

# Unified (Enough) Metasemantics for Expressivists

TEEMU TOPPINEN & VILMA VENESMAA

We discuss a challenge for expressivism in metaethics. According to expressivism, the meaning of normative sentences is explained by their playing a *practical role*, or by facts about what desire-like, or action- or attitude-guiding states of mind, normative sentences express. We first explain how expressivism can be understood as a view about the *metasemantics* of normative language (section 1). The challenge, which we may call *the problem of diverse uses* (Väyrynen 2022), is based on the simple observation that while terms such as “good” or “ought” plausibly have a unified meaning across a wide variety of different uses, not all uses of sentences that contain these terms seem to play a suitably practical role. How, then, can the expressivist explain the meaning of such sentences by appealing to the idea that they play a practical role (section 2)? We suggest that expressivists can deal with this challenge. Our response is based on two ideas. First, understanding expressivism as a view in metasemantics rather than in semantics creates space for the possibility that both the practical and the descriptive uses of normative terms might carry the same meaning. This requires adopting a metasemantics that has some complexity, which leads to what we may call *the problem of disunified metasemantics* (Wodak 2017, Väyrynen 2022). However, we argue that this problem may nevertheless be dealt with, given that the extra complexity is required in order to capture the relevant phenomena (section 3). Second, in order to avoid a remaining challenge that we may call *the problem of unexplained metasemantic coincidence* (Wodak 2017), the expressivist account should take a certain kind of form. We suggest that a

view called *relational expressivism* holds promise with regard to the prospect of giving a unified enough metasemantics for normative language that doesn't rely on unexplained coincidences (section 4). Finally, we briefly conclude (section 5).

## 1. Expressivism and the practical role of normative language

By the term "normative language" we may single out, roughly, those chunks of language that centrally deploy terms such as "good," "ought," and "reason," or whatever it is that these terms translate to in languages other than English. Examples of normative sentences would then be sentences such as "Knowledge is good in itself," "We ought to ban Nazi symbols," or "There's some reason to eat cars."

Normative terms, or terms such as "good," "ought," and "reason," play a *practical role* in our thought and talk. A term plays a practical role, we might say, when its use normally expresses the speaker's practical attitudes of some appropriate kind. But the meaning of "expresses," and consequently, the idea of a practical role, can be understood in very different ways. One possible view would be that normative language expresses practical attitudes as a *broadly semantic matter*, or in virtue of suitable linguistic conventions. According to an alternative view, uses of normative language express such attitudes as a matter of *pragmatics*.

As an example of a view of the latter sort, it might be suggested that normative claims carry a *generalized conversational implicature* to the effect that the speaker has certain attitudes. When some claim conveys, via a generalized conversational implicature, that the speaker has a certain attitude, the suggestion that the speaker has this attitude is not a part of what is said, and can be "canceled," but can nevertheless be assumed to be true in the absence of a special context. An example might be provided by the sentence "They drank some of the tea," which, in the absence of special circumstances, implicates, but doesn't say, that the person picked out by "they" did not drink *all* of the tea. (As noted, what is thereby pragmatically conveyed by the use of the sentence is, in this sort of case, cancellable in that it would not be linguistically inappropriate to say, "They drank some of the tea—indeed,

they drank it to the last drop.” On appeals to generalized conversational implicatures in relevant contexts, see Strandberg 2011, Fletcher 2014.)

Expressivism, by contrast, is an example of a view that construes the practicality of normative language as a broadly semantic matter. According to expressivism, the meaning of normative sentences is explained by their playing a practical role, or by facts about what desire-like or action- or attitude-guiding states of mind normative sentences express (e.g., Blackburn 1998, Gibbard 2003, Ridge 2014). This is plausibly understood as a thesis about *metasemantics*. That is, expressivism plausibly offers an explanation for why normative terms and sentences have the meaning that they do have, or an account of what it is in virtue of which they have their meaning.

Here the notion of expression is different from the pragmatic one mentioned above and illustrated with reference to generalized conversational implicatures. The expression relation, as understood here, in the context of expressivism, is not a matter of communicating a piece of information. Rather, it figures in the explanation of how normative language gets to have its meaning. What is it, then, for a sentence to express, in the relevant sense, some state of mind or attitude? Our project, here, is not that of developing an account of the expression relation. But perhaps something along the lines of the account proposed by Michael Ridge (2014, 109) will do for illustrative purposes:

A declarative sentence ‘*p*’ in sense *S* in a natural language *N* used with assertive force in a context of utterance *C* expresses a state of mind *M* if and only if conventions which partially constitute *N* dictate that someone who says ‘*p*’ in sense *S* in *C* with assertive force is thereby *liable* for being in state *M*.

This proposal requires some clarification. What is it to use a sentence with assertive force? Presumably, the idea is that using a sentence with assertive force just is a matter of using the sentence in such a way that one is liable or can be held accountable—in the light of the relevant linguistic conventions—for being in *some* state *M*, where this state must be of a certain type: a *belief-like* state or commitment, or a state or

commitment that has a *telos*, or function, of matching the world.<sup>1</sup>

Now, one might worry that this is not an expressivist-friendly proposal. For according to expressivism, normative sentences express practical states of mind or commitments, where these are to be somehow *contrasted* with representational beliefs, the *telos* or function of which is to represent the world as being this or that way. How, according to expressivism, could these sentences then be used with assertive force, if this is a matter of using a sentence in such a way that one is liable for being in a belief-like state, or for having a belief-like commitment? This kind of worry is easily disposed of. Expressivists should agree that the relevant states of mind or commitments are belief-like, in the relevant sense, despite not being representational beliefs the *telos* or function of which is to represent the world as being this or that way, normatively speaking. The relevant distinction can be drawn in many ways, but we do not wish to commit ourselves to any particular way of doing so, or to discuss this issue in any more detail here.<sup>2</sup> It suffices, for the present purposes, that expressivists may propose to understand the states or commitments expressed by normative sentences as being belief-like in some sense such that their expression amounts to an assertive use, and yet reject the view that these states or commitments could be understood simply as being in the business of representing the “normative bits of reality” (as Sinclair 2021 puts it).

Clearly more could and should be said by someone, somewhere, on what expressivists could and should say about the expression relation. But perhaps this suffices here (for more on the topic, see, e.g., Schroeder 2008, Sinclair 2021,

---

<sup>1</sup> Why “in sense *S*”? Presumably this is meant to restrict the set of assertive uses of the sentence, at issue, to those in which the sentence is used assertively in a certain type of way. So, one sentence, even when used assertively, might express one kind of state of mind in one type of use, and another kind of state of mind in another type of use. For example, “Allowing the use of Nazi symbols is wrong,” when used assertively and in a practical role, might express one kind of state of mind. And the same sentence might express a different kind of state when used, in a different kind of role, to report the mores of the surrounding society.

<sup>2</sup> For discussion, see, e.g., Sinclair 2007, Schroeder 2010, Ridge 2014.

Ch. 3). This rough account allows us to see how expressivism neatly captures the practicality of normative language. Normative language plays a practical role because that's what it's for. Normative terms mean what they mean because of their contribution to sentences that express practical attitudes, states, or commitments.

We have suggested that expressivism is best understood as a thesis in metasemantics, or as the thesis that the meaning of normative sentences is *explained* by facts about what states of mind these sentences express, in the relevant sense. By contrast, we may think of *semantics*, roughly, as giving a systematic account of what some set of sentences, *S*, mean, where this account is given in terms that we already understand. So, if we wanted to give a semantics for a snippet of the normative chunk of English, we could try saying, for instance, things such as the following:<sup>3</sup>

The meaning of "good" is a function from objects that gives the value *true* just in case the contextually relevant object has the property of goodness.

The meaning of "ought" is a function from agents and actions that gives the semantic value *true* just in case the contextually relevant agent has an obligation to perform the contextually relevant action.

A sentence such as "Knowledge is good" would, then, be true, just in case the thing picked out by "knowledge" – knowledge – has the property of goodness. And a sentence such as "Tove ought to turn down the volume" would be true just in case the person referred to by "Tove" would have an obligation to turn down the volume in the relevant context. Or perhaps we would rather wish to say something along the following lines:

The meaning of "good" is a function from objects that gives the value *true* just in case the contextually relevant object is highly

---

<sup>3</sup> Our formulations of the semantic theses below have drawn inspiration from the way some relevantly similar toy accounts are formulated in Chrisman 2016.

ranked by certain standards, where the relevant standards are fixed by the context.<sup>4</sup>

The meaning of “ought” is a function from a proposition that gives the semantic value *true* just in case the proposition is true in all of the worlds that are ranked highest by standards of a certain sort, among a set of worlds restricted in certain ways, where the relevant standards and restrictions are contextually determined.<sup>5</sup>

A sentence such as “Knowledge is good” would, then, be true, just in case the thing picked out by “knowledge” – knowledge – would be highly ranked by certain contextually determined standards. And a sentence such as “Tove ought to turn down the volume” would be true just in case Tove would turn down the volume in all of the worlds consistent with certain contextually determined background conditions and highly ranked by certain contextually determined standards.

To supplement our semantics for normative language, we could then add rules that allow us to use certain terms to create more complex sentences that have two or more simpler sentences as their parts. So, “and,” for instance, would contribute a function from propositions that gives the semantic value *true* just in case all of the relevant propositions are true. And so on.

Of course, this only gives us two candidate accounts of semantics for the tiniest of snippets of the normative chunk of English. Furthermore, these two accounts undoubtedly are poor candidates, too. But that’s OK. They are just toy models. The point, here, is simply that expressivists could sensibly think that *some* such truth-conditional semantics for the normative chunk of English can be given. Thus far, in providing a semantics for normative language, we wouldn’t have needed to invoke, for instance, anything like the idea that the ac-

---

<sup>4</sup> This is inspired by the account of the meaning of “good” in Ridge 2014, 26.

<sup>5</sup> This is a Kratzerian account of the meaning of “ought” (see Kratzer 2012), the formulation of which draws from Chrisman 2016, 86. For discussion of the semantics of “ought,” including Kratzer’s view, see also, e.g., Bronfman & Dowel 2018, Carr 2018.

ceptance of a normative sentence is a matter of representing the world as being a certain way, rather than a matter of having some desire-like attitude. We wouldn't have needed to appeal to any specific metaphysical account of the normative properties. The connection between semantics, in the sense outlined above, and many of the issues debated in metaethics doesn't seem to be especially close or direct (for a more detailed development of this point, see, e.g., Ridge 2014, Chrisman 2016, Sinclair 2021).<sup>6</sup>

We now are in a position to say some more about what expressivism might be taken to amount to. According to expressivism, again, the meaning of normative sentences is explained by their expressing certain suitably practical states of mind or commitments. That is, their meaning is explained by the fact that, when used assertively, the speaker of a sentence of the relevant kind can be held accountable, by linguistic convention, for being in some appropriate belief-like, yet practical, state of mind. On this kind of view, then, the sentence "Allowing the use of Nazi symbols is wrong" might express (appropriately belief-like) opposition to the use of Nazi symbols. This would be so in virtue of the fact that the relevant linguistic conventions would dictate that, when used assertively, the speaker of the sentence could be held accountable for being opposed (in an appropriately belief-like manner) to allowing the use of Nazi symbols.

## 2. The Problem of Diverse Uses

The appeal to the practical role of normative language in explaining its meaning gives rise to a challenge which we may, following Pekka Väyrynen (2022), call *the problem of diverse uses*. The challenge is based on the simple observation that not all uses of sentences that contain these terms seem to play

---

<sup>6</sup> The point is merely that expressivism seems compatible with such semantic views. This is not at all to suggest that expressivists thereby escape having to deal with various metaphysical issues. For instance, given that there really are normative properties, it seems fair to ask the expressivist what such properties are like (whether they are, for example, *sui generis*, or reducible to properties that may also be ascribed by descriptive judgments; see, e.g., Bex-Priestley forthcoming).

a suitably practical role. Väyrynen (2022) characterizes the problem as follows:

If the practical role of these terms were a part of their conventional profile in a language, it should not be subject to [...] exceptions but instead should be present in all literal uses in normal contexts. This raises what I will call the Problem of Diverse Uses: How do you reconcile the diversity of uses to which [...] normative terms may be put with the claim that their association with their normative roles is broadly semantic? The problem prompts a challenge: either offer some plausible explanation of cases where the relevant practical upshots are absent that reconciles these claims, or else do not build such upshots into our overall semantic theory for [...] normative terms (182–183).

We should look at some examples of candidates for non-practical uses of “ought.” Väyrynen (2022, 182–183) offers a useful selection. Consider, then, the following:

- (1) One ought to prioritize profit over fairness. But is that really the thing to do?

In (1), the ought-claim may, in a suitable context, be rightly understood as a claim about what follows from capitalist values or standards. It might be clear from the preceding discussion, or from the pins on the speaker’s jacket (Väyrynen 2022, 192), that such standards are not the speaker’s standards. Plausibly, the ought-claim in (1) need not, then, play any practical role for the speaker. Or consider:

- (2) Client: What is my legal obligation, and what do you expect me to do?

Lawyer: You have to report your liability, but I do not know if you will; you may prefer to push the limits of the law and just conceal it.

Here the client and the lawyer discuss legal oughts and musts, but they might not end up giving such oughts and musts any weight in their practical deliberation. Väyrynen (2022, 182–183) provides more examples:

- (3) It would be wrong to kill. But I’m ok with killing and do not feel bad about it.



- (4) I ought to finish grading. I have absolutely no intention to do so, though.
- (5) I should do the shopping today (as far as I know).

In all these cases, we may, given a suitable context, understand the relevant normative terms (“wrong,” “ought,” “should”) in relation to standards that may not engage the speaker, or have any motivational significance for them.

Now, one thing that the expressivist could say is, of course, that “ought” means different things in different contexts. Sometimes it has a descriptive meaning. In these descriptive uses, ought-claims just report how things are ranked according to certain descriptively specifiable standards. Such claims need not play any practical role. In other contexts, though, “ought” has a very different kind of, practically charged, meaning. Expressivists have indeed sometimes suggested just this. A. J. Ayer (1936, 105–106), for instance, distinguishes between the “normative ethical symbols” and the “descriptive ethical symbols,” which are “commonly constituted by signs of the same sensible form,” but make a very different contribution to the meanings of sentences. However, this is not a promising route for the expressivist. It is a striking feature of normative terms such as “good,” “ought,” and “reason,” that they *all* have both practical and non-practical uses, and that the patterns of their use are very similar across a range of languages. It is incredible that this would be due to normative terms being simply ambiguous. (Mackie 1977, 51, Chrisman 2016, Ch. 2.3; for warnings regarding a reckless postulation of lexical ambiguities, see also, e.g., Thomson 2008, Finlay 2014, Wodak 2017).

We do not wish to reject the possibility that normative words would turn out to be ambiguous in their meaning, possibly in a variety of interesting ways. But it does seem that the working hypothesis and the default position should be that there is significant unity to the meaning of the various uses of normative terms. In particular, and importantly in the present context, normative terms such as “ought” plausibly have a unified meaning across both practical and non-practical uses.

The desirability of a rather unified account of the meaning of normative terms provides one important reason for why it

makes sense to understand expressivism as a view in metasemantics.<sup>7</sup> For a metasemantic construal of expressivism creates space for a response to the problem of diverse uses. If expressivism is understood as a view in metasemantics, space opens for the following possibility: perhaps, in some cases, normative terms mean what they mean because of their contribution to sentences that express a practical state of mind; perhaps, in some other cases, normative terms have this same meaning because of their contribution to sentences that express another kind of, non-practical, state of mind.

For example, in the previous section, we mentioned the following toy semantics for “ought”:

The meaning of “ought” is a function from a proposition that gives the semantic value *true* just in case the proposition is true in all of the worlds that are ranked highest by standards of a certain sort, among a set of worlds restricted in certain ways, where the relevant standards and restrictions are contextually determined.

If this kind of semantics for “ought” is correct, then, when someone says, for example, that we ought to ban the use of Nazi symbols, their statement means that the use of Nazi symbols is banned in all of the worlds compatible with two contextual restrictions: First, the *modal base* restricts the set of worlds we are considering to those worlds that are compatible with whatever background conditions are determined by context *c* (for instance, those worlds in which it is possible for us to use Nazi symbols). Second, the *ordering source* further restricts the relevant set of worlds to those worlds which are ranked as best by some ordering over worlds, in accordance with whatever standards are determined by *c* (for instance, those worlds in the modal base that are best according to the correct standards of practical reason, or to give another ex-

---

<sup>7</sup> It's not the only one. We have already noted that going metasemantic seems to allow expressivists to adopt non-revisionary, standard views in semantics, which is nice also for other reasons. There are also reasons for understanding expressivism as a metasemantic view that are not narrowly semantic. One such reason, having to do with our knowledge of normative supervenience, is given in Venesmaa 2021.

ample, those worlds that are best according to the moral standards accepted in the speaker's community).

Now, this is all very toy-ish. But that's fine. What is important here is that expressivists can adopt some such semantic story with regard to "ought" (and the rest of normative language). Or better: it's not instantly obvious that they cannot do so. For they can suggest that while statements about what ought to be done, or about what ought to be, more generally, always have a meaning of this sort—a meaning captured by something like the toy account above—their expressivist account of normative language is entirely compatible with this. They may propose that expressivism helps to explain why the "ought"-sentences have the kind of meaning that they do have. More precisely, they may propose that the expressivist account explains whatever meaning "ought"-sentences have, according to the unified semantics, in some contexts, but not in others. This is why some of the "ought"-sentences are practical while others aren't. Or that's what metasemantic expressivism allows one to say.

This, by itself, is a very abstract point about the kind of structure that an expressivist proposal might take. It is one thing to point out that this kind of structure is, in principle, available to be utilized. It is another thing, entirely, to outline an expressivist view that has this kind of structure and that would be attractive. In the next two sections, we first present two challenges for the kind of expressivist response to the problem of diverse uses, according to which normative terms have a unified meaning, but this meaning is given an expressivist explanation only when the relevant terms are used in a practical role. We then articulate an expressivist view that is well-positioned to exploit the availability of this kind of response.

### **3. The problem of disunified metasemantics**

The idea that we might be able to combine a fairly unified semantics for normative terms—one that captures a wide range of both practical and non-practical uses—with an expressivist metasemantics that only explains the meaning of normative terms in some of their uses is not a new one. The space for this kind of move has been explored before (see, in

particular, Ridge 2014). But the idea has met with some skepticism. We next wish to address the interesting articulations of such skepticism by Pekka Väyrynen (2022) and Daniel Wodak (2017).

Väyrynen argues that the existing expressivist metaseantics for normative language “do not support the claim that the practical role of such language is a distinctive and particularly significant feature of its meaning” (2022, 200). He reaches this conclusion through considering the different ways in which the practical role of normative language might figure in its metaseantics. Let us suppose, then—along with Väyrynen—that “ought” has a Kratzerian semantics such as the one that we have used as our toy semantics above (cf. Väyrynen 2022, 189–190). There are two options, Väyrynen suggests, with regard to understanding the nature of the metaseantic work that is done by the idea that normative sentences (sometimes) play a practical role. The first one “has to do with the metaseantics of the context-sensitivity of *ought*”:

Perhaps its practical role contributes to explaining its semantic value specifically in its committal uses. [...] Whether a use is committal or not is a difference in context. We might then think that when *ought* is used in a committal way, this can make a difference to the values of its contextual parameters [the modal base and the ordering source]. In this way, the practical role of *ought* might contribute to explaining its semantic value in some cases but not others. (Väyrynen 2022, 190–191)

This is, indeed, one kind of metaseantic work done by the practical role of normative language. Whether an “ought”-sentence plays this kind of role contributes to determining its semantic value in a context. However, we agree with Väyrynen that this isn’t all the metaseantic work that expressivists should take to be done by the idea that normative language sometimes plays a practical role. The expressivist idea is not merely that the practical role of normative language sometimes helps to determine the values of the contextual parameters of an “ought.” This much could be agreed upon, for instance, by someone, according to whom the meaning of “ought”-sentences is always to be explained by their being expressive of robustly representational beliefs

about how various scenarios compare with regard to some contextually specified standards. This kind of representationalist, non-expressivist metasemantics could be combined with the idea that the relevant standards are sometimes determined by “ought” playing a practical role for the speaker in the context. The occasionally practical role of “ought” could then, in this way, contribute to determining the ordering source, and the semantic value, for some uses of “ought.”

However, this kind of work in the explanation of meaning would not be everything that an expressivist needs from the practical role of normative language. Rather, according to a metasemantic expressivist, “ought” (for example) has the kind of semantics that it has (a Kratzerian semantics, say) because it plays a practical role. This would seem to correspond to Väyrynen’s second proposal with regard to how the metasemantic significance of the idea that normative language plays a practical role could be understood. On this proposal, as Väyrynen puts it, the “practical role of *ought* is part of what explains why the dominant sort of formal models for modal language provide a good descriptive semantics for terms like *ought* in the first place.”

It is important to emphasize, though, that this proposal should not be understood as suggesting that the meaning of “ought” is always, in every context, explained by its serving a practical role. We have granted, in the previous section, that the uses of “ought” are diverse, and not always practical. So, the expressivist proposal should be that the practical role of “ought” is part of what explains why normative language has the kind of semantics that it has *in those cases in which it does play a practical role*. The thought would be, then, that the meaning of “ought” remains constant across both practical and non-practical uses, but is only explained (in part) by the practical role of “ought” in the uses of the former sort. In both sorts of uses, the given semantics would be a good model for “ought” “because it appropriately mirrors the structure of mental states that *ought* expresses” (Väyrynen 2022, 193). In the non-practical uses, the meaning of “ought” would be explained by the fact that the relevant sentences express certain representational beliefs or commitments; in the practical uses, it would be explained by the fact that the relevant sentences express practical commitments or states of mind.

Väyrynen is not happy with this kind of suggestion. He writes:

Explaining noncommittal uses only requires invoking theoretical commitments and cognitive states. By parity, that should suffice also for explaining committal uses. The standard semantics does not care about this distinction between uses. So, on the face of it, explaining why it is a good model for *ought* should not require invoking practical role. (It really is dialectically significant if committal and noncommittal uses of *ought* are uniform in their descriptive semantics!) If that is right, it would complete my case that nothing in the standard semantics for *ought* supports treating those uses that are associated with a practical role as semantically or metasemantically exceptional. (Väyrynen 2022, 193)

We have granted that in noncommittal or non-practical uses, the meaning of “ought” can be explained without appealing to the idea of a practical role. Väyrynen suggests that since it must then be possible to explain the basic semantic structure of “ought” without invoking practical role, and since we must, in any case, do so in the case of non-committal uses, we should also do so in the case of committal or practical uses. Why? Presumably, because this is the option favored by considerations of simplicity and uniformity; given that two metasemantic views do an equally good job in explaining why the standard Kratzerian view provides (what we are assuming is) a good model for “ought,” but one is more simple and unified in that it doesn’t invoke different kind of mental states in explaining its practical and non-practical uses, then this is a point in its favor.

As we understand Väyrynen’s objection, the same objection is raised also by Wodak (2017) who targets Ridge’s (2014) attempt to formulate an expressivist metasemantic theory that would vindicate unified semantics for “ought” and other normative terms:

First, Ridge must concede that there is a viable non-expressivist explanation of why “ought” means *Z* in a wide variety of uses. This explanation is non-expressivist insofar as it appeals to robustly representational beliefs. And it is viable in that it explains: why “ought” means *Z*; how context selects the relevant

ordering source; how competent speakers use “ought” to communicate, coordinate, and collect information; and how speakers disagree even in the face of systematic differences in their criteria for applying words (like “legally ought”). Once that viable non-expressivist explanation is on the table for some uses, why not offer it across the board? A unified meta-semantics is preferable, if only for the sake of parsimony. (Wodak 2017, 284)

However, what we want is not just the most simple and uniform metasemantic theory that explains why the meaning of “ought” has the kind of structure that it has on the Kratzerian view. Rather, what we want is the most simple and uniform metasemantic theory that can explain this and the other things that a metasemantic view should explain. For example, a plausible metasemantic theory should account for the data concerning normative disagreement. When someone accepts the sentence “We ought to ban Nazi symbols” and someone else accepts the sentence “We ought not to ban Nazi symbols,” this usually constitutes a disagreement. Different metasemantic views face different challenges in explaining why this is so, and the difficulty of the relevant challenges does not track the degree of simplicity and unity that such views enjoy in relation to explaining the semantic structure of “ought.” For instance, contextualist views, according to which ought-judgments express robustly representational beliefs about how things relate to certain contextually specified standards, may offer a very neat and simple explanation for why the meaning of “ought” would have the Kratzerian structure across different kind of uses, but struggle to accommodate the data concerning when we agree or disagree about normative issues (see, e.g., Finlay 2017).

In addition to the data concerning disagreement, the right metasemantic view also needs to explain whatever it is that needs explaining in relation to, for instance, the intuitions that fuel the open question argument, the relationship between normative judgment and motivation, and our plausibly conceptual knowledge of the supervenience of the normative on the descriptive (Venesmaa 2021). Undoubtedly there’s much more that needs explaining. But this sample suffices to make it clear that we shouldn’t be too quick to rule out the idea that capturing what needs to be captured by a metasemantic theory may require some complexity in the

metasemantics. This is not to say that we should make sacrifices with regard to simplicity and uniformity. We just need the most simple and unified theory that captures all the relevant phenomena.

This brings us naturally to Wodak's second objection to trying to occupy the space—one that we have proposed is available for an expressivist—of combining distinct metasemantic stories for practical and non-practical uses of normative terms with a unified semantic account.

Second, consider how [the] non-expressivist explanation interacts with its expressivist counterpart. Here [the metasemantic expressivist] is committed to an unexplained coincidence. There is one *explanandum*: that the word "ought" means Z. There are two radically different *explanantia*; the expressivist, after all, is emphatic about the differences between representational beliefs and non-representational conative states. If the *explanantia* are radically different, why is the *explanandum* exactly the same? Why don't the radical differences between the states that we are expressing translate to differences in meaning? And, relatedly, why would we employ one word to express such radically different mental states? (Wodak 2017, 284–85)

This second objection from Wodak doesn't concern the complexity of the expressivist's metasemantic account as such. Instead, the worry is that the expressivist's account leaves a striking coincidence completely unexplained. In order to appreciate the force of this objection, it is helpful to first consider an example of a very simple expressivist view.

According to what we may call *simple expressivism*, "ought"-sentences express desire-like states of mind that are quite different from the belief-like states of mind expressed by descriptive sentences. So, whereas a sentence such as "The use of Swastika by the Finnish Air Force is not historically unrelated to the use of Nazi symbols" expresses a belief, the job of which is to describe and match the way the world is, a sentence such as "We ought to ban Nazi symbols" expresses a different kind of state of mind, perhaps opposition to allowing the use of Nazi symbols.

We may now try combining this simple expressivist view with the attempt to explain the same semantic structure on the basis of different metasemantic accounts of practical and



non-practical uses of normative language. If we do this, we end up being committed to the idea that even though the practical and the non-practical uses of “ought” express radically different states of mind, both kind of uses involve the very same meaning for “ought,” where this uniform meaning is supposedly explained on the basis of this meaning somehow mirroring the structure of the very different mental states expressed by these sentences. Again, the problem, here, is not that the dualist metasemantics would be too complex. Rather, the problem is that the idea of a dualist metasemantics that appeals to very different types of states of mind and yet yields a completely unified semantic structure for “ought” commits one to an unexplained coincidence. An acceptable metasemantic theory doesn’t tolerate this.

It’s a good problem. It seems like a devastating problem for simple expressivism. There is no hope for an expressivist metasemantic view, such that would avoid the commitment to an unacceptable, unexplained coincidence, unless more structure—more structure suitable for being mirrored by the semantics of “ought”—is introduced in the expressivist’s account of the states expressed by normative sentences. This is a very interesting result. But there is a further interesting result that can be obtained here, namely, that there is a brand of expressivism that plausibly has the resources for providing the kind of structure that is needed. We next present an expressivist view that has this nice feature: relational expressivism.

#### **4. Relational expressivism and the problem of unexplained metasemantic coincidence**

Let us suppose that Alex and Blue both accept that Nazi symbols ought to be banned. Here’s what we believe is a plausible idea: Alex’s and Blue’s accepting that Nazi symbols ought to be banned is, very roughly, a matter of their being opposed to actions that have this or that property—who knows which one or which ones—and of their believing that failing to ban Nazi symbols has a relevant property (whatever property it is). Or perhaps we could say that Alex’s and Blue’s holding their view regarding the banning of Nazi symbols is, very roughly, a matter of their being opposed to

failing to ban Nazi symbols on the grounds of such failure's having some property that they would treat as relevant.

Alex and Blue may have very different normative perspectives. Perhaps Alex is a utilitarian who thinks that failing to ban Nazi symbols doesn't maximize happiness, and who is therefore opposed to failing to ban them. Perhaps Blue is a Kantian who thinks that failing to ban Nazi symbols is not compatible with the Categorical Imperative, and who is therefore opposed to not banning them. Or perhaps Alex and Blue are normal human beings and neither is very articulate about what their respective normative perspectives are like. Maybe Alex is opposed to things that are vaguely such and such—like *that* (mentally pointing, so to speak, toward *these* actions and policies), whereas Blue is opposed to things that are vaguely thus and so—like *that* (mentally pointing, so to speak, toward *those* actions and policies instead of these) (for appeal to the mental demonstratives of this sort, see Ridge 2014).

In any case, Alex and Blue are both opposed to things that have some—these or those—properties, and they both believe that failing to ban Nazi symbols has a relevant property. Even though their normative perspectives differ, they share an interesting similarity. They both are opposed to some type of actions—these or those—and believe, of the property that grounds their attitudes of opposition, respectively, that failing to ban Nazi symbols has that property. Their desire-like states of opposition and their suitably related representational beliefs (concerning Nazi symbols, in this instance) are related in the same way. They both are in the very same type of *relational state*, we may say, where this relational state is multiply realizable and differently realized by having some such desire-like state and a representational belief that are related in the relevant way.

As noted, we think that it is plausible to think that Alex's and Blue's holding the normative views that they hold is, very roughly, a matter of their being like this—a matter of their sharing this type of relational state. We also find it plausible that the sentence "We ought to ban Nazi symbols" expresses a relational state of roughly this kind. According to relational expressivism, normative sentences express states of this kind (see Schroeder 2013, Toppinen 2013, Ridge 2014).

Relational expressivism holds promise with regard to providing a suitably structured account of the states expressed by normative sentences – one that has the right kind of structure for the Kratzerian semantics to mirror. Above, we have characterized this kind of semantics as follows:

The meaning of “ought” is a function from a proposition that gives the semantic value *true* just in case the proposition is true in all of the worlds that are ranked highest by standards of a certain sort, among a set of worlds restricted in certain ways, where the relevant standards and restrictions are contextually determined.

So, for example, the sentence “We ought to ban Nazi symbols” would be true just in case, among a set of worlds restricted in certain ways (e.g., to those in which we are able to ban Nazi symbols), in all of the worlds that are ranked highest by some relevant standards, we ban Nazi symbols. How would this kind of semantic structure mirror the structure of the mental states that are, according to relational expressivism, expressed by “ought”-sentences?

There are two options that a relational expressivist might pursue here. First, it is quite plausible that the representational beliefs that partly realize the relational states of mind expressed by normative sentences (e.g., “ought”-sentences) are beliefs concerning standards. When we consider the non-practical uses of “ought”-sentences (about the requirements of etiquette, say), it is very natural to think that such sentences express representational beliefs about what is required (etc.) by certain standards. The expressivist may simply propose that this is a part of what’s going on in the case of the practical uses, too.

Some expressivists have adopted this kind of idea. Ridge (2014, Ch. 1), for instance, suggests that claims about what’s good express states that involve beliefs about what is highly ranked by certain standards. Likewise, claims about what ought to be done or about what must be done express states that involve beliefs about what is recommended or required by certain standards, and so on. That something like this is correct is, again, quite plausible. When we judge something to be good, it seems that we may always ask “By what standards?” When we classify actions as the ones that ought or

must be performed, we are committed to there being some grounds for why these are the actions that ought to, or must, be performed. And it is natural to think that we are thereby committed to there being some standards that help to articulate the relevant grounds.<sup>8</sup>

Above, we have suggested that when Alex and Blue accept that we ought to ban the use of Nazi symbols, this is a matter of their being opposed to things that have some—these or those—properties, and of their believing that failing to ban Nazi symbols has a relevant property. Perhaps, we said, Alex is a utilitarian who thinks that failing to ban Nazi symbols doesn't maximize happiness, and who is therefore opposed to failing to ban their use. And perhaps, we said, Blue is a Kantian who thinks that failing to ban Nazi symbols is not compatible with the Categorical Imperative, and who is therefore opposed to not banning them. Assuming that this is so, we may now understand this, somewhat more specifically, as follows: when Alex accepts that we ought to ban the use of Nazi symbols, this is a matter of their being opposed to actions that are not in accordance with the utilitarian standard (which requires that we maximize happiness), and of their believing that failing to ban Nazi symbols is not in accordance with this standard; when Blue accepts that we ought to ban the use of Nazi symbols, this is a matter of their being

---

<sup>8</sup> Does this rule out some forms of *particularism* that question the centrality of standards or principles to normative thought and talk? We think that particularists should agree that we are, in normative thought and talk, committed to the distribution of normative properties necessarily being determined by some necessarily true descriptive-to-normative principles, even if such principles are of no epistemic help. However, if the expressivist does need the standards to play a more ambitious role in normative thinking, then this may constitute a conflict with some interesting forms of particularism (for discussion on particularism, see, e.g., Dancy 2004, McKeever & Ridge 2006). Ridge (2014, 43–44) notes this issue in the context of his own favored formulation of a relational expressivist view and makes the interesting point that given that the appeal to standards is motivated by “quite general considerations in semantics for words like ‘ought’ that cut across both normative and non-normative uses,” this “puts pressure on the particularist to offer alternative unified semantics or defend a kind of ambiguous view of words like ‘ought’,” where neither of these moves seems “terribly promising.”

opposed to actions that are not compatible with the Kantian standard (which requires that we act in accordance with the Categorical Imperative), and of their believing that failing to ban the use of Nazi symbols is not in accordance with that standard.

Or consider the option of understanding Alex and Blue, more realistically, as being somewhat inarticulate about which properties are normatively relevant, or about what standards they endorse and reject. This, too, may be understood in terms of their endorsing or opposing standards of a certain kind. When Alex accepts that we ought to ban Nazi symbols, this can be understood to be, in part, a matter of their being opposed to actions that are *like that*, where having the relevant property (that is, being “like that”) amounts to being incompatible with certain standards. Which ones? Well, those that rule out, for instance, actions “like that.” In Alex’s judgment that we ought to ban Nazi symbols this state of opposition then combines with a belief that failing to ban Nazi symbols has the relevant property – that is, with their belief that failing to ban Nazi symbols it not compatible with standards such that rule out, for example, actions “like that” (whatever actions they think of as being “like that”).

In any case, no matter to what extent Alex and Blue are articulate about the standards that they have adopted, their accepting that we ought to ban Nazi symbols will be a matter of their being opposed to actions that do not meet standards of a certain kind, and of their believing that failing to ban Nazi symbols is not in accordance with the relevant sort of standards.

Above, we noted that there are two ways in which one might suggest, in line with relational expressivism, that the semantics of “ought” mirrors the states expressed by “ought”-sentences. The first was to appeal to the idea that “ought”-sentences express relational states that are always realized, in part, by representational beliefs concerning what is recommended or required by some standards. The second option that could perhaps be pursued here is the following. Instead of saying that the relational states expressed by normative sentences are always realized by beliefs that provide the structure for the Kratzerian semantics to mirror, one could suggest that the relational states themselves provide

the relevant structure. The idea would be that while the representational beliefs that partly realize these relational states need not concern standards, the relational states that they partly realize do so. Let us return, for example, to our first characterizations of the relational expressivist idea. We first proposed that when Blue accepts that we ought to ban the use of Nazi symbols, this could be construed, roughly, as a matter of their being in a state of being opposed to actions that have some particular property – that of being incompatible with the Categorical Imperative, say – and of believing that failing to ban the use of Nazi symbols has that property. Perhaps one could now propose that being opposed to actions that have some particular property amounts to acceptance of a standard. And perhaps one could then propose that if one now, in addition to accepting some standard, has a belief to the effect that some action has some property that suitably relates to the standard in question (e.g., has a property such that actions with that property are ruled out by the standard), then one's states of opposition and belief will be related to each other in a way that constitutes a judgment the content of which concerns a standard. In this way, even if the representational belief that partly realizes the relational state expressed by a normative sentence would not concern the relation of anything to any standard, the relational state itself could be taken to have the functional profile of a judgment or a belief that does concern such things. We shall not say more about this kind of idea here. But this kind of option may nevertheless be worth keeping in mind.

If the kind of relational expressivist metasemantics outlined above is correct, then it seems unsurprising that the practical uses of "ought"-sentences have the kind of semantics that we have assumed they have. The semantics now nicely mirrors the structure of the states of mind expressed by "ought"-sentences. We haven't said anything about *how*, exactly, the account of the states expressed by the relevant sentences explains the semantics. The talk of "mirroring" isn't perhaps very satisfying, ultimately. However, this seems acceptable in the present context. We have been operating here with the assumption that the broader expressivist project of explaining meaning in terms of the states of mind that sentences express is workable. This idea can of course be contest-

ed. However, in the present context, we have not been concerned with defending this idea. And the problems we have been addressing—the problem of diverse uses, the problem of disunified metasemantics, and the problem of metasemantic coincidence—are not supposed to be problems for this broader project of explaining the meaning of sentences in terms of the states of mind that they express. Instead, these latter problems have been raised for expressivism against the background of assuming, at least for the sake of the argument, the legitimacy of the expressivist metasemantic project.

In the context of this project, we may assume that if “Snow is white” expresses a belief that snow is white, this explains the meaning, or the truth-conditions, of the sentence. Likewise, in the context of this project, we may assume that if “One ought not to eat peas with a spoon” expresses a belief about what is required by certain standards, this explains the meaning, or the truth-conditions, of the sentence. That is, we may assume that this explains why the sentence (as uttered in an appropriate context) is true just in case peas are left uneaten in all the worlds restricted by certain background conditions and ranked highest by the standards of etiquette. If this is all correct, and if “We ought to ban the use of Nazi symbols” turns out to express the kind of relational state that this sentence expresses, according to relational expressivism, then this plausibly explains why this sentence, too, has a meaning similar to that of “One ought not to eat peas with a spoon.” That is, the relational expressivist account explains why this sentence, too, says of something—of banning the use of Nazi symbols—that that is what is done in all the worlds the set of which is restricted in certain ways and ranked in accordance with some suitable standards.

It’s worth emphasizing that, according to relational expressivism, when the sentence “We ought to ban the use of Nazi symbols” is used in a practical role, it does not *express* any representational belief concerning standards. Rather, the sentence expresses a relational state that is differently realized by different desire-like states and representational beliefs in different contexts of use. So, when Alex the utilitarian accepts “We ought to ban the use of Nazi symbols,” the sentence, as used in the relevant context, does not express a be-

lief that failing to ban the use of Nazi symbols is not in accordance with the utilitarian standards. These standards need not be, and are likely not to be, contextually specified, in this instance, in the way that the standards of etiquette might be specified in a context in which “One ought not to eat peas with a spoon” is used. What is contextually specified is just that what’s in play are the standards relevant to the practical use of “ought,” the standards of practical reason, we might say (with Ridge 2014), or of what to do—the “genuinely” or “robustly” normative standards. Whether these standards are utilitarian, or Kantian, or something completely different, is left for practical deliberation or normative theorizing to decide.

## 5. Conclusion

Expressivists wish to explain the meaning of (some of) the normative language by appealing to the practical role that such language plays. We have here addressed a problem for the expressivist proposal, *the problem of diverse uses*, which arises from the fact that normative terms often figure in sentences that do not play any interestingly practical role. The challenge is that of explaining, in the face of this fact, the meaning of normative language in a sufficiently unified manner.

We have proposed, as the first step toward responding to this challenge, that expressivism should be understood as a view in metasemantics. Expressivists may then suggest that they can make sense of normative terms as having a unified meaning across the practical and non-practical uses of normative language. It’s just that this unified meaning is sometimes explained by the practical role of normative language, and sometimes by its non-practical, representational role.

This kind of move gives rise to two further problems. First, *the problem of disunified metasemantics* draws our attention to the fact that the expressivist metasemantics is somewhat more complex and disunified than some of its alternatives. We have granted that we should seek a metasemantics for normative language that is as simple and unified as is possible, given that it allows us to explain all the semantic data that requires explanation. However, the practical uses of



normative language do differ, in many ways, from the non-practical uses. And some of the differences plausibly have to do with the meaning of the language used. We have suggested that the phenomena that are relevant here—the disagreement data, “open question” intuitions, etc.—may very well justify, or indeed require, some complexity in the metasemantics. We believe that expressivism is well-placed to capture complexity of the relevant kind, but determining whether or not this is so is beyond the scope of this discussion.

Second, though, *the problem of unexplained metasemantic coincidence* also needs to be addressed. Given the expressivist idea that the practical and non-practical uses of normative language express importantly different kind of states of mind, the expressivist account would seem to be committed to it being completely coincidental, and wholly unexplainable, that normative terms should have the very same core meaning across both practical and non-practical uses. We have granted that this seems like a devastating problem for a simple expressivist view, according to which normative sentences sometimes express representational beliefs (about, say, the requirements of the norms of etiquette), and sometimes desire-like attitudes of a wholly different sort (disapproval of the use of Nazi symbols, for example). However, there is no unexplained coincidence in how the meaning of normative terms is explained across their different uses, given the truth of one kind of expressivist view, relational expressivism. According to this view, normative sentences always express states of mind that involve representational beliefs relating things to certain standards. This offers promise with regard to providing us with a metasemantics that has a unified enough structure for a unified semantics to mirror.

Is there still some question that we would have alluded to, but that would have been left unaddressed? One of Wodak’s worries concerns the expressivist’s resources with regard to explaining why we would employ *one word* to express radically different mental states. Given the relational expressivist view, we may now replace “radically different” with “some-what different.” Still, one might wonder why it should be that judgments involving the term “ought,” or a word that “ought” translates to, would express two different kinds of judgments, practical and non-practical, about what is re-

quired or recommended by certain standards. We have not directly addressed this particular question. But as we see it, given the metasemantics for “ought” roughly outlined in the previous section, there shouldn’t be anything terribly surprising or mysterious about the fact that we use one word for expressing judgments or mental states of a somewhat different kind. The explanation for why we would have a word for expressing the relevant kind of mental states plausibly has to do with the usefulness, or necessity, of standards in the guidance of action and attitude management. Plausibly, the practical uses of “ought” have *priority* over the non-practical uses, in that the idea of a community that would only use “ought” in non-committal or non-practical ways seems very strange, while the converse doesn’t seem to hold. We need the practical or committal oughts to guide and coordinate our actions and attitudes. But we are also bound to have an interest in the ways in which those surrounding us guide and coordinate their actions and attitudes, even if their commitments differ from those of ours. And we are also bound to have an interest in the various possible ways of guiding and coordinating actions and attitudes. It would make perfect sense, then, to use the same words for relating things to various standards for choice and belief (etc.), regardless of whether the relevant language would have a directly practical use for us, or instead be used in tracking some “standard-involving” facts in the absence of a direct practical concern.

Even supposing that we have provided satisfying responses to the problems of diverse use, disunified metasemantics, and unexplained metasemantic coincidence, much more work remains to be done on related issues. For example, it would be nice to have an account of how, exactly, the expressivist metasemantics explains semantics. Also, it remains to be determined (as far as we know) what the correct semantics for the various normative terms is like. We have only toyed with one toy view regarding the meaning of one normative term (“ought”). Plausibly, we will only be in a position to figure out what the expressivist explanation for the meaning of various normative terms looks like once we know what the right semantics is for these terms. However, the relational expressivist view allows for a lot of variation in how, exactly, the relational states expressed by normative sentenc-

es should be understood. This makes us optimistic that whatever the right semantics for normative language is going to be, relational expressivism will provide interesting resources for finding the right kind of explanation for it. At the very least the challenges that we have here addressed give us no reason to be skeptical about the prospects of this brand of expressivism.<sup>9</sup>

*Tampere University  
University of Helsinki*

## References

- Ayer, A. J. (1936). *Language, Truth, and Logic* (2<sup>nd</sup> edition, 1946). London: Gollancz.
- Bex-Priestley, G. (forthcoming). "Expressivists Should Be Reductive Naturalists." In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*. Oxford: Oxford University Press.
- Blackburn, S. (1998). *Ruling Passions: A Theory of Practical Reasoning*. Oxford: Oxford University Press.
- Bronfman, A. & J. Dowel (2018). "The Language of 'Ought,' and Reasons." In D. Star (ed.), *The Oxford Handbook of Reasons and Normativity*. Oxford: Oxford University Press.
- Carr, J. (2018). "Deontic Modals." In T. McPherson and D. Plunkett (eds.), *The Routledge Handbook of Metaethics*. New York: Routledge.
- Chrisman, M. (2016). *The Meaning of 'Ought': Beyond Descriptivism and Expressivism in Metaethics*. New York: Oxford University Press.
- Dancy, J. (2004). *Ethics without Principles*. Oxford: Oxford University Press.
- Finlay, S. (2014). *Confusion of Tongues: A Theory of Normative Language*. Oxford: Oxford University Press.
- Finlay, S. (2017). "Disagreement Lost and Found." In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics, Vol. 12*. Oxford: Oxford University Press.
- Fletcher, G. (2014). "Moral Utterances, Attitude Expression, and Implicature." In G. Fletcher & M. Ridge (eds.), *Having It Both Ways: Hybrid Theories and Modern Metaethics*. Oxford: Oxford University Press.

---

<sup>9</sup> We thank an anonymous reviewer for a helpful set of comments on an earlier version of the paper.

- Gibbard, A. (2003). *Thinking How to Live*. Cambridge, Mass.: Harvard University Press.
- Kratzer, A. (2012). *Modals and Conditionals: New and Revised Perspectives*. Oxford: Oxford University Press.
- Mackie, J. L. (1977). *Ethics: Inventing Right and Wrong*. London: Penguin.
- McKeever, S. & M. Ridge (2006). *Principled Ethics: Generalism as a Regulative Ideal*. Oxford: Oxford University Press.
- Ridge, M. (2014). *Impassioned Belief*. Oxford: Oxford University Press.
- Schroeder, M. (2008). *Being For: Evaluating the Semantic Program of Expressivism*. Oxford: Oxford University Press.
- Schroeder, M. (2010). *Noncognitivism in Ethics*. Abingdon: Routledge.
- Schroeder, M. (2013). "Tempered Expressivism." In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics, Vol. 8*. Oxford: Oxford University Press.
- Sinclair, N. (2007). "Propositional Clothing and Belief." *The Philosophical Quarterly* 57, 342–362.
- Sinclair, N. (2021). *Practical Expressivism: A Metaethical Theory*. Oxford: Oxford University Press.
- Strandberg, C. (2011). "A Dual Aspect Account of Moral Language." *Philosophy and Phenomenological Research* 84, 87–122.
- Thomson, J. J. (2008). *Normativity*. Chicago: Open Court.
- Toppinen, T. (2013). "Believing in Expressivism." In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics, Vol. 8*. Oxford: Oxford University Press.
- Venesmaa, V. (2021). "Explaining Our Knowledge of Normative Supervenience." In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics, Vol. 16*. Oxford: Oxford University Press.
- Väyrynen, P. (2022). "Practical Commitment in Normative Discourse." *Journal of Ethics and Social Philosophy* 21, 175–208.
- Wodak, D. (2017). "Expressivism and Varieties of Normativity." In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics, Vol. 12*. Oxford: Oxford University Press.