(How) Is Ethical Neo-Expressivism a Hybrid View?

Dorit Bar-On, Matthew Chrisman, James Sias

1. Introduction

Mainline metaethical expressivists of the 1960s-1990s generally conceived of expressivism as a view in philosophical semantics. Roughly speaking, the idea is that the meaning of ethical sentences is to be given not in terms of what they represent or describe but rather in terms of the noncognitive attitudes they express.[1] But if this idea is combined with the traditional thought that the meanings of most other declarative sentences are given by the propositions they express, then it turns out to be impossible to give a systematic semantics for logically complex sentences with mixed ethical and non-ethical parts.[2] In response, many expressivists have endorsed an "ideationalist" conception of meaning across the board. That is, they suggest that *all* sentences mean what they do in virtue of the 'idea' (mental state type) they express; it's just that ethical sentences express a different kind of mental state from descriptive sentences.[3] Recently, a further epicycle of this debate has seen the articulation of various "hybrid" views which in some way seek to get the best of both of the views that ethical sentences express conative attitudes and that ethical sentences express beliefs by giving the meaning of ethical sentences in terms of both cognitive and conative states.[4]

---

[1] It's not obvious that this was how proto-expressivists like Carnap (1935), Ayer (1936/1946), and Stevenson (1937, 1944) thought of the view. However, under the influence of Hare (1952, 1963), who treated ethical sentences semantically as prescriptions and sought to align satisfaction conditions with ethical sentences as one might align truth conditions with descriptive sentences, the view became a semantic one. Many contemporary expressivists are quite explicit in locating (the expression of) conative attitudes in the semantics of ethical terms or sentences. See, for example, Gibbard (2003: 75). Even expressivism's most prominent critics typically conceive it as a semantic view. See, for example, Schroeder (2008: 87) and Jackson (2001: 10).

[2] This we take to be the main upshot of Geach's seminal 1965 paper.

[3] For a recent exposition of this idea, see Richard (forthcoming).

[4] It is important to note that, regardless of the innovations "hybrid" views purport to add to ethical expressivism, they nonetheless have tended to persist in conceiving of expressivism as first and foremost a take on the semantic content of ethical claims. According to Dan Boisvert's "Expressive-Assertivism," for instance, the claim "Tormenting the cat is bad" expresses both the belief that tormenting the cat has a certain (non speaker-

We find the initial ideationalist thought and its development in some hybrid views rather odd.[5] The main source of our puzzlement is that we find it odd to say that *sentences*—ethical or not—are in the business of expressing mental state types (whether cognitive, conative-motivational, or some suitable hybrid). We think the more natural and conservative idea is that sentences express propositions. At least this is often seen as the neutral framework in which to investigate how sentences are semantically composed and what it takes for good translations of a sentence in one language into other languages. Moreover, independently of metaethics, we'd be inclined to say that it's *acts* of making claims that express various mental states (or, better, that it's *people* making the claims who express their mental states). Appreciating these points in the context of trying to capture some special internal connection between ethical claims and motivation leads, we think, to something reasonably regarded as a hybrid expressivist view. But it's a view that – if it hybridizes anything – hybridizes the notion of expression, rather than the types of mental state semantically expressed by ethical sentences. This is the view that two of us previously defended under the label "Ethical Neo-Expressivism."[6]

According to this view, (declarative) ethical *sentences* do have *propositions* as their semantic contents—in this they are like any other (declarative) sentence. However, it is *acts* of

relative) property, and a negative attitude toward things with that property. The claim therefore has the following meaning, according to Boisvert: 'Tormenting the cat is F; boo for things that are F!' (2008: 172). Likewise, Ridge once described his "Ecumenical Expressivism" as "offering a systematic and unified *semantics* for both asserted and unasserted uses of normative predicates" (2007: 63, italics added; see also Ridge 2006). (Though more recently, Ridge, forthcoming, has moved to casting his hybrid expressivism as a view in *meta*-semantics, i.e., as a view about *that in virtue of which* claims have the semantic contents that they have. See Chrisman (forthcoming) for an argument that expressivism as a meta-semantic thesis is both a better version of expressivism and a plausible interpretation of the kind of quasi-realist proposal made in Blackburn 1984).

[5] It is especially odd in light of the widely received view among philosophers of language that the contemporary paradigm of ideationalist semantics, understood as a theory of meaning for natural language—Grice's theory—faces insurmountable difficulties. (For an attempt to develop a Gricean semantic theory that overcomes these difficulties, see Davis (2003).) For an alternative way of construing Grice's proposal (which is, however, consistent with a non-Gricean, propositionalist semantics), see Bar-On (1995).

[6] In Bar-On and Chrisman (2009).

making an ethical claim that are properly said to express mental states; and they do so 'directly' rather than by expressing somehow the proposition *that* one is in a particular mental state. Accordingly, the previous paper argued for two claims: (i) there are two different but antecedently quite plausible notions of *expression*, called s-expression (for 'semantic') and a-expression (for 'action');[7] (ii) if one wishes to capture a special internal connection between ethical claims and motivation in a way that accommodates the logico-semantic behavior of ethical sentences, it is best to articulate the idea that ethical claims express motivational attitudes in terms of the notion of a-expression rather than the notion of s-expression.

This raises two important questions:

- First, the traditional reason for denying that ethical sentences express propositions was that propositions determine truth-conditions, which are thought to be ways the world could be; so, if we reject all analyses of ethical terms in terms of natural properties, isn't anyone who says that ethical sentences express propositions (e.g. the neo-expressivist) committed to the Moorean conclusion that ethical sentences describe *sui generis* "nonnatural" ways the world could be?
- Second, if we reject the claim that motivational attitudes constitute any part of the literal or implicated meaning of ethical sentences, aren't we committed to denying the intuition behind motivational internalism after all? (Or, another way of putting it: isn't a-expression unsuitable for capturing internalism?)

We want to use this paper as an occasion to return to some of the main ideas in "Ethical Neo-Expressivism" in light of these questions. In particular, we want to explain in more detail the way in which ethical neo-expressivism adopts an appealing, metaphysically *neutral*

[7] The distinction takes its inspiration from Sellars (1969).

framework within which to think about the semantics of ethical sentences. In addition, we want to flesh out the notion of a-expression, explaining why we see it as more basic and more relevant to issues in metaethics than anything having to do with expressing or conveying the proposition that one is in a particular mental state. This will also put us in a position to address a number of critical questions we have encountered.

2. A Plea for Conservatism in Semantics

Independently of any metaethical debate, it's natural to say, of any declarative sentence 'S', that it expresses the proposition that-S (bracketing the subtle but, at present, irrelevant issues about context-sensitivity). The idea of a sentence's expressing a proposition provides an attractive framework for beginning to think about the literal meaning of sentences and several related issues.[8] This is for at least the following, connected reasons, all of which have been traditionally invoked in support of the need for invoking propositions.

First, the literal meaning of a sentence is surely preserved in a good translation of that sentence into another language. This is why it's natural to say that a sentence S1 in one language is a good translation of another sentence S2 in another language only if S1 and S2 express the same proposition. Expressing the same proposition may in some cases not be *sufficient* for good translation, but it's at least necessary.[9] So, for example, good translation cannot be achieved by

[8] The notion of 'proposition' we have in mind here is the ontologically neutral notion of *locutionary content*, rather than the more specific, and ontologically committal Fregean notion (viz. a Platonic *object* in a 'third realm'). In a model theoretic context, it might be represented as a function from the worlds postulated in a semantic model to truth values (or by some more complex function or set-theoretic object). In our view, whether using such theoretical models commits one ontologically is a further question beyond the question of whether such models are useful for understanding various compositional phenomena of the semantics of a language.

[9] Though see Bar-On (1993~~0~~) for some qualifications and complications that do not, however, bear on the issues of concern to us here.

simply preserving *truth-conditions*. Clearly, two sentences can have the same truth-conditions and not be good translations of one another.

Second, a declarative sentence-type can be tokened on its own in making an assertion, but it can also be tokened alongside other semantic elements, or embedded in various sentential contexts, and we need some generic way to keep track of the commonality of content across different contexts. For example, the sentence "It is raining" can be used to make an assertion, but it can also occur as the antecedent of a conditional or be appended to an epistemic modal – viz. "If it's raining, then I'll take an umbrella", and "It might be raining". What do all of these tokens of the sentence-type "It's raining" have in common? It may not tell us much, but it surely provides a convenient framework to say that they all share propositional content.

Third, declarative sentences can be used to articulate the object of various attitudes. For example, you may *suspect* that it is sunny in Edinburgh and I might *hope* that it is sunny in Edinburgh, in which case the sentence "It is sunny in Edinburgh" articulates what you suspect and what I hope. How does it do that? The straightforward answer is that it does that by expressing the proposition that it is sunny in Edinburgh, which is the object of both your suspicion and my hope.[10]

Now consider ethical sentences like "Charity is good," "Not giving to charity is wrong," and "Middle-class citizens are obligated to give to charity." Should we say that these sentences

[10] It doesn't seem plausible to us to suggest that the sentence "It's raining" expresses *the belief* (type) that it's raining without expressing the belief (token) of any particular individual. What is in common among attitude ascriptions like (i) John believes that it's raining, (ii) John hopes that it's raining, (iii) John fears that it's raining, (iv) John doubts that it's raining, and (v) John suspects that it's raining, it seems, is *some content*, and not a type of belief (or any other attitude) *with* that content. (Moreover, one might wonder: why single out *belief* as the relevant type of attitude that is held constant across attitudes?) What one wants is a notion of content that is *attitude neutral*, one that abstracts away from attitude type (as well as, relatedly, abstracting from speech-act type) – precisely something like the conventional notion of a proposition understood as Austinian locutionary content, or a Fregean thought (minus the Platonist ontological commitment). (Thanks here to Dean Pettit.) It is revealing that philosophers otherwise drawn to 'mentalist' conceptions of meaning have nonetheless found reason to resort to essentially abstract, and so in that sense non-mental notions of content (for recent examples, see Davis 2003 and Soames 2010).

also express propositions? Well, as far as syntax goes, these sentences are surely declarative sentences, and we can easily produce non-ethical sentences that would appear to have exactly the same logical form: e.g. "Charity is common," "Not giving to charity is legal," "Middle-class citizens are likely to give to charity." This suggests that like the non-ethical sentences, the ethical sentences fall under the generalization mentioned above that, for any declarative sentence 'S', 'S' expresses the proposition that-S (again, bracketing context-sensitivity). Moreover, ethical sentences in one language seem to admit of good translations into other languages just as much as non-ethical sentences. So insofar as natural criteria for good translation deploy the notion of two sentences' expressing the same proposition, we'll want that notion to apply in the ethical case just as much as in the non-ethical case. Likewise, it seems that ethical sentences can function just like non-ethical sentences in semantic embeddings and to articulate the objects of diverse attitudes. We can say, for instance, "Not giving to charity is wrong" but also "If not giving to charity is wrong, then I'll endeavor to give to charity" or "Not giving to charity might be wrong". Similarly, it seems that you may *suspect* that middle-class citizens are obligated to give to charity, while I *doubt* that middle-class citizens are obligated to give to charity, in which case the sentence "Middle-class citizens are obligated to give to charity" is a pretty good way to articulate what you suspect and what I doubt. Again, the straightforward account of how it does this is that this sentence expresses the proposition that middle-class citizens are obligated to give to charity.

Doesn't all of this speak strongly in favor of treating declarative *ethical* sentences just as we treat declarative *non*-ethical sentences, in terms of their expressing propositions? Obviously, yes. But, to be clear, we don't view a positive answer to this question as the end of philosophical semantics but rather as the insistence on working within a relatively neutral framework when it

comes to the metaphysics and psychology of particular areas of discourse, and for beginning to think more systematically about the literal meaning of declarative sentences of whatever category. There will surely remain interesting questions in semantic theory about how various syntactic, contextual, and logical factors contribute systematically to the propositions expressed by various sentences. There will still be questions about what semantic contribution is made by relevant sub-sentential components. There will also remain interesting questions in developmental linguistics/psychology about how beings like us acquire competence with the literal meanings of the relevant sentences and their components. Moreover, there may be further interesting questions in the metaphysics of semantics regarding what (if anything) a literal meaning *is* (abstract/concrete, structured/ unstructured, external/internal, etc.). Pursuing these further questions is perfectly consistent with the idea that declarative sentences, including those with ethical content, express propositions.[11]

Note that, on the conservative approach to theorizing about the meanings of ethical sentences that we are recommending, there is no expectation that it should be possible to provide a *meaning analysis* of ordinary declarative sentences, by paraphrasing them in other terms. A long history of failures in areas other than ethics to give paraphrases of sentences containing simple terms and analyses of atomic concepts should make us leery of any attempt to paraphrase sentences containing ethical terms in other terms (normative or not). Ethical sentences are not unique in this regard. According to the semantic conservatism we advocate, the unavailability of such paraphrastic analyses in the ethical case is not by itself indicative of some elusive 'is-ought' gap, but rather of the more general fact that meaningfulness does not require the availability of

[11] Price (1994, 2013: Ch. 1) defends a view he calls 'semantic minimalism' which we find congenial, but we don't think (as he does) that it supports global nonfactualism. The framework Price offers—in contrast to ethical neo-expressivism—is not designed to address, specifically, the motivational asymmetry between ethical and non-ethical discourse. We leave discussion of this and other differences for another occasion.

paraphrastic analysis. Nor do we need to suppose that the acquisition and competent use of such terms involves mastery of paraphrastic analyses. Sentences mean what they do in virtue of expressing propositions, and in many cases the propositions they express can be specified disquotationally, at least when the object language and the language in which the meaning is specified are the same. "John loves Mary" expresses the proposition that John loves Mary. And it is no different with ethical sentences: "Tormenting the cat is wrong" expresses the proposition that tormenting the cat is wrong. Thus the semantic conservatism we adopt allows us to be semantically neutral, in the sense of not being committed to the availability of a reductive meaning analysis for ethical terms and sentences. In general, we think that ordinary declarative sentences S can be said to express the proposition that-S, and this includes declarative sentences with ethical content.[12] Of course, a full compositional semantics must say much more about how the parts of ordinary declarative sentences interact to determine the proposition expressed, which will involve more than simple disquotation. But the view that declarative sentences express propositions that can be initially specified disquotationally comports with the most neutral framework for thinking about the compositional possibilities for ethical and nonethical terms and sentences.[13]

All in all, we think that there are good *linguistic* reasons not to make any special pleading for ethical sentences (or even normative sentences more generally). Such sentences do not constitute a distinct *semantic* category, i.e. one deserving a radically different treatment from other kinds of sentences, simply in virtue of containing ethical terms. (Specifically, there is

[12] Note that this is consistent with maintaining that meaning*fulness* requires much more than the mere possibility of disquotation or syntactic well-formedness.

[13] See Chrisman (2012) and Pettit (unpublished) for suggestions along these lines with respect to particular terms. The key point is to deny that semantic analysis must, e.g., result in an analytic paraphrase of some sort, involving lexical decomposition of the relevant terms, or spell out (nondisquotationally) necessary and sufficient conditions.

nothing in the ordinary use of these sentences that marks them as disguised non-declaratives, and thus as diverging in their linguistic behavior from declarative sentences that are plausibly taken to have propositional meaning. On the contrary.) Given the similarities in linguistic behavior between ethical and non-ethical terms, and given that ethical terms embed seamlessly in 'mixed' contexts, as well as figuring in various non-ethical contexts, we ought to prefer adopting the same basic semantic framework for all declarative sentences, ethical or not. And, for reasons articulated above, we think the propositionalist framework is a good place to start. (We note in passing that there's additional linguistic support for taking the link between descriptiveness and embedability – and the requisite notion of propositionality– to be relatively superficial. For, many sentences that are uncontroversially used (at least in part) to make genuinely descriptive claims – e.g. "What a loyal friend you are" – cannot be embedded in conditionals; "If what a loyal friend you are, then I can trust you" seems no less ill-formed than "If Loyal friend!, then I can trust you".)

So, what has motivated expressivists of many different stripes to depart from propositional semantics and undertake wholesale semantic revisionism? In the main, they have assumed this to be the only plausible way to accomplish two important things expressivism is intended to accomplish:

(a) Avoid commitment to 'spooky' irreducibly normative facts

(b) Capture the 'internal' connection between ethical claims and motivation

Re (a): Another traditional role assigned to propositions – one we haven't mentioned above – is *being bearers of truth-value*. Given that, and given an assumed connection between a sentence's being truth-evaluable and its serving to describe 'the way the world is', it might be thought that, if one concedes that an ethical sentence 'S' expresses the proposition that-S, then the only way

one could be an anti-realist would be to endorse some kind of unpalatable ethical error-theory or fictionalism.[14] Re (b): One of the key virtues of expressivism is supposed to be that it accommodates the motivational internalist's claim of a close connection between sincere ethical claims and motivation to act. But to do so, it is thought, one must build the relevant motivational attitudes into the *meanings* of ethical sentences.

However, we view the propositionalist framework we recommend as allowing one to be *neutral* on the metaphysical issue underlying (a) above—and neutral in a way that accommodates anti-realism. (See below.) Moreover, as regards (b) above, we think that whatever 'internal' connection there may be between ethical discourse and motivation, it should be captured not through the semantic content of ethical sentences—what *sentences* express, semantically speaking—but rather through what speakers (or thinkers) express in *acts* of ethical claim-making. In what follows, we take a step back and explain more carefully how we think of these two different expression relations and how they figure in defusing the motivations provided by (a) and (b) for expressivists to depart from standard propositionalist semantics. This puts us in a position to offer anyone moved by desiderata (a) and (b) new resources for meeting them without going in for anything like an expressivist semantics for ethical sentences.

3. A-Expression vs. S-Expression[15]

The class of behaviors ordinarily described as 'expressive' spans a wide range. At one end of the spectrum, we have so-called *natural expressions* – yelps, grimaces, and various gestures – where both the behavior and its connection to the expressed states are supposed to be

[14] See, e.g., Mackie 1977 and Kalderon 2005.

[15] Parts of the present section overlap, see Bar-On (forthcoming).

inculcated by nature. There are also mimicked or acquired facial expressions or gestures that become 'second nature', such as shrugging shoulders or tut-tutting. Then we have *conventional nonverbal expressions*, such as tipping one's hat or sticking out one's tongue. (The line here is not sharp; giving a hug, jumping for joy, stomping your feet, for example, all seem to fall somewhere in between, exhibiting both 'natural' and acquired elements.) Still in the conventional realm, we have *expressive verbal utterances* such as "Darn it!" or "Ouch!", "Sorry!", "This is great!" "I hate you!" and so on. We also find in the verbal domain utterances such as "There's a crow on the telephone pole," which (if sincere) expresses a speaker's present belief, or "Let it rain," which expresses the speaker's wish for rain. Finally, at the far end of the conventional side of the spectrum, we have *speech acts,* such as assertion, or promising, which are alleged to have the expression of certain mental states as part of their felicity conditions.[16]

Still, we can discern the following commonality among the expressive behaviors mentioned so far: they all express *states of minds*, as opposed to, say, propositions, concepts, or ideas. This is why we follow Sellars (1969) in conceiving of this kind of expression as a relation that holds between performers of acts and the mental states these acts directly express – "a-expression" – and contrast this with the relation that holds between meaningful strings (e.g., sentence-tokens) and their semantic contents. As we are thinking of it, a-expression is something a minded creature does, be it through bodily demeanor, facial expression or gesture, or else through speech, using a natural, conventional, and even idiosyncratic expressive vehicle.

Thus, for example, we say that the sentence "Snow is white" s-expresses the proposition that snow is white; and we say that the word "justice" s-expresses the abstract concept of justice. On the other hand, your dog, when he gets up and walks over to give you a lick, is a-expressing

[16] See Green (2007: 140-3).

his affectionate feeling. And when you give a friend a hug, or say "It's so great to see you," or alternatively "I'm so glad to see you," you a-express your joy at seeing your friend through an intentionally produced act (where some of the acts utilize sentences that in turn s-express propositions).

In general, when one a-expresses a state of mind using a sentence, the *sentence* uttered retains its linguistic meaning. "It's great to see you" and "I'm so glad to see you" each have their own meaning, in virtue of the linguistic rules governing the lexical items and their respective modes of composition. Each s-expresses a certain proposition. What proposition? Well, setting aside some nuances about the context-sensitivity of indexicals, it's most natural to say that the former sentence expresses the proposition that it's great to see the addressee, and the latter sentence expresses the proposition that the speaker is happy to see her addressee. It is because they express propositions that these sentences can partake in logical inferences and stand in systematic logico-semantic relations to other sentences (and, in particular, can be embedded in negation, conditionals, intensional contexts, etc.). For all that, we think that normal cases of producing unembedded tokens of these sentence types, in speech or in thought, are cases of one's *directly expressing one's joy*.

Intuitively, a-expression is more basic than s-expression. It's certainly more ubiquitous; nonhuman animals and prelinguistic children express states of mind through a variety of nonlinguistic means despite not having in their repertoire the ability to token sentences that s-express propositions. Now it may be that, ultimately (in 'the causal order of being'), s-expression has some (complicated, to be sure) relation to a-expression. Perhaps, for example, the conventional meanings of English sentences ultimately originate in Gricean speaker-meanings

that have 'fossilized'.[17] But once conventional meaning is in place, it is clear that we can separate what a given sentence s-expresses from what mental states speakers who use the sentence a-express, on a given occasion, or even regularly.[18]

Among linguistic creatures, a familiar acquisition process leads to the increasing use of more or less conventional means – both words and gestures – to give vent to present states of mind (making a disapproving face, thinking out loud, airing opinions, and so on). The distinction between expressive *acts*, on the one hand, and the expressive *vehicles* used in them, on the other, allows us to capture underlying action-theoretic similarities between expressing one's annoyance through a gesture, an inarticulate sound, or a full sentence, while still acknowledging significant differences.

Now, there's an intuitive contrast between acts of expressing states of mind and acts of merely *telling* about them. Anyone can say truly, and some can even tell reliably, *that* DB is feeling sad. But presumably only DB is in a position to *express* her sad feeling – for example, by letting tears roll down her cheeks, or saying "This is so sad". To use earlier terminology, we can say that, when you say "DB is feeling sad", you are employing a sentence that s-expresses the proposition that DB is feeling sad, and, if you are sincere, you are a-expressing your *belief* that DB is feeling sad, whereas DB's tears s-express nothing, though in letting them roll down she a-expresses the sad feeling itself; and her utterance "This is so sad" s-expresses the proposition that something sad is happening, and she can use it to a-express her sadness.

[17] As suggested in Blackburn (1984: Ch. 3). (See also Bar-On 1995.)

[18] An analogous point can be made regarding thought tokens. Such tokens have semantic contents. So they, too can be said to s-express propositions. (What the tokens s-express is something that an adequate (psycho)semantics will aim to specify.) But, when a given thought token is *produced*, we can separate what the individual producing the token a-expresses from what the token s-expresses.

What about DB's avowal: "I'm feeling so sad"? On the *neo-expressivist* view of avowals defended in Bar-On (2004), these are different from evidential reports on the presence and character of states of mind (whether others' or our own), in that they are acts in which we a-express the very state that is attributed to us by the proposition that is s-expressed by the sentences we use. One of the central points of defending this idea was that it can help explain the *epistemic asymmetry* between avowals and third-person reports of the same states without compromising the *semantic continuity* between avowals and other statements. The asymmetry is not between *types of sentences with certain semantic contents.* Rather, the contrast in play here is between *acts that directly express* one's mental state and reports of that state (whoever produces them, and however reliably). (For, notably, the sentence "I am feeling sad" can be used by me as a mere evidential report of DB's state, say at the conclusion of a therapy session.)

For present purposes, what matters is that the distinction between s-expressing and a-expressing be recognized as a distinction that applies across all areas of discourse, and regardless of what semantic, epistemological, or metaphysical analysis we adopt for the relevant domain. On the view we advocate, mental states are indeed the relata of *an* expression relation. And it's a relation that may well be relevant for understanding of various linguistic acts we perform. But the relation in question is what we're calling *a*-expression. It is *not* the expression relation that holds between sentences and propositions, or words and concepts (i.e., s-expression).

The relevance of this here is that we think that the expressive character of ethical claims and their apparently tight connection to motivation can be explained in partial analogy to the neo-expressivist treatment of avowals mentioned earlier.[19] According to avowal neo-expressivism, the epistemic asymmetry between avowals and (even self-)reports is due to the fact

[19] Following Bar-On and Chrisman (2009).

that, in acts of avowing, one gives direct vent to the very state that the sentence produced self-ascribes. Similarly, ethical neo-expressivism maintains that ethical sentences s-express straightforward propositions (which can be specified, at least initially, disquotationally); and this allows us to preserve their semantic continuity with other sentences. However, what we might describe as the *motivational asymmetry* between ethical discourse and ordinary descriptive discourse can—on analogy to the epistemic asymmetry between avowals and reports—be captured (according to ethical neo-expressivism) by appeal to the expressive character of *acts* of making ethical claims.

So, on our view, expressivists have been correct to identify the expressive function of ethical discourse and reflection; if internalism is correct, then it is in fact part of the 'job' of ethical claims (whether made in speech or in thought) to express certain of our motivational attitudes.[20] Armed with the distinction between s-expression and a-expression, we think we can retain this key idea while avoiding a host of difficulties faced by traditional expressivism and later developments of it. As in the case of avowals, we can distinguish between the act of making an ethical claim and the vehicle used in making it. The vehicle used in claiming "Tormenting cats is morally wrong" is a sentence- (or thought-) token that employs an ethical term (or concept); the token can be said to s-express a true or false proposition. However, as with avowals, what is s-expressed by a claim does not settle even what mental state is characteristically a-expressed by acts of making the claim, let alone what mental state is a-expressed on a given occasion of producing the token. With Humeans, one may think that purely cognitive states (such as beliefs) cannot by themselves motivate or explain action, and furthermore, one may think there is an intimate connection between sincerely making an ethical

[20] Below, we try to spell out how the neo-expressivist can respect the internalist take on this claim.

claim and being motivated to act (or refrain from acting) in accordance with it. The neo-expressivist maintains that in making ethical claims, we a-express the very same states whose presence is required for understanding the perceived motivational force of such claims.[21] As Ayer already saw,[22] this expressivist insight is best captured *without* supposing that the vehicles used in making ethical claims s-express propositions that self-ascribe those states. (So it's important to note that we are *not* claiming that ethical claims are themselves avowals!)

Thus, as long as we're talking about expressive acts, we can agree with Ayer that ethical claims betray or show motivational attitudes, *not* because they report them,[23] but because those who make them (a-)express them directly. What the *neo*-expressivist goes on to add is that we need not endorse Ayer's view that ethical claims are not truth-evaluable because they do not (s-) express *any* true or false propositions. According to ethical neo-expressivism, then

(i) As a species of evaluative claims, ethical claims understood as *acts* are different from ordinary descriptive claims in that agents making them (in speech or in thought) *a*-express motivational attitudes.

(ii) The *vehicles* used in making ethical claims – typically, ethical sentences – are semantically continuous with ordinary descriptive sentences in being truth-evaluable, embeddable in truth-functional as well as intensional contexts, logical inferences, etc. This is because they *s*-express true or false propositions.

[21] They do so, *whether or not* they *also* express a belief whose content is given by the s-expressed proposition. For some discussion, see Bar-On and Chrisman (2009).

[22] Ayer (1936/1946: 104-8).

[23] And, we might add, not even because they *imply* that one has them. (This, in contrast with implicature-style hybrid views of Copp 2001, 2009 and Finlay 2005. For discussion of the contrast, see Bar-On and Chrisman 2009.)

4. Ethical Neo-Expressivism: Anti-Realism, Internalism

In the previous section, we argued that the idea that ethical claims are distinctive in that they express motivational attitudes can be captured without building attitude expression into the *semantic content* of ethical sentences. But doing so was, of course, supposed to allow expressivists to accommodate ethical anti-realism and to capture motivational internalism. These were desiderata (a) and (b) from the end of section 2. How does a neo-expressivist view provide resources for satisfying them?

Re (a): Recall that above we distinguished propositions from truth-conditions, but said that propositions may *determine* the truth-conditions of the sentences that express them. There is a pervasive tendency in metaethical debate to combine this idea with a metaphysically inflated conception of truth and to interpret anyone who thinks that a sentence expresses a proposition as committed to there being a way the world might be that would make this proposition true. While we recognize that many metaphysically inclined philosophers are attracted to this picture, we deny that it is essential to the propositionalist framework for philosophical semantics. Notice that none of the reasons we gave earlier for assigning propositions to sentences as their semantic contents (translation, embedding, common content of diverse attitudes) trades on any particular conception of truth as correspondence with the world, or of truth-conditions as ways the world might be. In our view, semantics is not the place to explain how reality is or might be, it's the place to systematize our understanding of things like sameness of meaning across different languages, commonalities of content under embeddings in force-stripping contexts, and the way sentences can articulate contents towards which diverse attitudes may be taken.

Because of this, 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 assign, e.g., "Tormenting the cat is wrong" the meaning *that* tormenting the cat is wrong, for logico-semantic purposes, independently of determining what, if anything, constitutes the nature or essence of *moral wrongness* (or *torment*, or *cats*, for that matter), or whether moral wrongness constitutes a metaphysically genuine property (natural or otherwise). The important point is just this: an expressivist *can* accept the sort of semantic conservatism we're here encouraging and still avoid any commitment to 'spooky' irreducibly normative properties. Semantic conservatism is consistent with ethical anti-realism, since it implies nothing about the *existence* or *nature* of moral reality. Given this neutrality, it's left for further metaphysical investigation to determine whether there are ethical properties or states of affairs, and if there are, what their nature is and to what extent they are mind- (or judgment-) dependent.[24]

Some, however, will worry that the sort of semantic conservatism that we here encourage offers scant resources for explaining what sentences in a given area of discourse are *about*. To say that ethical sentences s-express 'innocent' disquoted propositions is to offer a very thin notion of the sentences' meanings. If one wanted to understand what such sentences 'really say', or what they're 'really about', then one would apparently have to beef up the account of the sentence's s-expressed content. But then it's no longer possible to maintain the sort of semantic conservatism that we prefer—either one goes realist, or one goes traditional expressivist, or subjectivist, or whatever. Or so the objection might go.[25]

[24] We are claiming that, in the ethical case (unlike in the avowal case; see Bar-On 2102), the possibility that there are no 'truth-makers' for the relevant propositions.should be left open, to allow for a meaningful dispute between ethical realists and anti-realists, who can agree on the availability of a disquotationally specified proposition as a semantic starting point. (Note that to say this is not to commit to a disquotational theory of the *truth* of ethical—or other—claims.)

[25] Richard (forthcoming) provides reasons for expecting semantics to yield assignments of meaning that go beyond 'innocent' propositions.

In response, we'd begin by pointing out that, as far as the semantics of ethical sentences is concerned, it is not as if *all we can say* is that they s-express 'innocent' disqouted propositions. As explained earlier, there is still plenty of interesting work left for the semanticist to do: arguably, she ought to explain how the proposition expressed by an ethical sentence is a systematic function of the semantic values of the parts of the sentence, how these are composed in its logical form, and so forth. She can seek systematic integration of that piece of her semantics with other pieces, such as her view about embedding in indicative conditionals, under epistemic modals, or in propositional attitude reports. She can posit covert context sensitivity or deny it. She can debate the logical form of the sentence and ones syntactically like it. What she cannot do, however—without abandoning the sort of ontological innocence that we'd like to maintain—is to say what *in reality* the sentence is *about*. Our view, however, is that the thought that a propositionalist semantics carries specific ontological commitments is a significant step beyond the truth-conditionalist framework that has proven so fruitful in compositional semantics.

We join Blackburn in thinking that "the ethical proposition is what it is and not another thing."[26] If we understand him correctly, his point is just that, when asked what a sentence like "Tormenting the cat is wrong" is *really about*, one should simply say that it's about tormenting the cat's being wrong. To say anything further would be to risk running afoul of the Open Question Argument and/or queerness worries. Taking such a conservative approach to the semantics of ethical sentences allows us to capture all of the virtues of traditional propositional-

[26] Blackburn (1998: 79). Blackburn here adapts a famous line from Butler via Moore: "Everything is what it is and not another thing".

compositional semantics while remaining totally silent on the appropriate metaphysics for ethical claims.[27] And it is this silence that allows us to accommodate ethical anti-realism.

Perhaps a further thing to say in response to this objection would be to question how coherent it ever was to ask what sentences in a given area of discourse are about. In particular, we wonder what sense can be made of the notion of "sentences in *a given area of discourse*," especially when it is noted that such sentences can typically be embedded rather easily into force-stripping contexts, some of which will be logically complex sentences with "mixed" parts—i.e., parts of which belong in the supposed domain, and parts that do not. For instance, "If he φ-ed, then he ought to be punished" seems like a sentence belonging to the domain of ethics, but what about "If he ought to be punished, then he φ-ed"? And what about sentences like "It's either morally wrong to φ or it isn't"? Such possibilities, and the ease with which they are entertained, may tell against the very notion of a *given domain*. And if sufficient sense cannot be made of this notion, then it would appear as if all that remains are *particular* meanings of *particular* sentences, to be determined lexically and compositionally—an outcome with which we are perfectly comfortable.

Furthermore, our semantic conservatism allows us to remain neutral on the nature of the *attitude expressed* when making ethical claims. Any number of attitudes can do, as long as it can be argued that the relevant attitude is suitably linked to motivation (a matter left to moral psychology). Again, we see it as an advantage of our view that it isn't forced to settle in the

[27] This is an instance of the sort of problem Price (2004) refers to as 'placement problems'. See also Price (2013: ch. 1 §4) for a helpful discussion of the origin of these problems and how to diffuse them.

semantics of ethical discourse *what* attitudes agents express when employing ethical vocabulary.[28]

Some may wonder, however: even if ethical neo-expressivism does not owe us an account of the states expressed when making ethical claims, does it not owe us an account of the *connection* between these states and ethical claims understood as *products*? After all, in the case of avowals, the connection seems pretty straightforward: the sentence typically explicitly mentions the very state that is a-expressed. So what accounts for the connection between a sentence like 'Tormenting the cat is wrong', and, say, disapproval, if the latter does figure into the semantic content of the former?

Here again, there are two things to say. First, it is a mistake to think that the connection between mental states and avowals is simply a matter of the sentences' typically explicitly *mentioning* the relevant states. Consider two sentences from our earlier discussion of a-expression:

(1) "It's so great to see you!"

(2) "I'm so glad to see you!"

As we explained before, these sentences can be used as vehicles for the a-expression of one and the same mental state—being glad, or feeling joy. However, if we were to explain the connection between the sentence and the mental state in (2) in terms of the sentence's explicitly mentioning the relevant state, we would apparently have to tell some *other* kind of story about how the same connection is forged in (1). Instead, we think this connection is forged in the same way for both (1) and (2), as well as for other vehicles for the a-expression of being glad, such as a hug. We

[28] This is the problem that has been called 'the Specification Problem' for expressivism. See Köhler (2013) and Björnsson and McPherson (forthcoming).

prefer to account for this connection by making reference to the conditions underlying *competence* with the use of expressive vehicles in performing relevant types of expressive acts. In the case of *both* (1) and (2), for instance, competence with the sentence requires a speaker to know that they are fit vehicles with which to a-express being glad, feeling joy, and the like. The same goes for acts such as hugging. (With a little imagination, the reader could envision various cases of misplaced, poorly timed, or otherwise inappropriate hugs, and ask what sorts of conclusions we'd draw about the hugger's competence.[29]) Explicitly mentioning the relevant state is neither necessary nor sufficient for establishing, mastering, or preserving the expressive link to it in an expressive vehicle.[30]

Secondly, going back to the question of how ethical claims in particular come to be linked to certain motivational states, we think it instructive to consider the case of *slurs* (and other 'linguistic expressives', since these, too, are terms (or phrases) with a clear connection to certain mental states (though, unlike avowals, the connection is not explicit). One can a-express anger toward John by saying, "I'm angry with John!" or by saying, "John is such an a-hole!" Here again, one of the sentences 's'-names the relevant attitude, but both are obviously effective (and common) as expressive vehicles for one's anger. And this is because competence with the term 'a-hole' requires implicit knowledge that it is a fit vehicle with which to a-express anger (or related sentiment). To this extent, we agree with hybrid expressivists who think that, e.g., pejoratives and slurs nicely illustrate how terms and sentences can be bound up with certain noncognitive attitudes, despite bearing no explicit connection to the relevant states.[31] Where we

[29] See Bar-On (2004: 320ff., 419ff.) for relevant discussion of what she calls 'expressive failures', which can arise even in the case of inadvertent expressive behaviors, such as yelps or grimaces.

[30] That it's not sufficient can be clearly seen from the fact that mental states are regularly named in sentences that are not used in acts of (a-) expressing them.

[31] One finds references to slurs throughout the literature on hybrid expressivism, as many think they nicely embody the basic point of hybrid theories. Copp, for instance, writes that ethical terms "are similar to pejorative

part ways, however, is in our denial that the noncognitive attitudes somehow figure into the *semantic contents* of expressive terms, or of the sentences containing them. Though there may be much to be said about the state of mind characteristically (a-) expressed by (proper) uses of "a-hole" (and though there may be interesting things to say about what *makes it the case* that someone is an a-hole), "John is an a-hole!" can still be said to s-expresses the proposition that, well, John is an a-hole. (And it is in virtue of this that the sentence can be embedded in a variety of contexts, partake in logical inferences, etc.[32])

Re (b): According to internalists, one who makes a claim such as "Charity is good" *competently* can't be completely *indifferent* to charity; whereas one who makes the grammatically similar claim "Charity is common" can. There is little consensus, even among self-avowed internalists, about the character of the 'can't be' here, but we'd still like to outline briefly how our view can accommodate an internal connection. No doubt there will be internalists who deem the connection insufficiently tight. But it's far from clear that any tighter connection is defensible.[33]

Earlier, we mentioned the neo-expressivist explanation of the epistemic asymmetry between avowals and evidential reports: someone who is *avowing* feeling annoyed—as opposed to reporting annoyance in consequence of, say, therapy—is taken to engage in an act of direct expression of annoyance (and *not* simply expressing the *belief* that she is annoyed). Though it's not *conceptually impossible* for someone to make a mental self-ascription and not be in the self-

terms in that their use can both ascribe a property and express a relevant conative attitude" (2009: 169-170). See also Boisvert (2008), Hay (forthcoming), and Schroeder (forthcoming).

[32] Note that, in the case of linguistic expressives, it may be plausible to suggest that, e.g. "John is such an a-hole" is more or less synonymous with "What an a-hole !" (applied to John). Yet the former can—but the latter *cannot*—be embedded in a conditional, or serve as a premise in an instance of modus ponens, and so on.

[33] See Björnsson, G. & Francén Olinder, R. (2013), Fletcher (unpublished) for critiques of stronger forms of internalism. See Björklund, et al. (2012) and Fletcher (unpublished) for helpful discussion of the many different forms and strengths of internalism and its dialectical role in metaethical debate.

ascribed state, it is a *propriety condition* on avowing that one is in the self-ascribed mental state. The neo-expressivist explanation of the motivational asymmetry between ethical claims and descriptive claims can be seen as analogous. Someone who is making an *ethical* claim (as opposed to producing a descriptive report of some state of affairs) is a-expressing the relevant motivational attitude – the very attitude whose presence would explain why she is suitably motivated. Even if it's not conceptually impossible for someone to make an ethical claim without having the relevant attitude, having the attitude can still be a *propriety condition* on making (genuinely) ethical claims. So someone who makes the claim while lacking the attitude violates a propriety condition on acts of ethical claim making. And adequate mastery of ethical discourse and ethical concepts requires grasping this propriety condition.[34] This arguably provides resources for capturing a fairly strong internal connection between ethical claims and action, as well as providing a more nuanced array of diagnoses of different ways the connection between making an (apparently) ethical claim and motivation can be broken.[35]

Propriety conditions lay down norms for what counts as doing things *properly*. For example, a propriety condition on making promises may be that one has an intention to do as the promise says. This means that, *if* one *has* made a promise without having the relevant intention, one has failed to *properly* issue a promise. Note that improperly making a promise is not the same as not making a promise at all. An actor on stage who says "I promise" is only making a pretend-promise; a parrot, or a not-yet-competent speaker who utters the words, is making no promise at all, rather than pretending to. Does this mean that it's *conceptually impossible* for one to say "I promise" meaningfully without having the relevant intention? Of course not, since it's

[34] Again, thinking of slurs and other linguistic expressives is instructive here.

[35] For further discussion, see Bar-On and Chrisman (2009: 143-150).

possible to say "I promise" *insincerely*. But notice that we *are* willing to credit the insincere speaker with understanding of and semantic competence with the relevant term, despite her failing to have the requisite intention. What licenses this? It is the fact that we have some independent grounds for crediting her with this receptive and productive competence. (Indeed, such competence must feature in the explanation of her insincere performance.) But now suppose we have a speaker who is *not* insincere; she says "I promise", but fails to have the relevant intention—perhaps due to some temporary or chronic psychological feebleness. Still, we may have reason to credit her with competence. For example, she may exhibit full awareness of the propriety conditions on the practice of promising; she just fails to meet them. Of such a speaker, we would want to say that she *has* made a promise, albeit an improper one.

In the ethical case, we are similarly inclined to say that it ought to be conceptually possible to have a speaker (or thinker) who has achieved competence with ethical vocabulary, but who makes an ethical claim while still failing to have the relevant attitude. This, however, doesn't commit us to saying that one can *achieve* competence with the ethical vocabulary without grasping the propriety condition mentioned earlier. Moreover, suppose we come across a whole community whose speakers appear to use ethical vocabulary, but who never have the appropriate attitudes. Given semantic conservatism, we should allow that the sentences used by members of this community may have the same propositional contents as our ethical sentences. But this doesn't settle the question whether they are actually making ethical claims (ever). If their discourse is governed by entirely different propriety conditions, then it's not ethical discourse. Compare: If in some twin community there's no connection whatsoever between the use of a type of sentence and trying to amuse/being amused, that tells against interpreting the use of that type of sentence as 'telling a *joke*'. Still, this wouldn't be reflected in the translation of

individual jokes. What *makes* an utterance an act of making a joke (involvement in a practice governed by certain norms) needn't be read off the meaning of individual sentences used as jokes. Similarly, we want to say, what makes a claim an ethical claim needn't be read off the meaning of individual sentences used as ethical claims.

An analogy may be helpful here. Suppose I'm playing a game of chess, and I move the rook diagonally. We might say either that I've made an *improper* chess move, or that I haven't made a *chess* move at all. There's good reason to say the former. To "make a chess move," is just to make a move while engaged in a certain sort of practice, which is itself governed by rules. This is the difference between me moving the rook diagonally while playing a game of chess, and say, 1-year-old child grabbing a rook and moving it elsewhere on the board (even if perfectly in line with the rules). The child is not *playing chess*; I am. Similarly, the psychopath, when he confesses that he behaved wrongly – despite not feeling guilty, regretting his actions, etc. – is still making an ethical claim, though he makes his claim improperly. Genuinely making ethical claims is to be contrasted with things like: using moral/evaluative language while acting. Regardless of whether the actor on stage has the relevant attitudes, she is not making a genuinely ethical claim—she is just *pretending* to engage in ethical discourse. Similarly for using moral/evaluative language sarcastically, playfully, etc. E.g., saying sarcastically, "Good job, Ryan!" Or saying playfully, "You're so bad." These are *not* cases of people using the language improperly, but rather cases of not making any genuine ethical claims at all. By contrast, it doesn't seem plausible to think that Ted Bundy was *pretending* to make ethical claims, nor was he using the language playfully or sarcastically. As far as he was concerned, he was just as engaged in the practice of moral evaluation as *we* are when we morally judge his behavior—and

so, we would perhaps affirm that he is *making an ethical claim*. But on the assumption that Bundy felt no remorse (etc.) at all, he was nonetheless doing so *improperly*. [36]

5. Conclusion: Neo-Expressivism and Hybrid Expressivism?

Proto-expressivism flounders on the challenge of developing a systematic account of the meaning of arbitrary sentences out of the essentially bifurcated thought that ethical sentences express noncognitive attitudes while non-ethical sentences express propositions. The response to this challenge from mainline expressivists of the 1960s-1990s was to suggest that the meaning of *all* sentences is to be given in terms of the attitudes they express; it's just that descriptive and ethical sentences express importantly different kinds of attitudes: beliefs vs. desires (or other motivational states). As we mentioned at the outset, more recent "hybrid theorists" have suggested that this approach still faces a number of difficulties stemming from the semantic continuities between descriptive and ethical sentences, and these might be overcome by suggesting that although the meaning of *all* sentences is to be given in terms of the attitudes they express, ethical sentences express a distinctive hybrid attitude, constituted by both belief-like and desire-like elements. Roughly speaking, the belief-like attitude is supposed to capture the semantic continuities between ethical sentences and descriptive sentences, while the desire-like attitude is supposed to retain the original benefits of mainline expressivism.

In light of this, it's worth asking, is the neo-expressivist view we've been discussing here a 'hybrid' view? It might seem that way, since, on the face of it, we maintain that when someone makes an ethical claim (in speech or in thought) two different things get expressed: a proposition

[36] Note that Bundy's case *could* be assimilated to the 'deviant community' case, if there were reasons to think that he simply hasn't cottoned on to the propriety conditions of ethical discourse. It depends on the case!

and the speaker's or thinker's motivational attitude. But, as should by now be clear, this is misleading, since on our view, the two things are expressed in two different senses—the proposition is s-expressed by the relevant sentence- (or thought-) token, whereas the motivational attitude is a-expressed by the individual making the claim. Importantly, in sharp distinction from other hybrid views, we do not maintain that the 'two things' constitute two parts of the meaning of ethical sentences. Nor do we locate the attitude expression in a conventionally implicated or conversationally inferable proposition that self-ascribes the attitude, as do implicature views[37]. Thus, rather than pursuing a non-conservative semantics for ethical sentences and then hybridizing the mental state said to be semantically expressed by ethical sentences, ethical neo-expressivism proposes that we use the more basic and less controversial distinction between what sentences express (i.e. propositions) and what people express (i.e. mental states) to capture the internalist intuitions, but in a way that allows us to remain neutral on the metaphysical issue of realism as well as on the moral psychological issue of cognitivism.

The notion of a-expression seems more general and more basic than any of the following: (i) semantically expressing a proposition, (ii) semantically expressing a mental state, (iii) stating, implicating or otherwise conveying the proposition that one is in a particular mental state. This is why we think it's the appropriate place to look to explain the supposedly distinctive way in which ethical claims express attitudes. Moreover, we take the notion of s-expression to establish a framework for thinking about semantic content within the constraints of compositionality, as well as semantic competence of speakers. However, it's a framework that we take to be attractively neutral on metaphysical and moral-psychological issues, which is why we think it's preferable for anyone seeking to accommodate ethical anti-realism and to capture motivational

[37] See Copp (2001, 2009) and Finlay (2004, 2005). Bar-On and Chrisman 2009: 150-158 compares and contrasts ethical neo-expressivism with these implicature views.

internalism. This is why neo-expressivism offers a nice way to retain one of the core insights of traditional expressivism (that acts of making an ethical claim 'have the job' of expressing attitudes distinctively connected to motivation) while avoiding one of the main pitfalls of many later versions of expressivism (resulting from abandoning propositionalist explanations of the semantic content of ethical sentences). We close by responding to three potential worries.

The first worry is that neo-expressivism doesn't meet desiderata (a) and (b) from p. XX above (avoiding spooky facts and accommodating the internal connection to motivation) in the semantics of ethical sentences, and it doesn't meet them in the pragmatics of ethical discourse, so it is unclear what part of the view of ethical language *does* allow us to meet them. Our answer to this is that these desiderata should be met in the theory of ethical claim making, including the propriety conditions on such acts. We are suggesting that there is a particular kind of act (both in public speech and in private thought) whose propriety is linked to having certain motivational dispositions, and this link is what explains whatever distinctive internal connection there is between ethical claims and motivation.

The second worry can be presented as a critical question from hybrid-expressivists like Ridge who think our view is not so different from theirs: "Even if you don't want to endorse it as your 'semantics' of ethical sentences, aren't you also committed to hybridizing not only the expression relation but also the mental state expressed (you'll say 'a-expressed') by ethical claims? In order to capture their *psychological* continuity with other claims, you'll have to say that ethical claims a-express *beliefs*, but in order to capture the difference in their motivational potentials, you'll have to say that they also a-express *motivational attitudes*." Our response is that ethical neo-expressivism itself is neutral on the psychology of motivation and on the extent to which ethical thoughts are psychologically continuous with other kinds of thoughts. It's open

to us to say ethical claims a-express beliefs with ethical content, which as a matter of their psychological nature have a distinctive motivational capacity compared to other sorts of belief; but we can also say that, in light of the Humean distinction between cognitive and conative attitudes in the psychology of motivation, acts of making ethical claims a-express both. Although ethical neo-expressivism was originally framed against the background of this Humean distinction, which is a key tenet of many forms of expressivism, the view is also compatible with anti-Humean accounts of motivation.

A final worry concerns ethical claims made in *thought*, rather than in speech. Many find mysterious the idea that inner thoughts are subject to propriety conditions about what mental state one should be in – for how would those propriety conditions get enforced, where would they come from, and what evidence could there be of their existence? In response, we'd stress that our view is that it's part of what it is to make an ethical claim, *whether in speech or in thought*, that one has done so inappropriately if one isn't in the relevant motivational state. This "comes from" the point and function of a practice of ethical claim making. Although the rules of that practice may be most easily *enforced* when it is public, we see no reason that these rules couldn't be internalized and be thought to govern 'silent' acts of ethical claim making. We take this to be a mundane phenomenon illustrated by other cases of expressing in thought, such as making promises to oneself (or to God), avowing, or using various expressives. Thus, all in all, we think that understanding the sense in which ethical claims express motivational states along our proposed neo-expressivist lines should allow expressivists to meet some of the main desiderata that originally motivated expressivists while avoiding the main difficulties besetting traditional expressivism and its heirs, and while maintaining neutrality on key debatable issues.

References

Ayer, A. (1936/1946). *Language, Truth, and Logic*. 2nd ed. London: V. Gollancz Ltd.

Bar-On, D. (1993). "Indeterminacy of Translation: Theory and Practice" *Philosophy and Phenomenological Research* 53: 781-810.

———. (1995). " 'Meaning' Reconstructed: Grice and the Naturalizing of Semantics," *Pacific Philosophical Quarterly* 76: 83-116.

———. (2004). *Speaking My Mind: Expression and Self-Knowledge*. Oxford: Clarendon Press.

———. (2012). "Expression, Truth, and Reality: Some Variations on Themes from Wright," in Coliva, A. (ed.). *Mind, Meaning, and Knowledge: Themes from the Philosophy of Crispin Wright*. Oxford: Oxford University Press.

_____. (forthcoming). "Expression: Acts, Products, and Meaning". In *Minimalism, Pragmatism, Expressivism: Essays on Language and Metaphysics*, S. Gross, M. Williams, and N. Tebben (eds.), Oxford University Press.

Bar-On, D., and Chrisman, M. (2009). "Ethical Neo-Expressivism," in Shafer-Landau, R. (ed.). *Oxford Studies in Metaethics, Vol. 4*. Oxford: Oxford University Press.

Björklund, F., Björnsson, G., Eriksson, J., Francén Olinder, R. & Strandberg, C. (2012) "Recent Work on Motivational Internalism", *Analysis,* 72**,** 124-137.

Björnsson, G. & Francén Olinder, R. (2013) "Internalists Beware -- We Might all be Amoralists!", *Australasian Journal of Philosophy,* 91: 1-14.

Björnsson, G., and McPherson, T. (forthcoming). "Moral Attitudes for Expressivists: Solving the Specification Problem," *Mind*.

Blackburn, S. (1984). *Spreading the Word: Groundings in the Philosophy of Language*. Oxford: Oxford University Press.

———. (1998). *Ruling Passions: A Theory of Practical Reasoning*. Oxford: Clarendon Press.

Boisvert, D. (2008). "Expressive-Assertivism," *Pacific Philosophical Quarterly* 89: 169-203.

Carnap, R. (1935). *Philosophy and Logical Syntax*. London: Kegan Paul, Trench, Trubner & Co. Ltd.

Chrisman, M. (forthcoming). "Attitudinal Expressivism and Logical Pragmatism in Metaethics". In *Pragmatism, Law, and Language*, G. Hubbs and D. Lind (eds.) New York: Routledge.

———. (2012). “On the Meaning of ‘Ought’”. In Shafer-Landau, R. (ed.), *Oxford Studies in Metaethics, vol. 7*, Oxford: Oxford University Press.

Copp, D. (2001). “Realist-Expressivism: A Neglected Option for Moral Realism,” *Social Philosophy and Policy* 18: 1-43.

———. (2009). “Realist-Expressivism and Conventional Implicature,” in Shafer-Landau, R. (ed.). *Oxford Studies in Metaethics, Vol. 4*. Oxford: Oxford University Press.

Davis, W. (2003). *Meaning, Expression, and Thought.* Cambridge, Mass: Cambridge University Press.

Finlay, S. (2005). “Value and Implicature,” *Philosophers’ Imprint* 5: 1-20.

Fletcher, G. (unpublished) “Motivational Judgment Internalism Explained (Away)”.

Geach, P. T. (1965). “Assertion,” *The Philosophical Review* 74: 449-465.

Gibbard, A. (2003). *Thinking How to Live*. Cambridge, Mass.: Harvard University Press.

Green, M. (2007). *Self-Expression*. Oxford: Oxford University Press.

Hare, R. M. (1952). *The Language of Morals*. Oxford: Oxford University Press.

———. (1963). *Freedom and Reason*. Oxford: Oxford University Press.

Hay, R. (forthcoming). “Hybrid Expressivism and the Analogy between Pejoratives and Moral Language,” *European Journal of Philosophy*.

Jackson, F. (2001). “What is Expressivism?” *Philosophical Books* 42: 10-17.

Kalderon, M. (2005). *Moral Fictionalism*. Oxford: Clarendon Press.

Köhler, S. (2013). “Do Expressivists Have an Attitude Problem?” *Ethics* 123(3): 479-507.

Mackie, J. L. (1977). *Ethics: Inventing Right and Wrong*. London: Penguin.

Moore, G. E. (1903). *Principia Ethica*. Cambridge: Cambridge University Press.

Pettit, D. (unpublished). “The Semantics of ‘Good’.”

Price, H. (1994). “Semantic minimalism and the Frege point”. In Tsohatzidis, S.L.(ed.), *Foundations of Speech Act Theory: Philosophical and Linguistic Perspectives*. Routledge, pp. 132-55.

Price, H. (2004). Naturalism without representationalism. In David Macarthur and Mario de Caro (eds), *Naturalism in Question*. Cambridge: Harvard University Press, pp. 71-88.

Price, H. (2013). *Expressivism, Pragmatism, and Representationalism*, Cambridge University Press.

Richard, M. (forthcoming). "What Would an Expressivist Semantics Be?" in *Minimalism, Pragmatism, Expressivism: Essays on Language and Metaphysics*, S. Gross, M. Williams, and N. Tebben (eds.), Oxford University Press.

Ridge, M. (2006). "Ecumenical Expressivism: Finessing Frege," *Ethics* 116: 302-336.

———. (2007). "Ecumenical Expressivism: The Best of Both Worlds?" in Shafer-Landau, R. (ed.). *Oxford Studies in Metaethics, Vol. 2*. Oxford: Oxford University Press.

———. (forthcoming). *Impassioned Belief*. Oxford: Oxford University Press.

Schroeder, M. (2008). *Being For: Evaluating the Semantic Program of Expressivism*. Oxford: Oxford University Press.

———. (forthcoming). "Tempered Expressivism," in Shafer-Landau, R. (ed.). *Oxford Studies in Metaethics, Vol. 8*. Oxford: Oxford University Press.

Sellars, W. (1969). "Language as Thought and as Communication," *Philosophy and Phenomenological Research* 29: 506-527.

Soames, S. (2010). *What is Meaning?* Princeton: Princeton University Press.

Soames, S. (2010) *What Is Meaning?* (Princeton University Press. Princeton, NJ.)

Stevenson, C. L. (1937). "The Emotive Meaning of Ethical Terms," *Mind* 46: 14-31.

———. (1944). *Ethics and Language*. New Haven: Yale University Press.