcessive expressivism; rather, like the Frege–Geach problem and Dreier's 'hiyo' example, it is a challenge to adequately formulate the central nonconcessive explanations. According to Horwich, the lesson of the Frege–Geach problem is to formulate the central positive claim of expressivism as a claim about constitution not identification. The lesson of Dreier's 'hiyo' example is that the attitudes whose expression are meaning-constituting must display sufficient structural complexity to allow for the possibility of ruling out. Horwich is right that conflicting attitudes with sufficient structural complexity will give rise to substantive disagreement. Perhaps he is right as well that substantive disagreement is fundamentally practical, but if it is, it is something more than a disagreement in attitude. The present difficulty further suggests that the structural complexity required for ruling out must involve the attitude/content distinction. Perhaps the expressivist can accommodate this. Just as a distinction can be drawn between believing and what's believed, perhaps a similar distinction can be drawn with respect to the relevant kind of noncognitive attitude. Thus for example, Gibbard (2003:72-4) distinguishes planning from what's planned. If plans are noncognitive contents, as Gibbard claims them to be, then planning might have the right kind of structural complexity. The point is that we now have a number of specific and substantive constraints on the nature and effects of the attitudes whose expression could be meaning-constituting. It is not obvious whether any noncognitive attitude could satisfy all these constraints just as it is not obvious that no noncognitive attitude could. The best way to establish the possibility of a nonconcessive expressivism is to give a detailed description of the attitude expressed and to argue that it satisfies all the relevant constraints. Establishing the possibility of a nonconcessive expressivism might involve actually giving the relevant nonconcessive explanations.

5 Noncognitivism, expressivism, and nonfactualism

One salutary feature of Horwich's discussion is the distinction he draws between expressivism and nonfactualism. After defending the possibility of expressivism, Horwich considers and rejects the suggestion that it might form the basis of a nonfactualist interpretation of the given discourse. This might seem odd if one is used to using these labels interchangeably, but a genuine distinction can be drawn here. Expressivism is primarily a positive claim about

the semantics of a given discourse, i.e. that the meaning of a word is partly constituted by its freestanding occurences expressing the relevant attitudes. In contrast, nonfactualism is primarily a negative claim about the semantics of the given discourse—that the meaning of the word is nonrepresentational, i.e. that sentences containing the word do not function to represent a special subject matter. These claims are conceptually distinct. While most expressivists have been nonfactualists, it is unclear why one could not be an expressivist without also being a nonfactualist (for relevant discussion see Rosen, 1998; Hawthorne, 2002; Gibbard, 2003; and Kalderon, 2005: chapter two).

Suppose that a nonfactualist about normative discourse denies that 'ought'-sentences express propositions that represent a practical modality. There would be, in some substantive sense, no normative belief because there would be no normative propositions to believe no matter how well normative acceptance otherwise mimics normative belief. Since 'ought'-sentences, according to the envisioned nonfactualist, do not function to represent a practical modality, their meaning must be constituted by some other function, say, the expression of pro-attitudes that can be individuated independently of normative propositions. Horwich rejects this familiar syndrome of ideas maintaining instead that 'even if the function of certain simple sentences is to express non-cognitive states, this does not preclude the possibility that they also (and simultaneously) express states of belief'. If utterances of freestanding 'ought'-sentences can express the relevant pro-attitudes and simultaneously express normative belief, then the expressivist semantics for normative discourse need not underwrite a nonfactualist interpretation for it.

I myself am skeptical about nonfactualism (see my 2005: chapter two). I believe that it is possible for utterances of freestanding 'ought'-sentences to simultaneously express noncognitive attitudes and normative beliefs; and indeed that it is possible that normative utterances express normative beliefs by expressing these noncognitive attitudes. (Perhaps this is implicit Butler's description of moral intuition as 'a perception of the heart and a sentiment of the understanding'.) I disagree, however, with Horwich's reasoning in support of these claims:

> On the contrary, the general form of a meaning-constituting property is (roughly) that we tend to accept '#w' in circumstances C (—for example, we tend to accept 'I ought to f' when we want to f). And there is nothing more to believing that p than accepting some sentence with the appropriate meaning. Thus it will be a trivial matter to say, even relative to a nonconcessive, non-cognitivist, use-theoretic account of w's meaning, which state is typically expressed by any declarative sentence, 'p', containing w. It typically expresses the belief that p.

The argument proves too much: If it were sound, it would preclude not only a nonfactualist understanding of the given discourse but a noncognitive understanding as well.

Just as expressivism and nonfactualism are primarily positive and negative claims about semantics, noncognitivism is primarily a claim about the nature and content of the attitudes involved in acceptance. Whereas the cognitivist claims that the acceptance involves belief in a proposition that represents the putative subject matter, the noncognitivist denies this—acceptance involves attitudes other than belief in a proposition that represents the

putative subject matter. This difference is partly a difference in the nature of the attitudes involved (whether they are cognitive or noncognitive) and partly a difference in the content of these attitudes (whether the beliefs involved, if any, take as their objects propositions that represent the putative subject matter).

Acceptance can be characterized in terms of its role in inquiry. Acceptance is both the object and grounds of inquiry. Acceptance is the object of inquiry in the sense that it is the terminal state of inquiry: If a competent speaker accepts S, he takes himself no longer have reason to inquire about S, and a competent speaker is justified in accepting S if he has sufficient reason to inquire no further. Acceptance is the grounds of inquiry in the sense that in accepting a sentence S, not only does a person no longer take himself to have a reason to continue to inquire about S, but he also relies on his acceptance of S as grounds for further theoretical and practical inquiry. (This latter aspect of acceptance is plausibly the source of Horwich's intuition that disagreement must involve a practical conflict.) Acceptance is the grounds of inquiry in the sense that a person relies on the acceptance of the sentence in theoretical and practical reasoning and takes himself to be justified in so doing. These two aspects of acceptance are related: A person could be said to accept a sentence only if he was prepared to rely on it in theoretical and practical reasoning over a wide range of contexts. Moreover, a person would be justified in relying on S in theoretical and practical reasoning if he were justified in accepting S, if he had sufficient reason to inquire no further.

So characterized it is an open question whether acceptance in a given domain essentially involves belief in a proposition that represents its putative subject matter or whether acceptance in that domain wholly involves other attitudes. Even if scientific realism were a correct description of scientific inquiry as it is actually conducted, there could be a scientific community of self-conscious constructive empiricists (see van Fraassen, 1980). There is nothing inconsistent or otherwise internally incoherent in imagining such a community, and hence no reason to believe that such a community is not genuinely possible. But notice that in the hypothetical community, the acceptance of a scientific theory does not involve belief in the unobservable structures described by that theory; rather, acceptance involves belief in that theory's empirical adequacy combined with a practical commitment to deploy that theory in the conduct of science. Scientific inquiry as conducted by the hypothetical community is noncognitive in the sense that the acceptance of a scientific theory consists wholly in attitudes other than belief in that theory. (Call it a weak noncognitivism, if you like, since it involves belief, if not beliefs about an unobservable domain, but it involves practical attitudes as well and even Ayer did not claim that ethical sentences were utterly devoid of factual content).

It is an open question whether acceptance in a given domain is cognitive in the sense defined above. It is easy to see why this so: Acceptance has been characterized in terms of its role in inquiry without making any cognitivist assumptions. However, if Horwich's argument were sound, the question would be closed. Nonconcessive expressivism is meant to be compatible with a cognitivist understanding of acceptance. I don't doubt that this is genuinely possible; I do doubt the grounds that Horwich advances for thinking that this is so. Whatever proves to be the case, progress in this area requires that we clearly distinguish the distinct claims of expressivism, nonfactualism, and noncognitivism. Whereas the former

are, respectively, positive and negative claims about semantics, the latter is a claim about the nature and content of the attitudes involved in acceptance.

Philosophy