
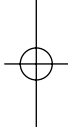




## 5

# Ethical Neo-Expressivism

*Dorit Bar-On and Matthew Chrisman*



A standard way to explain the connection between ethical claims and motivation is to say that these claims express motivational attitudes. Unless this connection is taken to be merely a matter of contingent psychological regularity, it may seem that there are only two options for understanding it. Either we can treat ethical claims as expressing propositions that entail something about the speaker's motivational attitudes (subjectivism), or we can treat ethical claims as nonpropositional and as having their semantic content constituted by the motivational attitudes they directly express (nonscognitivism). In this chapter, we argue that there is another option, which can be recognized once we see that there is no need to build the expression relation between ethical claims and motivational states of mind into the *semantic content* of ethical claims.

In articulating the third option, we try to capture what we think is worth preserving about the classical expressivist idea that ethical claims directly express motivational states, and separate it from the wrong semantic ideas with which it has traditionally been caught up. Doing so requires arguing for and deploying a distinction between claims considered as products—such as sentences—and claims considered as linguistic acts—such as utterances. In our view, the former are properly seen as standing in an expression relation to propositions, whereas the latter are properly seen as standing in an expression relation to mental states.

In the first section below, we use this act/product distinction to defend a “neo-expressivist” view of the way in which ethical claims express

For helpful feedback on earlier versions of this chapter, we'd like to thank the audience at the fourth annual metaethics workshop in Madison, Wisconsin (especially Daniel Boisvert, David Copp, Matthew Kramer, Stephen Finlay, Joshua Glasgow, and Mark Schroeder), the 2006 metaethics reading group at UNC-Chapel Hill (especially Ben Bramble, Geoffrey Sayre-McCord, and Teemu Toppinen), Collin Allen, Simon Blackburn, Peter Hanks, Grant Goodrich, Mitchell Green, Andrew McAninch, Ram Neta, Dean Pettit, Michael Ridge, and Carol Voeller.



motivational states. The core idea is to argue that ethical claims considered as acts of claim-making should be seen as expressing motivational states of mind, since this best serves the point and purpose of ethical thought and discourse. However, we insist that this idea should be kept separate from any thesis about the semantic content of ethical claims considered as products. We maintain that it is best to think of the products of ethical claims as expressing propositions. In addition to presenting a defensible version of expressivism, ethical neo-expressivism, we argue, also gives foundation to a new (and in our opinion highly plausible) form of internalism. In the second section below, we discuss another family of views that deny that the connection between ethical claims and motivational states is a matter of the literal semantic content of ethical claims (in the sense of being part of what such claims *say*). In opposition to the classical subjectivists, who maintain that ethical claims express propositions that ascribe motivational states to their speakers as part of their semantic content, Copp (2001) and Finlay (2004, 2005) have separately argued that we should instead see such propositions as only implicated, rather than directly asserted, by ethical claims—either conventionally (Copp) or conversationally (Finlay). We consider these “neo-subjectivist” views (as we call them) to be the principal competitors to neo-expressivism, and so we offer reasons for favoring neo-expressivism. In the final section below, we flesh out ethical neo-expressivism by addressing a number of potential objections.

#### NEO-EXPRESSIVISM

On a popular Humean view of motivation, being motivated to act is never merely a matter of having certain beliefs but also always requires the presence of some conative attitude. Given this account of motivation, if we are to regard an agent’s ethical claim as internally connected to her motivation, we have to see it as somehow ‘betraying’ the presence of conative attitudes. Perhaps Humeanism is the wrong view of the psychology of motivation (more on this in the final section below); but assuming that it isn’t, every view of moral discourse will want to capture at least *some* sense in which ethical claims betray conative attitudes.

Classical subjectivism offers one way to do this. The subjectivist argues that ethical claims *report* that the speaker has the relevant conative attitude. According to subjectivism, this is a feature of their semantic content. For example, on a crude form of subjectivism, “Tolerance is a virtue” just means that I (the speaker) approve of tolerance.

One of the original motivations for expressivism was to capture what is right about this view without succumbing to the obvious problem with it.

What the expressivist thinks is right is that sincerely made ethical claims betray the speaker's conative attitudes. The obvious problem is that, if they do so by reporting that the speaker has the relevant attitude, then it is wildly unclear how ethical disagreement is possible. Or, to put the point differently, it seems that we could only disagree with someone's ethical claim by saying "You don't feel that way." But this seems implausible. Whatever problems we may have with others' ethical views, that they are an instance of self-deception is not usually one of them. The expressivist's leading idea for avoiding this problem is to construe the way in which ethical claims betray a conative attitude not as a matter of reporting its presence but as a matter of *directly expressing* it.

Here is Ayer on the difference:

if I say, "Tolerance is a virtue," and someone answers, "You don't approve of it," he would, on the ordinary subjectivist theory, be contradicting me. On our theory, he would not be contradicting me, because, in saying that tolerance was a virtue, I should not be making any statement about my own feelings or about anything else. I should simply be evincing my feelings, which is not at all the same thing as saying that I have them. (1936/46: 109)

The subjectivist can, in a sense, agree with Ayer that conative attitudes are evinced by ethical claims. But he'll have to see this as, so to speak, an indirect consequence of what he thinks ethical claims assert; and it is precisely the subjectivist view of what ethical claims assert that gets the subjectivist into trouble. Ayer writes,

The distinction between the expression of feeling and the assertion of feeling is complicated by the fact that the assertion that one has a certain feeling often accompanies the expression of that feeling ... [However,] even if the assertion that one has a certain feeling always involves the expression of that feeling, the expression of a feeling assuredly does not always involve the assertion that one has it. And this is the important point to grasp in considering the distinction between our theory and the ordinary subjectivist theory. For whereas the subjectivist holds that ethical statements actually assert the existence of certain feelings, we hold that ethical statements are expressions ... of feelings which do not necessarily involve any assertions. (1936/46: 109–10)

We think there remains something of promise in Ayer's strategy for divorcing the way in which ethical claims express conative attitudes from what ethical claims assert. But it was precisely his view that (genuine and pure) ethical claims do not involve *any assertion at all* that led him to deny that (genuine and pure) ethical claims are truth-apt and to insist on a radically noncognitivist semantics for them, according to which their meaning, much like the meaning of expletives, is just a matter of what conative attitudes they express. In ordinary discourse, however, ethical claims exhibit undeniable

semantic continuities with ordinary descriptive claims. We often say things like “It’s true that tolerance is a virtue, but it’s also true that some things are over the line.” Moreover, we embed ethical claims in force-stripping contexts, e.g.: “If tolerance is a virtue, then one’s children should be raised to be tolerant” and “Either tolerance is a virtue or the doctrine of ‘turn the other cheek’ is a farce.” Ayer, it seems, is simply forced to explain away these appearances; and many philosophers have taken his (and others’) failure to do so to be the major stumbling block for ethical expressivism.

However, we think that contrasting Ayer’s noncognitivist semantics for ethical claims with the subjectivist’s implicitly reflexive semantics has encouraged a conflation in this debate that makes the expressivist’s leading idea seem much more problematic than it really is. To see this, it is helpful to consider a recent expressivist view in a different domain that is explicitly not a view about the semantic content of claims in that domain. This is the ‘neo-expressivist’ view of avowals defended in Bar-On (2004). Avowals—that is, ordinary first-person present-tense ascriptions of mental states, such as “I’m feeling anxious,” “I’m hoping that you’ll come to the party”—are distinctive for their security. It is usually misplaced to challenge an avowal or ask for the speaker’s reasons for thinking that it is true. However, avowals manifest undeniable semantic continuities with ordinary descriptive statements, such as “I am bleeding,” or “I am walking down the street.” We treat avowals as truth-evaluable, saying things such as “It’s true that I am feeling anxious, but I’ll pull through.” And we embed avowals in force-stripping contexts, saying things such as “If I am feeling anxious, you should be concerned.”

This may seem to militate against treating avowals as directly expressive of the underlying mental states they ascribe; however, that appearance rests on a conflation of different senses of ‘express.’ The first sense is:

**(a-expression)** the *action* sense: a *person* expresses a *mental state* by intentionally doing something.

For example, when I give you a hug, or say: “It’s so great to see you,” I express, in the action sense, joy at seeing you. In this sense, one can express dis/approval of *x* by saying (to others or to oneself): “I don’t/like *x*,” as well as by using a dis/approving verbal epithet for *x* or displaying nonverbal dis/approving behavior toward it, just as one can express dis/agreement with *p* by saying: “I dis/agree that *p*” as well as by saying simply “(not) *p*” or by nodding in dis/agreement. It is important to note that a-expression is a relation between a person and a *mental state*. It is also worth noting that the notion of a-expression requires that a person do *something* intentionally. It does not require that what one does intentionally is *express*. Though one *can* intentionally express a mental state—for example, by deciding and setting

out to give vent to a present emotion, instead of suppressing it, it seems that the more basic case is one in which a person gives spontaneous expression to a present state of hers by performing some intentional act, such as giving a hug or uttering a sentence, that doesn't have expression as its intentional aim. The second sense of 'express' is:

**(s-expression)** the *semantic* sense: for example, a *sentence* expresses an abstract proposition by being a (conventional) representation of it.

Thus the sentence "It's raining outside" expresses in the semantic sense the proposition that it is raining at time *t* outside place *p*. This is what this English sentence has in common with its counterpart in other languages, such that they mean the same thing. (A speaker who utters this sentence will typically be a-expressing the belief that it's raining at the time and place of her utterance.) We can also speak of a thought-token as expressing a proposition (and a thinker who produces the token will typically be a-expressing the relevant belief).<sup>1</sup> The important thing to bear in mind is that s-expression is a semantic relation, holding between linguistic expressions (and their analogues in thought) and their meanings, whereas a-expression is a relation between subjects and mental states.

The distinction between a-expression and s-expression lines up with a further distinction between the *act* of making a claim and the *product* of this act. The act of making a claim a-expresses certain mental states, whereas the product of this act s-expresses a particular proposition. And these distinctions can be applied to claims made both in language and in thought. An act of saying "This painting is interesting" produces a sentence token, whereas a mental act of thinking it has as its product a thought token, where both tokens s-express the proposition that a designated painting is interesting. In producing these tokens, a person may (only) be a-expressing her state of believing that the painting is interesting. She may (also) be a-expressing her state of feeling intrigued by the painting (or whatever).

<sup>1</sup> The distinction between expression in the action vs. semantic sense is due to Sellars (1969: 506–27). Sellars identifies a third sense of 'express': in the *causal* sense, an *utterance* or piece of outward behavior expresses an internal episode by being the culmination of a causal process beginning with that episode. For example, one's unintentional grimace or trembling hands may express, in the causal sense, one's feeling pain or nervousness, respectively. We will set aside expression in the causal sense in what follows. Bar-On (2004: chs. 6–8) makes extensive use of the threefold distinction, in order to capture the distinctive security of avowals. Kalderon (2005: 64) also deploys a distinction between conveying and expressing, which is similar to the distinction between what we are calling a-expression and s-expression when arguing that proponents of noncognitivist semantics are guilty of what he terms the 'pragmatic fallacy,' which is, in essence, the mistake of confusing s-expression and a-expression.

It is often uncritically assumed that what we are calling a-expression requires engaging in a *communicative* act, even if not speech. But we think this is wrong. Given the definition of a-expression, it seems to us that one can a-express one's affection to a beloved one by caressing her photo in the solitude of one's room, one can cuss to oneself without making a sound, one can think "This room is a mess!" as one walks into a child's room, giving (silent) vent to one's annoyance while suppressing both vocal and facial expressions, and one can reflect on the steps of an argument, expressing one's thoughts to oneself all along the way. In all of these cases, the a-expression/s-expression distinction applies since there is both an expressive act and a product of this act. For example, the thought: "This room is a mess!" s-expresses the proposition that the room is a mess, but for all that, in the act of thinking it, one may be a-expressing one's annoyance at the mess.<sup>2</sup>

Properly understood, the a/s-expression and the act/product distinctions can be accommodated by any number of views on the relationship between sentence meaning and mental states. In particular, suppose one adopts a roughly Gricean view of meaning. Using our terminology, the Gricean view is that we should explain what a sentence in a public language s-expresses in terms of what mental state one expresses with sincere utterances of this sentence. For instance, and roughly, the sentence "Grass is green" on this view, would be taken to s-express the proposition that grass is green in virtue of the fact that one conventionally a-expresses one's belief that grass is green with an utterance of this sentence. This may appear to collapse the distinction between s-expression and a-expression, by reducing the former to the latter. But this is a mistake. The Gricean account must presuppose, rather than reduce away, the notion of s-expression *as it applies to mental states*. So *at best* it would allow us to reduce s-expression *as it applies to items in a public language* to a-expression of items endowed with powers of s-expression all on their own. Moreover, even setting aside this point, it seems clear that a viable Gricean view would have to (and should be able to) accommodate the s/a-expression distinction. Take for example a declarative sentence in a public language—"There's a thief in the yard"—that, on the Gricean account, is linked by convention to the expression of a belief state with the relevant content. On a given occasion, a speaker may utter the sentence intentionally and comprehendingly while failing to have the relevant belief—say, if she is lying, or dissimulating. In such a case, we might say, the speaker expresses the belief that there's a thief in the yard without expressing *her* belief to that effect. (The distinction between expressing *a*

<sup>2</sup> For relevant discussion and further examples, see Bar-On (2004: esp. chs. 6 and 7) and Green (2007).

*mental state* and expressing *one's mental state* can then be seen as a first approximation to a Gricean equivalent of the distinction between *s-* and *a-expression*.<sup>3</sup>)

Now, consider a simple expressivist account of avowals, according to which an avowal such as “I am in pain” is just like a natural expression of it (say, a wince, or a grimace) in that it does not serve to *ascribe* pain to the subject who avows; it serves *only* directly to express the feeling of pain itself. This view is often thought to explain the distinctive security of avowals but at the cost of disrespecting the obvious semantic continuity between avowals and other mentalistic ascriptions. However, if we apply the act/product distinction here, we can generate a highly plausible explanation of the distinctive security of avowals that has the explanatory advantages of simple expressivism without its disadvantages. As acts, avowals, like acts of natural expression (such as wincing, or smiling) directly *a-express* the mental state avowed. Consequently, avowals, like natural expressions, enjoy protection from epistemic criticism (we would not challenge a sigh of relief, ask for reasons for it, etc.); however, unlike acts of natural expression, avowals have as their products sentence- (or thought-) tokens with genuine truth-conditions.

Here is not the place to explore further the neo-expressivist view of avowals.<sup>4</sup> Our point in bringing it up is rather this: it seems as though a parallel neo-expressivist treatment of ethical claims promises to retain Ayer's leading idea for understanding the way in which ethical claims betray their speakers' motivational attitudes without going in for the dubious noncognitivist semantics and the denial of semantic continuity that it entails. The key is that we can separate what speakers (or thinkers) *a-express* through their acts from what the products of these acts *s-express*. This makes room for a view about what mental states are *a-expressed* by a claim (act) that does not imply any particular view about the appropriate semantics for that claim (product). In the case of avowals, we saw that one can issue a claim that *s-expresses* a proposition about the avower (“I feel excited”) but where what the avower *a-expresses* is the self-ascribed state

<sup>3</sup> For relevant discussion of some of these points, see Bar-On (1995). And for a recent comprehensive development of the Gricean approach that would accommodate our distinctions, see Davis (2003). Thanks to Michael Ridge for prompting us to make the clarification in this paragraph. (For more on the Gricean account in relation to ours, see n. 5 below.) Since we do not take ourselves to be providing an account of the meaning of ethical claims, we will not take a stand here on the appropriate semantic theory for ethical sentences. Suffice it to say that the Lockean/Gricean approach is not without viable rivals in this as in other cases. One well-known alternative is the Davidsonian view. (And see Chrisman (2008) for the beginnings of an inferential role approach to the meaning of ethical claims.)

<sup>4</sup> The view is developed in Bar-On (2004; see esp. chs. 6–8).

itself (excitement), rather than, or perhaps in addition to, her self-ascriptive belief. One can also a-express a mental state (feeling joy at seeing one's addressee, or feeling intrigued by a painting) that *isn't* self-ascribed by the proposition that one's utterance s-expresses ("It is great to see you," "This painting is interesting," respectively).<sup>5</sup> Similarly, we can distinguish between the act of making an ethical claim and the product of this act. An ethical claim considered as a product—that is, a sentence- (or thought-) token that essentially employs ethical terms (or concepts)—can be said to s-express a proposition. It is this feature of ethical claims considered as products that can be used to explain their semantic continuity with ordinary, non-problematic descriptive claims.

However, as with avowals, the issue of what is s-expressed by an ethical claim shouldn't be taken to prejudice the issue of what mental state one who makes the claim characteristically a-expresses by it. Here, independently of any more specific view about the semantics of the products, we want to suggest that the original expressivist idea about what someone making an ethical claim a-expresses is, at least within a Humean framework, quite plausible. It's plausible precisely because it provides an explanation of the apparently 'internal' connection between sincerely making an ethical claim and having certain conative attitudes. On the expressivist view, people who issue ethical utterances characteristically a-express the very same states whose presence is required for understanding the necessary connection of such utterances to motivation. But they do so without their token sentences or thoughts s-expressing propositions that *self-ascribe* those states. Thus, as long as we're talking about expressive acts and working within a Humean framework, we can agree with Ayer that ethical claims betray conative

<sup>5</sup> What if one were to insist (along the lines of the Gricean approach mentioned earlier) that what the sentence "It is great to see you" s-expresses must be given in terms of the mental state one a-expresses (perhaps in context) by uttering it? We think there are powerful (Frege–Geach style) considerations that would require us to give at least *a* meaning assignment to the sentence that was a straight function of its compositional semantics (capitalizing on the rules governing "great," "see," "you," etc.). Moreover, we think that our understanding of what mental state a speaker a-expresses when uttering the sentence is partly parasitic on this meaning, rather than the other way around. (This could remain so even if one thought that the rational reconstruction of how the sentence *came* to have the linguistic meaning that it has requires reference to the mental states that speakers of the language have somehow converged on a-expressing. For relevant discussion, see Lewis (1970) and (1975), and see our earlier distinction between expressing a mental state and expressing *one's* mental state.) However, our present point can be taken to be the more modest one: that once we have the *a/s*-expression and the *act/product* distinctions in place we at least *can* separate what a given sentence 'linguistically says' or 'means' (= s-expresses) from what mental state a speaker uses the sentence to express (= a-expresses) on a given occasion. (See above, n. 3.)



attitudes, not because they report them, but because speakers who make the claims express them directly. However, given the distinction between *a*-expression and *s*-expression, accepting this does not force us to agree with Ayer's noncognitivist account of the meaning of ethical claims. For it seems open to us to maintain that ethical claims, considered as products, *s*-express truth-evaluable propositions.

Since it is modeled on the neo-expressivist view of avowals, we shall call this view of ethical claims *ethical neo-expressivism*. A bit more carefully stated, the view's two main tenets are these:

(a) Ethical claims understood as *products* are semantically continuous with ordinary descriptive claims in being truth-evaluable, embeddable in conditionals, negation, logical inferences, etc. This is because ethical claims *s*-express true or false propositions.

(b) Ethical claims understood as *acts* are different from ordinary descriptive claims in that an agent who felicitously makes an ethical claim (in speech or in thought) essentially *a*-expresses a motivational state—a state whose presence is normally required for the explanation of the agent's relevant actions. This—rather than any semantic feature of the propositions speakers communicate when making ethical claims—is what allows us to explain the 'internal' connection between making an ethical claim and being motivated to act in accordance with it.

Before moving on, it may prove helpful to elucidate each of these claims a bit further.

In our view, ethical claims can be made in speech or in thought; the notion of *a*-expression applies to expressive acts, and we see no reason to presume that all expressive acts have speech as their medium. Because of this, the neo-expressivist view of ethical claims, like the neo-expressivist view of avowals, is not a view about acts made exclusively in *speech*. The view is that agents who make ethical claims in speech *or* in thought produce sentence-tokens (in speech) or thought-tokens (without speech) that possess truth-evaluable content and thus *s*-express propositions. What propositions? For present purposes, the proposition *s*-expressed by an ethical claim can be specified disquotationally. For example, the claim: "Murder is wrong" (understood as product) can be taken to *s*-express the proposition that murder is wrong. In saying this, we're not committing to a disquotational theory of the *truth* of ethical (or other) claims. We're just prescinding from any paraphrastic semantic analysis thereof. Moreover, we are not suggesting that the meaning of *every* sentence is to be specified disquotationally. Many sentences (including ones that use ethical vocabulary) are semantically complex and do admit of

semantic analysis. What we are resisting, however, is the idea that all ethical sentences must admit of such analysis.<sup>6</sup>

We think that to say that ethical claims s-express propositions is to make a logico-semantic point concerning ethical discourse and reflection, which does not have any automatic implications regarding moral psychology or the nature of moral reality. Consider a nonmoral sentence (or thought): “John loves Mary.” Within a broadly Davidsonian semantic framework, the sentence (or thought) has as its truth-condition *that John loves Mary*. By contrast, “John and Jason love Mary” will have a non-disquotational truth-condition—to the effect that John loves Mary *and* Jason loves Mary. But assigning this truth-condition is consistent with any number of views on the nature of love and the worldly conditions that constitute John’s loving Mary. (Perhaps, indeed, there is no such thing as love; so “John loves Mary” will always be false.) This assignment of truth-conditions is also silent on the characteristic psychological states of those who go in for the use of the vocabulary of love. It may be interesting to speculate on the reasons human beings have come to possess the concept of love, or the role this concept plays in the cognitive economy and lives of creatures like us. But it’s very unclear how to translate the results of such investigation into an analysis of the literal meaning of sentences containing the word “love,” or the semantic content for propositions involving the concept. (For another case, consider a well-known metaphysically problematic case: sentences of the form ‘X caused Y,’ whose Davidsonian truth-conditions, even if not disquotational, will still not serve to unpack the meaning of ‘cause.’)

To be sure, if one thought that a proposition is true just in case it corresponds with a fact, and one thought that some (positive atomic) ethical sentences are correct because they express true propositions, then the neo-expressivist view would combine with these commitments to yield the realist view that there are ethical facts. However, we take the commitment to the correspondence account of truth and to the success-theory of ethical sentences to be open to debate within the philosophy of language and the metaphysics of morals, respectively. So, strictly speaking, neo-expressivism is neutral on those debates. We see this as an advantage of the view.

Crucially, unlike traditional expressivism, neo-expressivism does *not* offer a *meaning analysis* of ethical sentences in terms of the attitudes they express. Nonetheless, we think that our account can capture a suitable sense in which there is an internal connection between ethical claims and motivation to

<sup>6</sup> We discuss this issue further in our concluding section. Fodor (1998) offers powerful considerations in support of the view that very few, if any, atomic sentences admit of semantic analysis.

act. The connection is captured by maintaining that what is distinctive about *acts* of making ethical claims (in speech or in thought) lies precisely in their expressive import: such acts are essentially linked to the expression of motivational attitudes. The precise characterization of the link is a delicate matter to which we now turn.<sup>7</sup>

Although an ethical claim, considered as product, does not s-express a proposition that ascribes a motivational state or attitude to the person making the claim, we think that for a person to make a genuine ethical claim is for her to a-express such a state or attitude. (Insofar as one accepts the Humean explanation of motivation, this encourages the view that ethical claims a-express conative states.) In other words, if we want to locate an internal connection between ethical claims and motivation, we should look to the nature of acts of making ethical claims and the character of a-expression. Consider again the case of avowals. On the neo-expressivist account, there's a difference between a person who *avows* and a person who simply *reports* feeling disappointed. Though they may both say (or think) the same thing, namely: "I'm feeling disappointed," the person who avows is engaging in an act whose point is to express (vent, give voice to, show) the very state whose presence would make the avowal true, namely, the feeling of disappointment. By contrast, the reporter is only giving voice to her *belief* that she is disappointed. This characterization gives us a way of separating genuine avowals from non-avowing acts of self-ascription—notably, self-reports—*not* in terms of semantic content but rather in terms of the kind of act performed. We want to carry over this idea to the ethical case. We suggest that the connection to motivation that is thought by internalists to be an essential characteristic of ethical claims should be located in the kind of acts performed when making ethical claims rather than in their semantic content. The suggestion is that *for a person to make a genuinely ethical claim is for her to a-express a motivational state or attitude*, and so *having* this attitude is a condition on *properly issuing the claim*. Notice that, given the distinction between a-expressing a mental state M and a-expressing *one's* M, it is possible for one to make a genuine ethical claim and thereby a-express a motivational state *without* having this motivational state; however, our suggestion is that it is not possible to do so without violating one of the *propriety conditions* on making ethical claims. In such a case, one will be giving voice to a motivational state or attitude that one doesn't have, which, on the present suggestion, stands in the way of *properly* making an ethical claim.

<sup>7</sup> We wish to thank Dan Boisvert, Stephen Finlay, Josh Glasgow, Matthew Kramer, Mark Schroeder, and an anonymous editor for correspondence and discussion that helped us to clarify the characterization of our version of internalism that follows.

Denying that the connection to motivation is forged in the semantic content of ethical claims makes room for the following conceptual possibility: a person could issue a claim such as “Torturing cats is wrong,” or “It is good to rescue a drowning person” sincerely and while understanding what it says, yet without being motivated to act in accordance with it. Making theoretical room for this possibility is important, not only because it seems to be conceptually possible, but also because various cases that are thought to undermine motivational internalism—of psychopathic, or extremely lethargic, or satanic individuals—may well provide actual illustrations of this possibility. It seems to us implausible to legislate, along the lines of a strong motivational internalism, that such individuals are either insincere or fail to have mastered the semantic rules governing the sentences they use. On the other hand, we don’t believe that acknowledging this requires us to agree with a strong motivational externalism that holds that such individuals are *only* psychologically abnormal and/or morally deficient. We think that the neo-expressivist framework is attractive because it affords a more nuanced picture of the relationship between ethical claims and motivational states than is afforded by *either* subjectivism and classical expressivism, on the one hand, *or* externalist views, on the other. This framework also allows us to capture intuitive differences between the ways the link to motivation can be broken.

To see this, notice that on the neo-expressivist view we are proposing, the connection to motivation is forged in the character of ethical claims understood as acts (made in speech or in thought), rather than in the semantic content of ethical claims understood as products. As we saw, the neo-expressivist about avowals proposes that what is distinctive about avowals is that a person who avows does not *report* the presence of an occurrent mental state of hers, but rather is supposed to give vent, air, share, or give voice to the relevant mental state using an articulate linguistic (or language-like) vehicle that self-ascribes that mental state. This is not simply a generalization about what happens when people avow; rather it is a characterization of a certain category of acts, avowals, in terms of the *point* of avowing, which distinguishes it from other kinds of acts that may be performed using mental self-ascriptions, and has implications for their *proper* performance. Somewhat similarly, ethical neo-expressivism makes room for the following internalist claim: that what is distinctive about ethical claims—what renders them *ethical* claims—is the fact that a person who issues an ethical claim is supposed to give voice to a (type of) motivational state using a linguistic (or language-like) vehicle that involves ethical terms or concepts. This too is not offered simply as a generalization about what regularly happens when people issue ethical claims; rather it is a characterization of a certain category of acts—acts of

making ethical claims—in terms of their point, which distinguishes them from other kinds of claim-making acts, and has implications for their proper performance.

Issuing a claim is an intentional act; issuing an ethical claim, according to ethical neo-expressivism, is issuing a claim whose point is in part to give voice to a motivational state. But, unlike traditional expressivism and subjectivism, ethical neo-expressivism allows that the *vehicle* of expression—the sentence (or sentence-like) token used to express the motivational state—has truth-conditions and is not itself a *self-ascription* of the motivational state. (As we mentioned, the view is neutral on what *satisfies* the truth-conditions.) Of course, intentional acts can fail, and expressive acts are no exception. We think that one advantage of locating the expressive import of ethical claims in their character as acts is that we can recognize a varied range of *expressive failures*, where an (apparently) ethical claim can be issued in the absence of the relevant motivational states. This, we believe, is the key to articulating a more nuanced view of the connection between ethical claims and motivation than is afforded by either an implausibly strong internalism or an equally implausibly strong externalism.

It is agreed on all hands that the possibility of issuing ethical claims insincerely does not refute internalism. That a person can insincerely say “It’s good to help the poor” while lacking all motivation to help the poor does not show that there is no internal connection between making an ethical claim and being at least somewhat motivated to act in accordance with it. But why not? The answer can’t be that the insincere person says what she believes to be false. For one thing, traditional expressivists deny that ethical claims express beliefs in the first place; they locate the internal connection to motivation precisely in the fact that ethical claims express conative states rather than cognitive (belief-like) states. But then what is the expressivist to make of the failure of insincere ethical claims to express the relevant conative states?

We’d like to suggest that a better way to understand insincere ethical claims, one that could accommodate noncognitivist as well as cognitivist construals of such claims, is along the lines of a certain model of insincere promises. The person who says insincerely “I promise to take you to dinner” isn’t typically thought to have said what she believes to be false. Rather, she is thought to have made an *empty* promise, as we say. This is because she has failed to fulfill a *different* propriety condition on making promises, namely, having the intention to do as the promise says. Her promise fails to ‘come from’ the relevant intention. To the extent that the promise comes from the *wrong* intention—to mislead, for example—her act may still be criticizable. And there may be good reasons to insist that she *ought* to do as the promise says. But we can still recognize

the act as disingenuous. Similarly, we could think of the insincere person who says “It’s good to help the poor” as having *failed to properly make an ethical claim*, not because she doesn’t believe (or understand) what she has said, but rather because she (knowingly) fails to meet another propriety condition on making such claims, namely having the appropriate motivational state.

The case of avowals offers a similar model for understanding expressive mismatches that result in a failure to meet the propriety conditions for a specific kind of expressive act. A person can say insincerely “I feel awful” without feeling awful, perhaps just to gain sympathy. Such a person, we may suppose, understands what she has said, and the sentence she utters is a genuine self-ascription of a mental state. But given the description just given, it is reasonable to maintain that this person has failed properly to issue an avowal (indeed, it may even be suggested that she’s only issued a *pretend-* or *mock-*avowal) though she has said of herself that she feels awful, she hasn’t ‘spoken from’ the relevant mental state. But now consider a different case. Having undergone psychoanalysis, a person may conclude: “I feel intimidated by my sister.” The self-ascription in this case is sincere—it represents the person’s belief about her state, one she has acquired on the basis of her therapist’s testimony. Here it seems very reasonable simply to deny that this person has issued a genuine avowal at all. What the person has issued is an evidential self-report, albeit using a sentence that could have been used in issuing an avowal.

We have seen that the possibility of saying “I promise” without the right intention doesn’t show that there isn’t an internal connection between making a promise and having an intention to do as the promise says. We have also seen that the possibility of saying “I am in M” without expressing M doesn’t show that there isn’t an internal connection between avowing and expressing the avowed state. We now want to argue that the possibility of issuing a claim such as “It’s wrong to torture cats” in the absence of the relevant motivational state, even when done sincerely, doesn’t show that there isn’t an internal connection between making an ethical claim and a-expressing such a state. It is open to the internalist to maintain that the person in such a case fails to be *properly* making an ethical claim *because she fails to a-express the relevant motivational state of hers*, since a-expressing one’s motivational state is a condition on proper performance of acts of making ethical claims.

So the idea is that one of the norms governing the proper performance of acts of making ethical claims is that a person making such a claim is giving voice to the relevant motivational state. If a person says something such as “Stealing is wrong” but lacks the relevant motivational state, she may be sincerely and comprehendingly issuing a claim with the same semantic

content as an ethical claim. But, on our view, she should not be counted as properly making an ethical claim. This could be because she fails to recognize the point of making ethical claims, which need not mean that she has failed to master the relevant ethical concepts—she may be as adept as anyone at tracking the extensions of such concepts. (This may be a reasonable diagnosis of at least some versions of the case of the psychopath.) Or it could be because, though she recognizes the point of making ethical claims, she suffers some psychological impairment that stands in her way of *a*-expressing *her* motivational state (as may be the case with the lethargic). In these sorts of cases, the person issues a claim with the *semantic content* of an ethical claim in the absence of the right motivational state. And we may accept that she is issuing an ethical claim all right. But it is still open to the internalist to insist that she cannot be regarded as making the ethical claim *properly*, since she violates one of the propriety conditions governing the making of such claims.

The satanic case is interestingly different. In this sort of case, the person supposedly does recognize the point of making ethical claims, but she appears to be out to subvert it or deliberately flout it. Depending on the details of the case, we may want to assimilate it to the case of insincerely making an ethical claim or instead to one of the foregoing cases. Either way, the internalist can insist that the satanic issuer of an ethical claim fails to make such a claim *properly*.

The above diagnosis allows us to respect the internalist intuition that the psychopathic, lethargic, satanic, etc. are more than just psychologically and/or morally deviant. However, unlike the strong internalist, we do not regard this as failure to have mastered the semantic rules for the application of ethical terms/concepts or the semantic conventions underlying ethical language. On the neo-expressivist construal of the internal connection between making an ethical claim and being appropriately motivated, the link is not merely a matter of contingent psychological regularity, nor is it simply a condition on being a morally virtuous agent. At the same time, the view does not build the link into the *semantic content* of individual ethical claims, considered as products, nor does it portray it as a pre-condition on the very mastery of ethical concepts. Rather than exposing the internal link to motivation by offering a revealing, paraphrastic analysis of the *meaning* of individual ethical sentences, the neo-expressivist account proposes investigating the conditions on properly performed acts of making ethical claims. Since the general sort of link we are pointing to has some analogs, it will be useful to consider them.

One kind of example comes from the use of color vocabulary. Color concepts and predicates have extensions that can be tracked in a variety of ways. Suppose we come across an individual who is able to track the

extension of the predicate “looks blue” as well as any of us, but who lacks our ability to track colors through vision. In other words, such an individual systematically fails to have the sorts of visual experiences that underwrite the normal use of color concepts among us. Arguably, such an individual is not a full participant in our color discourse. If she utters: “This piece of paper looks blue” when asked for the color of a piece of paper in front of her, then even if we are willing to credit her with deploying the concept *looks blue*, we can still recognize her utterance as infelicitous or inappropriate. She has successfully identified an item within the extension of “looks blue.” But although she has managed to *say* (and think) truly *that the piece of paper looks blue*, she did not—indeed could not—properly issue a *visual report* about the color of the paper. There may be good reasons for crediting her with having mastered and mobilized the relevant color concepts. Still, arguably, *properly making a visual color report* (using the ‘looks’ locution) requires making it on the strength of a visual experience. Consequently, a person who under ordinary circumstances issues a claim about what color something looks to be in the absence of the relevant visual experience will be guilty of an impropriety, even if she issued the claim sincerely, and even if what she claimed (the product) is correct.

Perhaps a closer kind of example comes from acts of producing a gesture, facial expression, or turn of phrase that are conventionally associated with the expression of certain sentiments (e.g. shaking hands, tipping one’s hat, showing the ‘thumbs up,’ rolling one’s eyes, saying “Have a nice day,” “Pleased to meet you,” or “Sorry!”). It seems uncontroversial that there are conventional rules connecting the use of these diverse expressive vehicles with certain sentiments; the *point* of using them is to express the relevant sentiments. (In these cases, however, the relevant products—the utterances, gestures, or facial expressions—do not all *s-express* propositions that are *about* those sentiments. Gestures and facial expressions do not *s-express* anything. And “This is great!” is not *about* the expressed sentiment. Nor does it seem plausible that producing such expressions contextually *implies that* the speaker has the relevant sentiment, courtesy of the speaker’s intention to communicate that information.) That there are such rules is of course no guarantee that one could *only* use the relevant expressions when one has the relevant sentiments. Tipping one’s hat is used to express respect; but it may fail to express *my* respect. One can deliberately tip one’s hat to deceive or mislead one’s audience into thinking that one feels respect. But under deviant circumstances, one may tip one’s hat sincerely and *still* fail to express *one’s own* feeling of respect (just as one may sincerely emit “Ouch!” without actually feeling pain when primed to anticipate pain). It thus seems that the best way to think of the conventions as laying down norms for *proper* acts of producing the gestures, utterances, etc. Given



the conventional link, the actual presence of the relevant sentiments is required for the act to be *proper*. Thus, one who says: “Pleased to meet you” without feeling at all pleased to meet her addressee, or “Sorry” while having absolutely no regret for what she is apologizing about, is guilty of a certain kind of impropriety. And total absence of the correct sentiments will make it difficult to view someone as having full mastery of the point of the discourse in question.

Finally, consider someone’s saying: “That’s ugly” pointing to a sofa in a furniture store. The product of this act (the sentence token) ostensibly ascribes an aesthetic property (being ugly) to a certain item (an ostended sofa). But for all that, it is compelling to take the speaker (or thinker) who produces the sentence token to be giving voice to a (negative) aesthetic attitude. Suppose she proceeded to purchase the sofa, put it in her living room, proudly show it to her friends, or even added: “I really like it”; in other words, suppose it became amply evident that she felt no aesthetic repulsion toward the sofa whatsoever. Then even if we came to think of her original pronouncement as sincerely made, we would question its propriety. The sofa may well *be* ugly (whatever makes that the case), our speaker may well be very good at tracking ugly things and separating them from the non-ugly ones, so she may be credited with mastery of the relevant aesthetic concepts, and may also have no intention to deceive anyone when saying “That’s ugly.” Nonetheless, given the absence of any negative aesthetic attitude or sentiment it seems that ‘she has no business’ *using* the epithet ‘ugly.’ Her act seems not unlike that of someone who *puts on* a smile or a sad face.<sup>8</sup>

The neo-expressivist view invokes a similar sort of expressive failure by way of explaining the internal link between ethical claims and motivational states. When it comes to *ethical* discourse and thought, absence of the appropriate motivational state is presumably connected with *more* than just failure to be a fully competent participant in the discourse. It will often also attest to criticizable moral failures and/or lamentable psychological failures of various sorts. But the expressive failure is what matters to capturing an internal connection between ethical claims and motivational states. On the present proposal, the internal link to motivation is to be captured through the appeal to the requirements on proper engagement in ethical discourse and reflection. A person *could* sincerely say (or think) that, for example, it would be morally wrong not to help someone in trouble on a given occasion, and yet not feel in the least inclined to offer the needed help. This may not be conceptually impossible; nor can a proponent of the neo-expressivist

<sup>8</sup> Or, consider someone like Star Trek’s Data, incapable of normal human affective states, who learns to classify correctly some jokes as *funny* without ever feeling amused.

view insist that such a person cannot possibly understand what she is saying. Still, the person who says/thinks “This is the morally good thing to do” without any motivation to act accordingly will fail to properly issue an ethical claim just as the person who says “Congratulations!” without feeling in the least pleased for the addressee will fail to properly issue a congratulation or the person who says “This is beautiful” but has no positive aesthetic attitude toward the relevant item will fail to properly issue an aesthetic claim. For, in keeping with the version of internalism we have articulated here, the norms governing ethical discourse as well as reflection dictate that issuing genuine ethical claims (whether in speech or in thought) involves *a*-expressing motivational states or attitudes, and thus *properly* issuing such a claim requires the speaker (or thinker) to *have* the relevant state or attitude.

#### NEO-EXPRESSIVISM VS. NEO-SUBJECTIVISM

The strategy we have pursued for saving the expressivist insight while avoiding the core problem for noncognitivist semantics is, in effect, to enforce a separation between the semantic content of a claim considered as a product and the mental state that gets expressed by the person who makes the claim. There is, however, a more familiar distinction in the philosophy of language that might seem to be put to similar use. This is the distinction between asserted and implicated content. Often we can distinguish between the asserted content of what someone said and the content of what, due to linguistic conventions or conversational dynamics, they implicate. For one familiar example, a professor writing “Mr. Smith has good handwriting” in a letter of recommendation might be thought of as having asserted the proposition that Mr. Smith’s handwriting is good and implicated the proposition that Mr. Smith is not a good student. And, it has been suggested that, when someone says, for example: “The krauts are at it again” she might be thought to have asserted the proposition that the Germans are gearing up for war again and implicated the proposition that the speaker dislikes the Germans.

This distinction provides the basis for a rival strategy for capturing the internal connection between ethical claims and motivations to action. The basic idea is subjectivist, but not in the traditional sense. Traditional subjectivists capture the way in which ethical claims are expressively loaded and, in turn, the internal connection between ethical claims and motivations to action by taking ethical claims to have something like *I disapprove of x* as their asserted content. The proposal presently under discussion utilizes the distinction between asserted and implicated content to suggest that

ethical claims do not assert propositions such as *I disapprove of x* but rather implicate them. Because we take this ethical neo-subjectivism (as we might provocatively label it) to be the main rival to ethical neo-expressivism, we will consider it more carefully in this section and argue for the superiority of neo-expressivism.<sup>9</sup>

There are two ways to deploy the notion of implicature to develop a neo-subjectivist account. These are distinguished by two different kinds of implicature—conventional and conversational. In defense of the conventional-implicature view, David Copp writes:

it seems plausible to me that the term “morally wrong” has a coloring such that, other things being equal, a person who asserts that an action is “morally wrong,” using the term literally in asserting a moral proposition, *conventionally* implicates that she subscribes to a standard that prohibits the action ... This is explained by the fact that the use of “morally wrong” is governed by expressive linguistic conventions of the kind I have described. (2001: 36–7)

Since he sees subscription to a moral standard as a conative mental state, there is a sense in which, on Copp’s view, ethical claims express a conative mental state. He writes, “it is matter of linguistic convention that in asserting a basic moral proposition by uttering a sentence in which a moral term is used, a speaker ‘expresses’ a relevant conative state of mind, other things being equal” (2001: 14–15). However, this is not the direct sort of expression of a mental state that seems plausible in the case of avowals and that we used as our model in developing the neo-expressivist view of ethical claims above. Rather, it’s an indirect sort of expression which Copp defines as “Frege-expression”: “a speaker’s use of a term Frege-expresses a state of mind just in case it is a matter of linguistic convention governing the use of the term that other things being equal, if a speaker asserts a simple isolated subject-predicate sentence in which the term is used literally, the fact that the speaker used the term conveys that the speaker is in the state of mind” (2001: 20). The important point is that, unlike expressivists who think that a conative mental state is *directly* expressed by an ethical claim, Copp thinks that the conative state is conveyed courtesy of the fact that an ethical claim, due to the linguistic conventions of moral language, *implicates* the self-ascriptive proposition *I (the speaker) subscribe to such-and-such moral standards*.

<sup>9</sup> In what follows, we discuss Copp’s view as it’s presented in Copp (2001). In his contribution to this volume, which is in part a response to an earlier draft of our chapter, Copp clarifies and amends the account offered in the earlier paper, and addresses some of the difficulties we raise below (as well as ones raised by Finlay (2004, 2005), which we also briefly discuss below). For obvious reasons, we cannot here undertake a detailed response to Copp’s more recent paper (but see n. 14 below).

Stephen Finlay (2004, 2005) defends a similar view, but he uses the device of *conversational* implicature to explain “the fact that a speaker’s uttering ‘T is good’ is a sign *by itself*, independent of extra-linguistic information about social values, human motivational nature, etc., that she possesses (and expresses) corresponding motivational attitude” (2004: 222). Finlay’s view is that it is in virtue of the dynamics of conversations in which an ethical term occurs, rather than the linguistic conventions surrounding the term, that this internal connection between ethical claims and motivation should be explained. More specifically, he thinks that value terms such as ‘good’ are always indexed to interests or ends, but this index is often omitted in normal conversation because speaker and audience both typically know what the index is. He writes, “We are conversationally licensed to omit constituents of our assertions precisely when their presence and identity can be presupposed by our audience. Nonexplicitly indexed value judgments are thus conversationally appropriate just when the audience can be relied upon to presuppose or identify the desire-index without verbal directives” (2004: 217). So it is due to the conversational dynamics whereby we omit what can be assumed that Finlay thinks value claims can convey that their speaker has a particular motivational state. He writes,

In some paradigm contexts ... of value judgement, when speakers neglect to specify the relevant interests, their audiences are conversationally justified in assuming them to be motivated by the interests their speech acts presuppose ... They therefore express pragmatically that they are relevantly motivated, hence that they approve of T, which they judge to satisfy those interests. (2004: 219)

Again, the conative mental state is not directly expressed; rather it is expressed because the proposition (roughly) *I (the speaker) am in this conative state* is, due to the conversational dynamics, implicated.

What then is the proposition *asserted* on the neo-subjectivist view? Copp and Finlay take different lines. For a sentence such as ‘A is good’ Copp seems to think that the asserted content is (roughly) that the relevantly justified moral standard M positively values A, which he takes to be compatible with many different views (relativist and objectivist) of what the relevantly justified moral standard is. Finlay, on the other hand, argues for a relational analysis on which such a claim would assert (roughly) that moral standard M positively values A-ing, where the value of M is context-dependent (2004: 11–13).

Whatever their views on the propositional content of ethical claims, Finlay has an argument for why his way of working out the neo-subjectivist view of the connection between ethical claims and motivational states in terms of conversational implicature is superior to Copp’s way of doing it in terms of conventional implicature. Briefly, Finlay notes that a key

feature of conversational as opposed to conventional implicatures is that they are *cancelable*. And, as many examples typically used to argue against internalism suggest, what is implicated by ethical claims *is* cancelable. For instance, when Satan says, “Evil be thou my good,” arguably he cancels any implication that he is motivated to do what is good and avoid doing what is evil. If he were to go on and say, “Torturing babies is evil,” with a satanic grin, we wouldn’t suppose him to be implicating that he dislikes the torture of babies. And, more generally, Finlay argues that there are clear cases where we adopt another’s point of view in making value claims, and in so doing we cancel the implication that we have the relevant conative attitude. For example, one might be discussing a new Pope and his plans for the church, and, given an implicit understanding of his interests and convictions, when the question of gays in the clergy comes up, one might say, “Letting gays into the clergy would be a very bad thing.” Here it seems that the conversational setting cancels the implicature that the speaker is motivated to adhere to a moral standard that forbids letting gays into the clergy. For the speaker is presumably speaking from the point of view of the Pope.

Perhaps Copp can reply to this argument, but what is important to notice for our immediate purposes is that, if Finlay is right, this represents a major drawback of the neo-subjectivist strategy as compared to our neo-expressivist strategy. For conversational implicature typically requires a *conversation*, which means that Finlay’s view can explain only the way that ethical claims made *in conversation* are expressively loaded. Yet the expressive character of ethical claims and the attendant internal connection between making ethical claims and being motivated to act appears to cross the speech–thought divide. We don’t have to be engaged in a conversation with someone to express our ethical commendation or condemnation of some action; indeed, we don’t even have to be talking out loud. We can just *think* “Stealing is wrong,” and it seems that there will be just the same sort of internal connection to motivations not to steal. On our neo-expressivist view, the notion that ethical claims directly express underlying motivational states applies both to ethical discourse and to ethical reflection. And this, we think, represents an advantage of our view over Finlay’s. Of course, Finlay might argue that even in thought we can conceive of ourselves as, so to speak, conversing with ourselves or with an imaginary audience; and it is for the benefit of this audience that we conversationally implicate things in thought. However, though we don’t doubt that this happens, we do doubt that this is the normal case of thinking, namely making claims in thought.

Since conventional implicature clearly occurs in thought as well as speech, this may provide fodder for Copp to argue that the notion of conventional

implicature is a better way to work out the neo-subjectivist view after all. Moreover, perhaps, *pace* Finlay, cancelability does *not* speak in favor of the conversational-implicature view. Copp argues that the implicature carried by ethical claims is *not* cancelable in the strong sense adumbrated by Grice (1989: 44) and deployed in Finlay's argument.<sup>10</sup> However, he suggests that it is cancelable in a weaker sense that is consistent with conventional implicature. This is the sense in which one can use a statement *intelligibly* even when the implicature carried by the use of a contained term is canceled either explicitly or contextually, so "[i]t would not be self-contradictory, and it would be fully intelligible as an assertion despite the nonstandard use of the term" (2001: 18). What he seems to have in mind are cases where there is some recognizable conversational point to using the term, which typically triggers the implicature, in a nonstandard way, which on this occasion does not trigger the implication. For instance (to use Copp's own type of example involving 'coloring'), an American who feels no particular contempt for the English might mock his Australian friends by saying, "You can't stand the Pommies, and you love beating them at cricket. But at the end of the day you all are very similar; you play all of the same sports." Here, the contempt implicated by the standard use of the colored term 'Pommie' is canceled by the context.<sup>11</sup>

Does this provide a plausible explanation of the cases of the amoralist, Satan, or anyone who apparently sincerely and comprehendingly makes ethical claims while lacking the relevant motivational state? The idea would be to argue that these individuals use ethical terms in a nonstandard way that is nonetheless intelligible, given the conversational point of so doing. For instance, this seems to be a plausible account of the case where Satan says, "Torturing babies is evil" with a satanic grin. We remain unsure whether this move can be generalized to save the conventional implicature way of working out the neo-subjectivist view. But what is worth noticing is that, even if it does, this move *too* trades on features of *conversation*, thus exposing Copp's view to the difficulty we pressed above against Finlay's view. For, given that the amoralist, Satan, and others can make (even if not completely felicitously) ethical claims *in speech* without having the relevant motivational state, it seems that they should be equally well able to make such claims *in thought*. However, in thought there are not the same sorts of conversational points which may make nonstandard uses of a

<sup>10</sup> For Finlay's response see his (2005: 13–17).

<sup>11</sup> We're here following Copp in supposing that the phenomenon of coloring is properly explained through implicature. This is not obvious to us, however. It may well be that the use of derogatory or approbatory terms is *also* better explained within a neo-expressivist framework. But we cannot explore this issue here.

term intelligible. So, again, a view which captures the expressive character of ethical claims in a unified way across speech and thought is preferable. And the neo-expressivist strategy of construing ethical claims, whether made in speech or thought, as directly expressing a motivational state does just that.

Perhaps there is a way for Copp to meet this worry by just insisting that the implicature carried by ethical claims is not really cancelable.<sup>12</sup> We are sympathetic to something like this idea because we think that making ethical claims while failing to be in the relevant motivational state represents an impropriety: one cannot sincerely and *properly* make an ethical claim while failing to be in the relevant motivational state. However, we want to argue that, even if there is some way to save Copp's view from the criticisms pressed so far, the neo-expressivist view still has a distinctive advantage over neo-subjectivist views (of whatever kind).

This advantage traces back to the original debate between subjectivists and expressivists. In this debate, virtually everyone agreed that subjectivism is faulty because it cannot explain why saying "No, you don't dislike torture"—as opposed to "No it isn't wrong"—is not a cogent way to disagree with someone who says, "Torture is wrong." Expressivism was thought to be superior to subjectivism precisely because it can explain this. According to expressivism, in saying "Torture is wrong" one doesn't report a negative attitude towards torture; one directly expresses it. To disagree, one must express an attitude that conflicts with the negative attitude; that's why one's conversant says something like "No, it isn't wrong."

Now, it would seem that the neo-subjectivist view, which has it that the asserted content of an ethical claim is (roughly):

Action A conforms (or doesn't conform) to moral standard M<sup>13</sup>

avoids this problem with traditional subjectivism. For this proposition doesn't have anything to do with the speaker's attitudes, so to respond to a claim that asserts this proposition with "No, you don't" is clearly inappropriate. Moreover, this account of the asserted content of ethical claims provides for a clear sense in which someone who responds to "Torture is wrong" with "No, it isn't wrong" is disagreeing. The first claim asserts the proposition that torture doesn't conform to some moral standard M, while the second claim asserts the proposition that torture does conform to M. These are contradictory propositions.

<sup>12</sup> This seems to be his current view. See Copp's contribution to this volume (Ch. 6).

<sup>13</sup> Again, as we said above, there are several available accounts of what determines the value of this M. Copp suggests that it is a matter of the "relevantly justified standards," which is meant to be vague and neutral between many different accounts. Finlay suggests that it is a matter of the conversational dynamics in which the sentence is uttered.

However, in our view, this can only be a special case of disagreement between ethical claims and it is unclear whether it constitutes genuinely *ethical* disagreement at all. To see this, recall the example of talking about the Pope and his policies. There is surely a genuine question about whether it is bad according to Catholicism to allow gays in the clergy. You may think that it is, and I may think that it isn't. Where we are discussing things from the Pope's point of view, we might then disagree by saying, "Letting gays into the clergy would be very bad" and "No, it wouldn't be very bad." But that is no *ethical* disagreement; it's rather a straightforward factual disagreement about what is good or bad *according to the moral standard of Catholicism, which standard we may both reject*.

A genuinely ethical disagreement or, at least, the more standard sort of disagreement we encounter in ethical discourse is one where we accept different ethical standards (or moral principles) and because of this evaluate some action differently. For instance, the committed pacifist says "War is always wrong" and the skeptical hawk says "No, it's not always wrong." Here, in order to capture disagreement, the neo-subjectivist would say that what the pacifist is asserting is that war never conforms with some moral standard M, whereas the hawk is asserting that war does sometimes conform to M. But this is surely implausible. Rather, if we follow neo-subjectivists in thinking that the asserted propositions are of the form:

Action A conforms (or doesn't conform) to moral standard M

we should say that the pacifist is asserting that war doesn't conform to moral standard P and the hawk is asserting that war sometimes conforms to moral standard H. However, there's no disagreement in that. Nor is there disagreement in the putatively implicated propositions, respectively, that the one speaker is motivated to adhere to P and that the other is motivated to adhere to H. We've lost the apparent disagreement.

So neo-subjectivism appears to be subject to the same sort of difficulty that has beset classical subjectivism—the difficulty of explaining the full range of genuine ethical disagreement. By contrast, neo-expressivism has two avenues for locating potential disagreement. Like the neo-subjectivist, the neo-expressivist thinks that ethical sentences *s-express* propositions, and it is surely possible to use sentences which *s-express* contradictory propositions. But the neo-expressivist is not committed to construing ethical propositions as descriptive of ethical standards or codes. It is open to other construals, which would make it possible to have disagreement even between individuals who agree on what a given ethical standard prescribes. Moreover, for the neo-expressivist, there are also the motivational attitudes that are directly *a-expressed* by one who uses ethical sentences to make



ethical claims. These attitudes can stand in conflict, which represents a different sort of disagreement.<sup>14</sup>

<sup>14</sup> Strictly speaking, Copp's view about the proposition implicated by ethical claims could be neutral with respect to the proposition *s*-expressed by moral sentences. So, for instance, he could allow that moral claims about something being wrong express the proposition not that some action is forbidden by standard *M*, but that some action has the *sui generis* nonnatural property of wrongness. Or, he could adopt our disquotational tack from above and say that these claims express the proposition that the action in question is wrong. In either case, he could then claim to explain what we are calling genuine ethical disagreement at the level of asserted content rather than at the level of implicated proposition. However, insofar as his account is supposed to remain neutral about the proposition expressed by ethical claims, it won't yet explain the appearance of disagreement. Indeed, in his contribution to the present volume, Copp explicitly embraces this sort of neutrality, and retracts the epithet "realist expressivism" in favor of "cognitivist expressivism." But note, first, that this would allow him to accommodate 'cognitive' disagreement only on *some* construals of the content of ethical sentences; there will remain those construals (e.g. *A is forbidden by contextually determined standards M*) that will still not afford *ethical* disagreement. And, secondly, it looks as though cognitivist expressivism could at best accommodate the appearance of one sort of disagreement. For note that, on the amended view, it is still the case that the proposition implicated by an ethical claim (or rather, 'simplicated,' to use Copp's now-preferred term) is a self-ascriptive subjective proposition—a proposition about what one's motivational attitude is. But such propositions, viz. "I don't approve of murder" (s)implicated by one person vs. "I approve of murder" (s)implicated by another do *not* conflict. By contrast, the neo-expressivist view can still explain the appearance of conflict between e.g. "Murder is wrong" and "Murder is not wrong" along traditional expressivist lines, in terms of the *mental states a-expressed* by those making the ethical claims, *regardless* of how one construes the proposition *s*-expressed by these sentences.

The source of the remaining difficulty, as we see it, is that Copp is still working within a subjectivist framework insofar as the link he wishes to forge between ethical claims and motivation is made via a proposition *about the speaker's own mental states*, which the speaker communicates (albeit indirectly, by 'simplicating' it). We are told that the speaker who says "Murder is wrong" indirectly conveys the information ('simplicates') that she disapproves of murder whereas the one who says "Murder is not wrong" conveys that she doesn't disapprove of it. But from an expressivist perspective such as Ayer's (see quotation above p. 135), this does nothing to explain the disagreement in attitude *unless* we add that in conveying this information the two speakers directly *express conflicting attitudes* (as they would if, for example, they were to *avow* their attitudes in self-ascribing them). Copp helps himself to this idea by defining "Frege-expression" in such a way as to collapse Ayer's distinction between a speaker directly expressing a state of mind and asserting or even (s)implicating, *that* she has it. (To wit: "a speaker's use of a term Frege-expresses a state of mind just in case it is a matter of linguistic convention ... that the fact that the speaker used the term *conveys that* the speaker is in the state of mind" (2001: 20, emphasis added). In the contribution to the present volume, it is amply evident that Copp thinks we can move freely between the idea of a speaker expressing her disapproval and her conveying the information *that* she disapproves. However, it is not clear to us how simply appealing to the fact that speakers convey information about their motivational states *indirectly* (by 'simplicating' it), as opposed to directly stating it (as subjectivists would have it) helps capture the kind of ethical disagreement Ayer points to. Moreover (and relatedly), it is also not clear to us how Copp's 'cognitivist

So, in sum, we suspect that there are problems for either way of developing the neo-subjectivist view having to do with its dependence on *conversation* to explain the internal connection between ethical claims and motivation. However, even if these can be overcome, we think that the neo-subjectivist will still have problems accounting for the full range of ethical disagreement.

#### SUMMARY, OBJECTIONS, AND REPLIES

Ethical neo-expressivism is motivated by a desire to recognize the semantic continuities between ethical and nonethical claims while capitalizing on the expressivist explanation of the connection ethical claims bear to motivation. Semantic continuity can be preserved as long as ethical claims *qua products* are taken to s-express truth-evaluable propositions. The connection to motivation, on the other hand, is accommodated by considering ethical claims *qua acts* and supposing that when making such claims we a-express motivational attitudes. We think the view is sufficiently broad and neutral that it should be acceptable to opposing parties of many central debates in ethics and metaethics.

One standardly central issue we have elided is the question of what sort of mental states play the requisite motivational role in creatures like us. As far as we're concerned, there is an interesting and substantive debate here concerning *moral psychology*, a debate that is not to be automatically settled by one's views about what ethical claims regarded as products s-express. Much of this debate has proceeded as though Humeanism is true. If that's right, then the relevant states cannot be purely cognitive states such as beliefs, in which case it may very well be true that when making ethical claims we a-express *conative* states. These are the kind of states whose presence (in combination with the presence of beliefs) is invoked by the Humean in explanations of actions.<sup>15</sup> However, the important point is that, whatever one's account of moral motivation is, it need not rule out the possibility that the products of acts of making ethical claims still have propositional character and possess truth-evaluable content.

We have allowed that these products, ethical sentence- (or thought-) tokens, may not admit of any illuminating paraphrastic analysis, that ethical propositions may have any number of different kinds of 'truth-makers'

expressivism' can genuinely claim to do justice to the internalist link that the expressivist seeks to capture. However, spelling out these concerns goes beyond what we can do in this chapter.

<sup>15</sup> These explanations are the sort pursued in action theory, rather than the sort we might want from a fully physicalistic theory or from a social-political theory.

(or none), and that in acts of making ethical claims, speakers may a-express either cognitive or noncognitive states (or both). Thus the account of ethical claims that we have presented purports to remain relatively neutral on the nature of the *meaning* of ethical claims, of their *metaphysics*, and of the nature of the *motivational states* they express. How could a substantive (and successful) account of the connection between ethical claims and motivation remain neutral on these key issues? In the remainder of this section, we try to answer this question by addressing a few potential objections to the alleged neutrality of our account.

Ethical claims, we have said, should be taken to s-express truth-evaluable propositions. But which propositions do they express? What *do* ethical sentences say or mean? We think that decades of trying to offer revealing meaning analyses of ethical sentences should have taught us that the attempt may be futile. Indeed, if there's a clear lesson to be learned from Moore's (1903) Open Question argument, it is that there's little hope for reconstructing the meaning of ethical sentences *in other terms* (whether normative *or* nonnormative). Here we are in agreement with Blackburn who writes, "Moore's negative argument against identifying moral propositions in different terms remains convincing. I believe he was absolutely correct about the content of the moral proposition. It is what it is, and not another thing" (1998: 86). That is, there is not much to be said here as far as the *semantics of English* (or any other language) is concerned, at least not by way of offering a response to the question: what do competent speakers mean when they say (or think): "Murdering innocent people is morally wrong," "Helping the poor is morally good," "One ought not to torture animals," "Taxing the rich so little is unfair" (or their translations into other languages). Ethical sentences, however, are not unique in this regard. There is not much hope for a paraphrastic meaning analysis of "John loves Mary," or "Tom is a handsome fellow," "Grass is green," and even of "Energy equals mass times speed of light squared," "The wreckage was caused by the explosion," or "Whales are mammals." Even if we deny that there can be such an analysis, we can still point to conceptual or inferential connections between the relevant sentences and other sentences. So we are not left without things to say about the meaning of the relevant sentences. However, we *are* giving up on the possibility of *saying in other words* what a sentence in the given class is used to say. This is what we had in mind when suggesting that the proposition an ethical claim s-expresses should be specified disquotationally; that, for example, "Tolerance is a virtue" s-expresses the proposition that, well, tolerance is a virtue.

We do not maintain that meaning can be assigned disquotationally to *any* syntactically well-formed sentence (what with sentences such as "Colorless green ideas sleep furiously" or "Twas brillig and the slithy toves did gyre

and gimble” around). To say that we should assign meaning disquotationally is *not* to say that for a sentence to *have* meaning it is *sufficient* that it be possible to disquote it. *Meaningfulness* requires much more than that; but we would deny that it requires the possibility of providing a synonymous sentence, analytic paraphrase, or translation. Obviously more could be said here about deep issues in the philosophy of language; but we must leave further discussion for another occasion.

In addition to semantic neutrality, we want to allow that it makes sense to speak of an ethical sentence expressing a proposition prior to settling on the correct metaphysics for ethical discourse. Thus, we are assuming that it can make sense to say that by “Tolerance is a virtue” we mean *that* tolerance is a virtue, even though we may not (yet) know what, if anything, constitutes the nature or essence of tolerance, or virtue. We realize, however, that on some views of meaning, if there is *no* property for “tolerance” or “virtue” to denote, sentences involving these terms will lack all meaning (so there will be no proposition expressed by them and no thoughts they can properly be said to articulate). We find comfort in the fact that this view of meaning is controversial and will face great difficulties even outside the ethical context. We trust that if the difficulty with assigning meaning to “Phlogiston is combustible,” given the (perhaps necessary) non-existence of phlogiston, can be overcome, so can the difficulty with “Tolerance is a virtue.”

We think that a virtue of the view we are proposing is that it locates the link between ethical discourse and motivation in *acts* of making ethical claims rather than either in any aspect of their semantic content *or* in their truth-makers. This is what saves us from the traditional dilemma of either subscribing to non-cognitivist semantics, which Geach argued is hopeless, or embracing necessarily motivating properties, which Mackie argued is queer.

This brings us to another worry. If the point of making ethical claims is indeed as the neo-expressivist maintains, and if, moreover, it is allowed that there are no facts that make ethical claims true, why is it that ethical claims need to be understood as s-expressing propositions at all? Why not go back to Ayer’s expressivism? Put differently, if we begin with the (anti-realist) metaphysical assumption that there are no ethical facts for us to describe, then, even granting that ethical claims exhibit propositional form, why should we suppose that this form is anything more than merely superficial?

Of course, on one way of understanding the classic Frege–Geach problem, it’s necessary to see ethical sentences as s-expressing propositions in the same way as any other indicative sentence does in order to explain the logical interaction between the sentences that form *inter alia* valid *modus ponens* inferences. However, we think it is important to note here that

the issue of the propositionality of ethical sentences is orthogonal to the issue of what someone who utters these sentences a-expresses. For compare: utterances of “Yuck!” typically serve to a-express an utterer’s state of disgust. “Yuck,” it seems, is not a descriptive term whose function is to denote an observed property of things. For all that, “Yuck” can be replaced by utterances that have propositional form: “This is yucky!”<sup>16</sup> Now, having recognized the roots of “This is yucky” in the purely expressive “Yuck,” couldn’t we leave it at that, notwithstanding the semantic continuity of this sentence with sentences such as “This is 5 feet long”? Must we insist that it has to s-express a truth-evaluable proposition, especially given how implausible it is to think of the world as furnished with the property of yuckiness? Or think of other sentences whose characteristic use is to give voice to present states of mind: “It’s nice to meet you,” “I’m much obliged.” These too exhibit semantic continuity with ordinary descriptive claims such as “It’s hard to pick up this suitcase,” and “I’m walking down Franklin Street.” Despite their expressive function and origin, once possessed of propositional dress, they admit of significant negation, and can presumably be embedded in conditionals and other propositional contexts (including logical inferences).

But this just seems to be grist to our mill. Reflection on these other clear cases of *purely* expressive uses of language suggests that semantic continuity (and the features that give rise to most statements of the Frege–Geach problem) comes with acquired propositional form, and does not require descriptive origins or grounding in objective facts.<sup>17</sup> This is all the more reason why we should have recourse to a relatively innocuous notion of a proposition, one that does not require analytic paraphrase and does not commit us to ‘robust’ truth-makers. It is precisely this notion that we invoked when claiming that ethical claims s-express propositions, by way

<sup>16</sup> This replacement process seems to be very productive. And the expressive term need not always be replaced by an adjective; it can be replaced by a noun. One hears parents saying things like: “Do you have an owie?” In this context, it is also instructive to consider utterances such as: “This is so *funny!*” “*Great* to see you!” “I *hate* you!” “This is *gorgeous!*”, and so on, which can be thought as replacements of non-propositional expressions (and which, moreover, have distinctively expressive uses in thought, and not just in conversation).

<sup>17</sup> Although we don’t have space to discuss it fully here, we would say something similar about the sort of worry pressed by Dorr (2002) about the nature of the mental states involved in accepting an ethical claim. Dorr argues that one must agree that these states are belief-like states on pain of otherwise endorsing bad forms of inference such as wishful-thinking. However, even if they must be counted as beliefs in the sense that they are the proper kinds of mental states to reason to and from, this doesn’t imply that they have to be beliefs in the sense of being mental states that aim to represent/describe the objective features of the world. We can similarly reason to and from states of accepting propositions of the form ‘This is yucky.’

of avoiding the difficulties posed by semantic continuity for traditional expressivism.

This is not to deny that there may be an interesting story yet to be told about the emergence of ethical propositional content, *if one denies* that ethical terms/concepts are genuinely descriptive and serve to pick out objective properties of things (or even just ‘response-dependent’ properties). Here’s one way such a ‘genetic’ story might go. In keeping with the idea that ethical language and thought have their origins in expressive behavior, we can perhaps think of ethical terms such as “good/bad,” “right/wrong” as initially serving as verbal replacements for ‘reactive’ verbal emissions of the “Yum/Yuck” or “Boo/Hurray” variety.<sup>18</sup> We have seen that even an uncontroversially expressive term such as “Yuck” can acquire adjectival form and work its way into indicative sentences, much like color terms (which, at least on some views, begin their lives as externalizations of human visual responses). What could be the point of endowing an expressive term with a grammatical form that allows it to be used predicatively? Speculation: doing so affords what we may call ‘discursive negotiation.’ When we move from “Yuck” to “Yucky,” and thereafter to “X is yucky,” we are poised to draw others’ attention and attend ourselves to features of X relevant to our yuck-reactions; and our own ‘yuck’ reactions become ripe for intersubjective give-and-take. (Something like this may happen when terms like “ought” and “should” come to replace verbal “Do/Don’t” injunctions.)

In the case of “yuck” it may, however, be argued that the gain is purely verbal. Nothing will be lost if we gave up on the advantages of being able to say/think: “X is yucky.” This is not completely clear, since it seems as if we can, once they have taken on a propositional guise, embed our claims about what is yucky in complex logical constructions. Still, we recognize the intuition that moral discourse involves more than trade in subjective reactions. However, a neo-expressivist account can respect that intuition. For an expressivist may be able to locate a relevant difference between ethical discourse and yuckiness discourse in the function of the very states of mind to which (she claims) we initially give voice when using ethical terms. Ethical terms are, arguably, not merely associated with approbation

<sup>18</sup> For a story of this sort concerning the expressive character of avowals, see Bar-On (2004: ch. 7). For ‘quasi-genetic’ stories in the ethical case, see Gibbard (1990: ch. 3) and Blackburn (1998: ch. 3). Bar-On (1995) briefly discusses the philosophical status of such stories in connection with a ‘genetic’ interpretation of Grice’s account of nonnatural meaning. Whichever way one wants to spell out the details, we could imagine a stage at which ethical terms were introduced into language to replace specific sorts of expressive verbal emissions (in the style of Sellars’ (1956) “Myth of Jones”). We should emphasize that we are not ourselves committed to this particular story. We are here offering it only by way of forestalling objections to our claim that the neo-expressivist account can remain neutral in the ways we have indicated.

and disapprobation *tout court*, but with *commendation* and *condemnation*. Perhaps, then, we should think of the initial states of mind to which the use of these terms serves to give voice as rather more complex than states of disgust, enjoyment, dis/pleasure, etc. Expressing such states does not just play the role of displaying merely subjective reactions but rather serves to display reactions and behavioral tendencies that one expects others to share. Perhaps, for example, they are states that incorporate a person's perception or apprehension of potential more or less immediate consequences of a perceived act or conduct, and involve a person's desire for herself or others to alter the perceived situation or conduct, to promote or prevent them.<sup>19</sup> If so, then this may give rise to a difference in what happens at the predicative stage. Whereas "X is yucky" can (perhaps) continue to be interchanged with the purely expressive "Yuck" with no appreciable loss, arguably, "X is good/bad; right/wrong" and "You ought/not to do A" may not. Whereas gustatory discourse could remain at (or return to) a primitive emotive stage without much loss, ethical discourse has additional point and purpose, which is only fully served by a propositional discourse.<sup>20</sup>

Obviously, a lot more needs to be said. We think that we have said enough, however, to indicate in broad outline the kind of 'genetic' story one could try to tell to explain why ethical discourse naturally takes on propositional form while remaining expressive in its function. What is important for us as neo-expressivists is to insist that, even if one wants to tell such a story about ethical discourse, there is no reason to expect that we should be able to read the story off the *semantic* behavior of ethical sentences. Instead, we should locate whatever 'traces' are left by the 'genesis' of ethical discourse in the account of what those who make ethical claims a-express rather than in what the claims s-express.

Perhaps there are other views concerning the expressive character of ethical claims yet to be articulated. However, we think that the neo-expressivist view we have outlined here has considerable advantages over traditional forms of expressivism and subjectivism as well as over the neo-subjectivist views that have been defended in more recent years. The relative neutrality of ethical neo-expressivism allows it to avoid some of the main difficulties that beset other views of ethical discourse and thought. And this very same neutrality, we believe, should make it acceptable to members of otherwise opposing camps in the contemporary metaethical debate.


<sup>19</sup> This is in line with Gibbard (2003: 65), who argues that disagreement is key. You can't disagree with a state of disgust, but you can disagree with a moral commitment.

<sup>20</sup> The above suggestion was couched in terms of *discourse* that affords intersubjective negotiation. A fully adequate story would need to show us how to extend the notion of discursive negotiation from the intersubjective case of ethical discourse to the intrasubjective case of ethical thinking.

## REFERENCES

- Austin, John (1962) *How to Do Things with Words* (Cambridge, MA: Harvard University Press).
- Ayer, A. J. (1936/1946) *Language, Truth and Logic* (London: Gollancz).
- Bar-On, Dorit (1995) 'Reconstructing "Meaning": Grice and the Naturalization of Semantics' *Pacific Philosophical Quarterly* 76: 83–116.
- (2004) *Speaking My Mind: Expression and Self-Knowledge* (Oxford: Clarendon Press).
- Blackburn, Simon (1984) *Spreading the Word: Groundings in the Philosophy of Language* (Oxford: Clarendon Press).
- (1998) *Ruling Passions: A Theory of Practical Reasoning* (Oxford: Clarendon Press).
- Chrisman, Matthew (2008) 'Expressivism, Inferentialism, and Saving the Debate' *Philosophy and Phenomenological Research* 77(2).
- (Ms) 'Normative Thought and Inference'.
- Copp, David (2001) 'Realist-Expressivism: A Neglected Option for Moral Realism' *Social Philosophy and Policy* 18: 1–43.
- Davidson, Donald (1967) 'The Logical Form of Action Sentences,' in Nicholas Rescher (ed.), *The Logic of Decision and Action* (Pittsburgh: University of Pittsburgh Press).
- Davis, Wayne (2003) *Meaning, Expression and Thought* (Cambridge: Cambridge University Press).
- Dorr, Cian (2002) 'Non-Cognitivism and Wishful Thinking' *Nous* 36: 97–103.
- Finlay, Stephen (2004) 'The Conversational Practicality of Value Judgement' *Journal of Ethics* 8: 205–23.
- (2005) 'Value and Implicature' *Philosophers' Imprint* 5: 1–20.
- Fodor, Jerry (1998) *Concepts: Where Cognitive Science Went Wrong* (New York: Oxford University Press).
- Geach, P. T. (1965) 'Assertion' *Philosophical Review* 74: 449–65.
- Gibbard, Allan (1990) *Wise Choices, Apt Feelings: A Theory of Normative Judgment* (Cambridge, MA: Harvard University Press).
- (2003) *Thinking How to Live* (Cambridge, MA: Harvard University Press).
- Grice, H. P. (1989) *Studies in the Way of Words* (Cambridge, MA: Harvard University Press).
- Green, Mitch. (2007) *Self-Expression* (Oxford: Oxford University Press).
- Kalderon, Mark (2005) *Moral Fictionalism* (Oxford: Clarendon Press).
- Lewis, David (1970) 'General Semantics' *Synthese* 22: 18–67.
- (1975) 'Languages and Language,' in Keith Gunderstone (ed.), *Minnesota Studies in the Philosophy of Science* (Minneapolis: University of Minnesota Press).
- Moore, G. E. (1903) *Principia Ethica* (Cambridge: Cambridge University Press).
- Ridge, Michael (2006) 'Ecumenical Expressivism: Finessing Frege' *Ethics* 116(2): 302–36.





- Sellars, Wilfrid (1956) *Empiricism and the Philosophy of Mind* (Indianapolis: Bobbs-Merrill).
- (1969) 'Language as Thought and as Communication' *Philosophy and Phenomenological Research* 29: 506–27.
- Stevenson, Charles L. (1937) 'Emotive Meaning of Moral Terms' *Mind* 46: 14–31.

