Deflating the Many Attitudes Problem[1]

Derek Baker

Note: This is a penultimate draft of a paper forthcoming in Philosophical Quarterly. Please quote or cite the published version.

Abstract

Attitudinal embeddings, such as 'I hope that murder is wrong' or 'she is glad that eating meat is not wrong' are a less substantial problem for expressivists than is standardly thought. If expressivists are entitled to talk of normative beliefs, they can explain what it is to for an attitude to be semantically related to a normative content in terms of being functionally related to a belief with a normative content.

Keywords: Expressivism, Quasirealism, Frege-Geach Problem, Many Attitudes Problem

1. Introduction

'I hope that murder is wrong'. 'He's glad that murder is wrong'. 'She wonders if eating meat is wrong'. Attributions like these are thought to pose a serious problem for expressivists, being especially tricky looking instances of the general problem of embeddings (M. Schroeder 2008 and 2013; and Carr 2015). What is the person who hopes that murder is wrong hoping for, if the normative term 'wrong' is a device for

expressing disapproval? Is the attitude picked out by the word 'hope' in 'I hope that murder is wrong' the same as that picked out in 'I hope it is raining'? Contemporary quasirealists already hold that normative beliefs are a different kind of attitude from nonnormative beliefs. Will we also need normative versus nonnormative versions of every attitude marked out by folk psychology? This has been identified as the many attitudes problem (M. Schroeder 2008 and 2013; Shiller 2017; and Beddor 2020).

I will argue that, despite appearances, the problem is actually not particularly difficult; or at least, the difficulty is simply the difficulty of the Frege-Geach problem, the problem of logical embedding. If the expressivist can solve *that,* solving the problem of attitudinal embeddings is easy. The trick is to keep in mind that the expressivist does not owe much by way of an account of the attitudes themselves, only of the normative content. She only needs to tell us what the clause 'murder is wrong' contributes to 'Bob hopes that murder is wrong'. She does not owe us an account of what hope is, beyond a general commitment to some form of functionalism about attitudes in general.

My proposal is that for any propositional attitude other than belief, the attitude takes p as its content in virtue of standing in the right functional relations to the belief that p. Sebastian Köhler (2017) and Bob Beddor (2020) have already proposed solutions of this type. However, their solutions come with additional theoretical commitments we need not accept. Köhler's solution, to be fully realised, depends on explaining what it is to believe, desire, hope, or fear that p without making any appeal to propositional content at all in the explanation. This is an ambitious project, requiring a purely functional, content-free way of individuating each attitude type. In short, for Köhler, the expressivist does owe us an account of what hope is. Beddor proposes that the expressivist offer an account of what it is to desire that p when p is normative, and then to reduce all remaining attitudes to combinations of beliefs and desires. This is quite obviously to give an account of what hope is. My proposal will be simpler, requiring only minimal commitments, commitments implicit in much of expressivism already.

To anticipate one source of potential dissatisfaction, the solution may seem overly simple. The basic idea is that to hope that murder is wrong is to be in a mental state that interacts with the belief that murder is wrong in the same way that hope that it is raining interacts with the belief that it is raining. Some extra precision will be added, and some extra sophistication. What I will try to establish in this paper is that this very simple and minimal answer is sufficient for addressing the specific problem created by expressivism. What remains of the problem is not a specifically expressivist problem, but a general problem in philosophy of mind (and empirical psychology) about the nature of intentional attitudes. As evidence of this analysis, I will show that both Beddor and Köhler's proposals are more specific versions of the solution I will offer, and that they are only able to offer more detail in virtue of taking on auxiliary hypotheses in philosophy of mind which an expressivist could consistently deny. In short, much of the apparent difficulty surrounding the many attitudes problem is simply that people have expected the expressivist to solve problems that have nothing to do with expressivism. The only commitment the expressivist needs is, as noted, a general commitment to functionalism—and most expressivists were already implicitly committed to that.

I will start with an analysis of what the problem really is, with the aim of establishing that the simple answer will suffice.

1. The Problem

Beddor begins his discussion of the many attitudes problem with a statement of its importance to the broader expressivist project:

> …[N]oncognitivists aspire to give an account of all normative
> thought. Given this lofty ambition, the standard practice of focusing
> on one species of normative thought—the doxastic attitudes—seems

parochial. And it is worrisome, because if noncognitivists cannot give

a plausible account of the nondoxastic attitudes, then their entire

project is doomed.

<div align="right">(2020: 2791)</div>

Beddor is correct that 'noncognitivists aspire to give an account of all normative thought'. But there are two ways to interpret the explanatory demand the aspiration imposes. Does the expressivist need to explain what it is *to think* about normative stuff, or does she need to explain what it is to think *about normative stuff*? To put this more clearly: does she need to provide an account of what attitudes such as hope, wonder, and gladness are, or does she need an account of what it is for these attitudes to have a normative content? It may be possible to answer one without the other. Based on his proposal, Beddor seems to think the expressivist must do both, and that she will solve the problem of what it is for a thought to have normative content by providing a reductive account of what these thoughts are.

Beddor begins by providing a reductive account of desire in terms of dispositions. A desire is made up of a disposition to act in ways that, according to one's beliefs, make it more likely that content of the desire will be realised, and also a disposition to experience pleasure when one believes the desire's content obtains, displeasure when one believes it does not. Beddor then reduces all further propositional attitudes to combinations of beliefs and desires. Along with the standing expressivist commitment to reduce normative beliefs to desire-like states with nonnormative contents, this represents a commitment to giving a reductive explanation of the nature of all normative thought. The explanation is an explanation of what it is to think—a reductive account of what it is to desire, hope, fear, intend, and even to believe when the belief in question is normative. By way of this reduction of normative attitudes to dispositional clusters, we can also say what it is to have thoughts with normative content without making reference to anything normative in the explanation.

Beddor's proposal may be correct. But it rests on an extremely contentious thesis in philosophy of mind, that all other attitudes reduce to beliefs and desires. Imagine that we conclude that Bratman (1999) and Holton (2009) are right: intention must be understood as a *sui generis* attitude. It would be very surprising if it followed from this that expressivism were unworkable. But then the solution to the many attitudes problem should not depend on Bratman and Holton being wrong.

More generally, Beddor's solution does not improve the expressivist's dialectical position. It establishes that the expressivist can solve the many attitudes problem—but only at the price of taking on contentious theoretical commitments on a whole range of debates where we would have thought expressivism could be neutral.

The psychologist Lisa Feldman Barrett has, for example, challenged the idea that *any* of our emotions are natural kinds, holding instead that our emotion terms ('angry', 'sad', 'frightened') pick out a disjunction of different psychological responses, held together family resemblances that are culturally specific (Barrett 2006, 2012, and 2018; and Gendron et al. 2014). Even if this is wrong about some emotions, it may well be true of others. It would be unsurprising if 'amused that', 'annoyed that', and 'elated that' failed to pick out natural kinds. In any case, right or wrong, Barrett's thesis seems orthogonal to the truth of expressivism. But if it is true, being in an emotional state cannot be given necessary and jointly sufficient conditions, as Beddor's reductive thesis would have it.

Similarly, Beddor proposes a unified account of the nature of desire, in terms of the dispositional states to which it reduces (2794ff.). This may be correct, but it is one very specific analysis. Timothy Schroeder has argued that desires should not be understood dispositionally; rather, 'desire' is a natural kind term for a representation in the reward system of the brain (2004). He contrasts this with another kind of view which treats 'desire' as a generic term for picking out a variety of different motivating states, which may differ in many of their psychological properties, despite having a few key

5

ones in common (4-5). Again, it would be surprising to find out that the truth of expressivism depended on this debate. Beddor's proposal wins one battle for the expressivist, but only by dragging her into many more.

What the expressivist should want is a solution in more general terms, so that it does not depend on highly specific and controversial theses in philosophy of mind. This means that expressivists should not try to give reductive accounts of what hope, fear, gladness, and wonder are. It would be no objection to a metaethical realist that they lack a reductive story of what it is to hope that murder is wrong, even though the lack of such a story would mean that the realist has in some sense failed to account for all of normative thought.

As Beddor notes, the realist has an answer unavailable to expressivism. The realist can simply say: 'Give me your favourite story about what's involved in hoping for a promotion, or world peace, or what have you. On my view, to hope that lying is permissible is to be in the same sort of psychological state, just with a normative proposition as its content' (2020: 2790). Expressivists cannot appeal to normative content in their explanation of what it is to hope that murder is wrong. Quasirealists do allow for talk of normative propositions, but as we will see, talk of them is a device of communication and generalisation. Such propositions are not supposed to bear explanatory weight.

It does not follow, however, that the expressivist must give a reductive account of what hope is. What she owes us is an explanation, in nonnormative terms, of what it is for an attitude to have normative content. She needs to explain what it is to think *about normative stuff,* not what it is *to think*. I will spell out here how the expressivist can do this, while making only minimal commitments about the nature of the attitudes themselves.

2. Believing Murder Is Wrong

How do expressivists handle the attribution 'Bob believes that murder is wrong'?

It is now commonplace that expressivists will talk of normative truths, properties, and beliefs, so long as these are understood in a *deflationary* sense. A deflationary account tells us that even though some term is grammatically a noun or a predicate, it should not be understood in terms of its reference, whether an object or property or intension. Rather, the meaning of the term is given by some other use that it serves (Köhler 2017: 194; and 2018: 342 and fn. 11). For example, despite the fact that it is a predicate, 'true' does not refer to a property of *being true,* at least not according to the deflationist. Rather, the meaning of the predicate is exhausted by the T-schema:

'S' is true if and only if S.

The predicate is used for disquotation, indirect assertion, and quantification.[2] For example, with the predicate, I can say 'Everything Kant says is true', and thereby commit myself to everything Kant said without repeating it all, or without even knowing everything that Kant said.

A similar rule provides the meaning of 'proposition':

'S' expresses the proposition that S.

Such a term allows us a convenient way of making generalisations, claims about translatability, and how our attitudes relate to what's linguistically communicated.[3]

---

[2] For a classic statement, see (Horwich 1998).

[3] For an excellent account of how *that*-clauses can serve these functions, see (Köhler 2017 and 2018). While Köhler is working within the framework of metasemantic expressivism, much of what he says about propositions should be available to semantic varieties of expressivism as well.

On the other hand, Billy Dunaway (2010) identifies potential difficulties with deflationary propositions. See (Köhler 2017) for a reply.

Attitudinal embeddings are one aspect of that last use, of how propositions let us relate our mental states to what's linguistically communicated. To see how they do this, consider a similarly minimalist account of belief. To believe that p is simply to be in a mental state that would be expressed by sincerely asserting the sentence 'p'. Furthermore, whether a given sentence is assertoric or not is simply a matter of its grammatical and logical form. 'Murder is wrong' can thus be asserted, and its assertion expresses a belief —that murder is wrong (Timmons 1999: 135ff.; and Dreier 2004).

So we have a story about how propositional-clauses serve to link one attitude, belief, to what's asserted. The problem is extending this story to include other propositional attitudes. We can say what it is for a belief to be related to the proposition p: it is for the assertoric sentence 'p' to express the belief. But what is it for desire or hope or fear to be related to p?

3. The General Form of Solutions to the Frege-Geach Problem

The second step is to look at what most solutions to the Frege-Geach Problem have in common. The problem, remember, is that of offering a systematic and constructive account of the meaning of arbitrarily complex sentences, that predicts the right kind of semantic relations between these sentences.[4] The conditional sentence 'If murder is wrong then encouraging murder is wrong' needs to have the right kind of logical relations to 'Murder is wrong', for example.

What most solutions propose is to explain the semantic (including logical) relations sentences stand in to each other in terms of the functional relations between the beliefs expressed (Blackburn 1984; Gibbard 1990; Björnsson 2001; Sinclair 2011; and

---

[4] See (Woods 2017) for overview of the problem.

Baker & Woods 2015).[5] Let's say the belief that murder is wrong is type-identical with disapproval of murder. Then the belief that if murder is wrong then encouraging others to murder is wrong will be type-identical to some state bearing appropriate functional or rational relations to disapproval of murder. It might be, for example, a higher-order attitude: the belief in the conditional might be type identical to some negative attitude toward disapproving of murder while not disapproving of encouraging murder (Blackburn 1984; and Sinclair 2012). Alternately, it might be identified as a (possibly complex) state with the functional role of producing a state of disapproval of encouraging murder when the agent already disapproves of murder, and which prevents the disapproval of murder when the agent already approves (or tolerates) encouraging others to murder (Gibbard 2003; and Baker & Woods 2015).[6] In any case, a sentence of the form 'if p then q' expresses a mental state which bears appropriate functional relations to the mental states expressed by 'p' and 'q', for any 'p' and 'q'.

4. The Solution

Expressivists can tell a similar story about embeddings under attitudinal terms. How does the meaning 'Murder is wrong' relate to the sentence 'Bob is glad that murder is wrong'? The first expresses the belief that murder is wrong. The second ascribes to Bob a state functionally related in the proper way to the belief that murder is wrong.

The rough idea is this, to be glad that murder is wrong is to be in a psychological state that depends on the belief that murder is wrong, and goes away when this belief goes away. Similarly for being sad or angry that murder is wrong. And because we want our account to be uniform, we must say the same about being glad, sad, or angry that

---

[5] An exception to this is Mark Schroeder, who tries to account for the semantic relations between sentences by appeal to further semantic relations between the contents of the relevant attitudes (2008 and 2013).

[6] Note that Gibbard would analyse *disapproval* of X-ing in terms of a plan to feel angry or guilty about X-ing.

grass is green. Attitudinal relations to propositions, other than the belief-relation, are mediated by a belief in the same proposition.

A bit more rigorously:

> For any propositional attitude other than belief, A, the content of A is p if and only if A stands in the right kind of functional relations to the belief that p.

A does occur on both sides of the proposal, but remember the thesis is not intended to be reductive of A. Giving reductive accounts of attitudes is a project in psychology, cognitive science, and philosophy of mind. It is, as noted before, outside the purview of metaethics or semantics. The target of analysis is not A, but having p as a content.

Admittedly, p also appears on the right side of the biconditional. But the biconditional is not intended to stand alone. The assumption is that the expressivist already has an account of what it is to believe that p when p is normative, that can be given in terms in which p no longer appears. On the other hand, if p is not normative, the expressivist does not necessarily want to analyse it away. (The statement above is completely trivial in the nonnormative case, but the expressivist never promised to say anything interesting about nonnormative cases.) To finalise the proposal, the expressivist need only plug in her preferred account of normative belief.

Unfortunately, for some attitudes, standing in the right kind of functional relation to the belief that p is not sufficient to determine that the agent holds that attitude toward p. It may be that to desire that p one must be disposed to feel happy that p and frustrated that $\sim p$. But note, first of all, that we are now individuating the desire that p only partly on the basis of its relation to the belief that p. The belief that $\sim p$ is also part of how we individuate it. Nor are these the only important functional relations for determining the

desire's content. The desire that *p* must also stand in functional relations to the relevant means-end beliefs, to beliefs that *x*-ing is a necessary means to *p*.

Of course, one might argue, as Köhler does, that desires with normative contents do not motivate action. So they cannot stand in functional relations to means-end beliefs (2017: 206). But the proposal here is neutral on whether 'desire' picks out a unified class of mental states, or a disjunction of different kinds of mental states, held together by family resemblance or certain practical or theoretical interests. Köhler may be right that some desires do not have the functional role of motivating action. But some clearly do, and these desires must be partially individuated by their relation to the belief that *x*-ing is a necessary means to *p*, and not simply by their relation to the belief that *p*. So we need to tell a story about these.

The belief that *x*-ing is a necessary means to *p* stands in some semantic relation to *p*, as does the belief that ~*p*, just as the belief that if *p* then *q* does. In the case of the last belief, we explained this semantic relation in terms of the functional relations between the conditional belief and the belief in the antecedent. While the functional relation between the means-end belief and the belief that *p* is presumably more complicated, we can say the belief that *x*-ing is a necessary means to *p* is a state that has the functional role of preventing the agent from believing that *p* so long as she does not believe that she has or will *x*.

This suggests the following. The content of *A* is *p* just in case *A* is appropriately functionally related to the belief that *p*, or to other attitudes that are appropriately functionally related to the belief that *p*, or to attitudes functional related to attitudes functionally related to…

So the improved solution would go like this:

For any propositional attitude other than belief, *A*, the content of *A* is *p* if and only if *A* stands in the right kind of functional relations to other attitudes which are semantically related to *p*.

And we can combine this with the following inductive definition:

A is semantically related to *p* if and only if

(i) *A* is the belief that *p*; or

(ii) *A* is functionally related in the right kind of way to some attitude, *B*, which is semantically related to *p*.

With this in place, we can accommodate other potentially troubling attitudes. Hoping or fearing that murder is wrong are likely functionally incompatible with the belief that murder is *not* wrong. To hope murder is wrong presumably involves a disposition to become happy when presented with evidence for the belief that murder is wrong, and to become unhappy when presented with evidence that it is not. That is, many attitudes are constituted by functional relations, not to a belief that shares their content, but to other mental states functionally related to such a belief.

Unfortunately this still does not account for all the ways an attitude may be semantically related to *p*. Köhler, for example, calls attention to the attitude of entertaining the thought that *p* (2017: 206). He proposes that 'to entertain a thought is to simulate the corresponding belief, where "simulation" is understood as running the belief "off-line"' (*ibid.*). But a simulation of the belief that *p* is not obviously functionally related to the belief that *p,* at least in the sense of 'functionally related' at work in (ii) above. In that context, A is functionally related to B when A and B are disposed to affect or interact with each other causally.[7] Köhler does not say much about what is involved in

_____

[7] Thanks to a referee for pointing out this problem.

simulation, but an obvious thought is that we simulate a belief by imagining that we have it. This *can* causally interact with the relevant belief—imagining what it would be like to believe that p may play a role in suppositional reasoning, and so make it more or less likely that I come to believe that p. But it is questionable whether any such relations to belief are essential to being a state of imagining.

This is actually an instance of a wider problem. Higher-order attitudes and metalinguistic attitudes are all semantically related to various contents. If I believe that Bob believes that it is raining, my belief is semantically related to <Bob believes that it is raining>, but also to <it is raining>. It is unclear, though, that there are any functional relations between my belief about Bob's belief and my belief that it is raining. More importantly, even if there are, this seems like an unnecessarily roundabout way of explaining why my belief is semantically related to <it is raining>. Here is a more natural alternative: my higher-order belief is semantically related to the relevant proposition because it is about a functional state which is itself semantically related to the relevant proposition.

This also helps address metalinguistic attitudes. Let's say I believe 'Mord ist schlecht' means that murder is wrong. Again, this belief is semantically related to <murder is wrong>, but it is unclear that it bears any essential functional relations to the belief that murder is wrong. Again, the obvious explanation is that belief either itself represents the sentence as a conventional expression of a functional state semantically related to <murder is wrong>, or else is functionally related to a tacit understanding that the sentence conventionally expresses such a functional state. (People do not need to explicitly believe expressivism in order to have metalinguistic attitudes; however, expressivism is a theory of our natural languages, and so, if it is true, mastering a language must involve at least tacit understanding of how sentences are used to express attitudes.)

This suggests the following modification to the inductive definition:

A is semantically related to p if and only if

      (i) A is the belief that p; or

      (ii) A is functionally related in the right kind of way to some

      attitude, B, which is semantically related to p; or

      (iii) A is about an attitude, B, which is semantically related to p.

With this addition we can handle entertaining the thought that p. To entertain the thought that p is to imagine what it would be like to believe that p; it is to imagine what it would be like to be in a certain functional state.

4.1. Commitments of the View

Before moving on to the second objection, let us review what the solution is meant to accomplish, and what it is not. As Beddor noted, the realist can give the following answer to what it is to hope that murder is wrong. Take your favourite story about what it is to hope that p when p is nonnormative. To hope that murder is wrong is to stand in that same kind of psychological state, only to the content *murder is wrong*. My position is that the expressivist should try as far as possible to emulate that kind of answer. So, take your favoured story about how hope that p is functionally related, both directly and via the mediation of other attitudes, to the belief that p, when p is nonnormative. To hope that murder is wrong is to be in a state with the same kind of functional relations to the belief that murder is wrong. This is compatible with the hope reducing to a complex of belief and desires which are in turn individuated functionally. It is also compatible with 'hope' being a label for a disjunction of different kinds of functional states held together by a family resemblance. It is neutral, as far as possible, on what it is *to think*.

14

This is not the same as being completely neutral, however. The solution does assume that some form of functionalism about propositional attitudes is correct. Attitudes such as hope and fear are defined in terms of their functional relations; attitudes have their propositional content in virtue of their functional relations. So the solution is incompatible with *strong phenomenal intentionality,* the view that the intentional content of a mental state is fully determined by its phenomenology (see [Bourget & Mendolovici 2019] for overview).

Is the solution, then, vulnerable to the objection I raised against Beddor and Köhler, that it forces expressivists to take a stand where we would have thought she could remain neutral? First, both Beddor and Köhler's solutions explicitly presuppose functionalism as well (Beddor 2020: 2791ff.; and Köhler 2017: 199ff.). In short, all extant solutions presuppose functionalism. The proposal on offer has the advantage, then, of weaker commitments than its rivals. Even if it drags the expressivist into battles she would rather have avoided, it drags her into fewer. It should also be kept in mind that many expressivists had no intention of remaining neutral on this particular point. Functionalism is a presupposition of the dominant strategy for solving the Frege-Geach problem, considered in section 3. So, for the majority of expressivists, functionalism is already a commitment of the theory. In summary, the solution brings no extra theoretical commitments for most expressivists; and for the remainder, it comes with the weakest commitments currently on offer.

A final clarification should be made. To say that the solution presupposes functionalism and is incompatible with strong phenomenal intentionality is not to say that it must deny that mental states have a phenomenology. Beddor and Köhler both define desires partially in terms of phenomenology—roughly the way that believing or imagining the desire's satisfaction causes pleasure (Beddor 2020: 2796; Köhler 2017: 206). The disposition to produce pleasure, pain, or some other felt experience in interaction with other attitudes may well be part of the functional role of a given mental state. What

is ruled out are views on which the content of an attitude is solely a matter of how it feels, on which functional roles are not necessary for determining the content of an attitude.[8]

5. Objections

5.1. Doesn't the Expressivist Owe Us Something More?

A series of related objections stem from disbelief that the problem can be addressed this easily.

As was noted earlier, the metaethical realist does not owe us much of an explanation of what it is to hope that murder is wrong. This is because the question of what hope is, or what propositions are, are questions for philosophy of mind and language in general. There is no special metaethical problem here. So why would the expressivist owe us more?

Due to her commitment to an unorthodox theory of meaning the expressivist needs to show that she can offer a systematic story about the contribution a normative clause makes to an attitudinal ascription in a way that is clearly compatible with her expressivist commitments (M. Schroeder 2013). As Beddor notes, this means an explanation that makes no reference to normative propositions in its *explanans*. But this is all that is special to the expressivist: she needs an explanation of what it is for an attitude to have normative content which does not invoke normative content in the *explanans*. The proposal here provides a general recipe for how to do that. So far no reason why the expressivist needs to answer more general problems in philosophy of mind, such as what it is to hope, has been given.

---

[8] For an example of such a view, see (Mendolovici 2018).

It's worth repeating in this context that Beddor's proposal is a special cases of the more general solution proposed here. Beddor reduces all attitudes to various combinations of beliefs and desires, with desires analysed in terms of dispositional properties to produce certain events in the presence of beliefs with the right kind of content. Beliefs with normative content are then given the expressivist treatment. This is clearly one way of specifying the functional roles relating each attitude type to beliefs with the relevant content. But, again, suppose Bratman and Holton are correct: intention is a *sui generis* attitude, irreducible to beliefs and desires. This is still compatible with individuating the content of a given intention on the basis of its functional relations to beliefs. The intention that one does something wrong, for example, could be a state with the function (i) of being relatively stable, (ii) of causing, in conjunction with the belief that *X*-ing is a necessary means to doing something wrong, an intention to perform action *X*, and (iii) of going out of existence when one believes that doing something wrong is not possible.[9]

In short, the most obvious way of providing more detailed answers to the many attitudes problem is to simply combine the general proposal given in this paper with auxiliary hypotheses *that have nothing to do with expressivism*. But since there are many possible auxiliary hypotheses, it is unclear why the expressivist should need to pick one, rather than simply noting that the details depend on debates in philosophy of mind, and that she has an answer compatible with a number of ways of such debates might turn out.

Köhler's proposal illustrates the same point. Köhler starts with a bifurcation in the nature of desire. Desires with normative contents have a different functional role from those with nonnormative contents. The latter have the motivate action—

---

[9] For discussion of to what extent intentions must be consistent with beliefs that satisfying the intention is possibly, see (Holton 2009: 40ff.). Note again that the expressivist could give many different answers to the functional role of intention, even after agreeing that intention is a *sui generis* attitude.

specifically they have the functional role of motivating us to pursue those features of the environment that nonnormative beliefs have the functional role of tracking (205). Normative desires, on the other hand, dispose us 'to take pleasure in entertaining the thought that p' (206). As noted above, entertaining the thought that *p* is then analysed in terms of simulating the belief that *p*. Thus both kinds of desires are individuated by their relation to the relevant beliefs—mediated by the state of entertaining. But again, as the case of Beddor shows, accepting a bifurcated account of desires is orthogonal to expressivism. This is an independent thesis in philosophy of mind.

Perhaps the complaint is that this solution amounts to a mere list of the properties a state of hoping, desiring, or fearing that *p* for some normative *p* needs to have.[10] This is reminiscent of Schroeder's objection to Gibbard's expressivism: it provides a mere list of what the attitudes expressed by our sentences must be like, and so is not explanatory (2008: 52). The problem is that this objection is itself unclear. As Baker and Woods note in response, providing a recipe that specifies the functional role of the attitude expressed by any arbitrary sentence seems constructive, at least until we see an explicit argument for stricter conditions on constructiveness (2015: 411-12, fn. 69). Likewise, those who think the problem of attitudinal embeddings must be hard need to argue that the solution must go beyond a description of the functional relations that ground having some normative *p* as a content. To repeat, a metaethical realist is allowed to give a placeholder solution: 'To hope that murder is wrong is to have the attitude of hope—whatever that turns out to be—to the proposition that murder is wrong— whatever that is'. In light of this, it is hard to make out why an expressivist is forbidden from giving answers that similarly leave the details to other parts of philosophy.

The worry may be that even if we can specify the functional role such an attitude would have to play, we have no reason to think that there is any state that plays that functional role. But the reason to think that there are states playing the relevant roles is

---

[10] Thanks to a referee of an earlier draft for raising this problem.

the intuitive data that motivated the many attitudes problem in the first place. We say that people hope that murder is wrong or worry that eating meat is wrong. Given the general acceptance of folk psychology, we presumably have some defeasible justification for assuming these commonsensical attitude ascriptions are sometimes true. But then that is reason to think that there is some state fulfilling the relevant functional role.

This may seem to get the cart before the horse. Doesn't it assume that the expressivist analysis of the functional roles of the relevant attitudes is correct? It does, but recall the dialectic. I am arguing that insofar as the many attitudes problem is a distinct problem, is not a very substantial one. *If* the expressivist can solve the Frege-Geach problem, *then* the many attitudes problem is easily solved. Assume she can solve the Frege-Geach problem. Then take your preferred story about how hope that murder is wrong is functionally related (directly or indirectly) to the attitude expressed by 'Murder is wrong'. She will tell the exact same story. Your objection can't be to the functional story she just told about the nature of hope—it's your own story. Is your objection then to her analysis of the attitude expressed by 'Murder is wrong'? Okay. But in that case the objection is that the many attitudes problem cannot be solved—*if* expressivism is false. This is true, but now the many attitudes problem no longer serves as a distinct objection to the theory.

5.2. Massive Coincidence

A more substantial problem is presented by Derek Shiller (2017). Why do our normative beliefs have the same functional relation to hope, fear, anger and so on as our nonnormative beliefs? The objection is that for the expressivist, 'belief' names a disjunction of mental states. Normative beliefs are not the same kind of attitude as nonnormative beliefs—the former are really a kind of desire-like state. But why would the attitude-type hope interact with a desire-like state in the same way that it interacts

19

with a state that is not desire-like, namely nonnormative belief? The expressivist seems committed to a massive coincidence.

In reply, we should first note that the expressivist can deny that 'belief' names a disjunction of states on her view: belief is a genus, with normative and nonnormative species. What all beliefs have in common is that they are expressed through assertoric speech acts, that they can be true or false, and that they interact with other beliefs to maintain logical coherence. This does not fully answer Shiller's objection on its own, but it does force us to be clearer about what the puzzle is. After all, the expressivist can maintain, at least initially, that there is no massive coincidence. We have a unified attitude-type, belief, and all beliefs interact with token cases of hope or fear or anger in the same kinds of ways. Why does the fact that the category of belief divides into two subcategories introduce a mystery?

There are two obvious concerns. First, the commonalities which hold this genus together seem relatively superficial. That both species are expressible through assertions may explain, for example, why we categorise them together in everyday language. But normative beliefs are, according to the expressivist, desire-like, and so they seem more naturally classed with desires and intentions than with nonnormative beliefs. The second reason for concern is that we normally understand fear, anger, hope, and so on as responsive in various ways to our picture of what the world is like. But according to the expressivist, normative beliefs are not part of our picture of what the world is like (save for some deflationary sense, which is more a matter of linguistic convenience than anything else).

The answer to these concerns is that if normative attitudes share in these seemingly superficial commonalities then a person's normative perspective will share important features in common with their picture of what the world is like. We can see this by first considering motivating attitudes which we will stipulate are not assertorically expressible, then stipulate that another class of functionally similar attitudes is

assertorically expressible. We will see that to be assertorically expressible, the attitudes will need other, "deeper" properties in common with nonnormative beliefs.

Start with intention, stipulating that it is not expressed assertorically. The intention to murder Bob and the intention not to murder Bob are functionally incompatible. But notice that if we divide the two incompatible attitudes among two different people, those people need not disagree or regard the other's intention as mistaken. If I intend to murder Bob, I may well regard your intention not to murder him as unproblematic. People have different life projects, after all. One and the same person cannot have both intentions, but if they are split up among multiple people, we can each go our own way.

Imagine instead that there are intention-like states, with the same kind of functional incompatibilities, but now expressible through normative assertions. Imagine that I have the intention-like state toward murdering Bob, and you have the intention-like state toward refraining. Now there must be disagreement between us, at least on the general strategy for solving the Frege-Geach problem discussed in section 3. On that strategy, remember, semantic relations between normative assertions are explained in terms of functional relations between the attitudes those assertions express. The attitudes are functionally incompatible, and so the assertions expressing them must be semantically incompatible. So we must disagree. We cannot both be right. The superficial commonality—being assertorically expressible—predicts a more significant commonality. Beliefs, whether normative or nonnormative, commit us to understanding incompatible beliefs as mistaken.

Note that this claim of a link between assertoric expressibility and disagreement is neutral on the order of explanation. It may be that attitudes are assertorically expressible because it is the nature of these attitudes to place their holders in relations of disagreement. There is something special about the functional role of our intention-like states, some functional difference between them and ordinary intentions, in virtue of

which they commit those who hold them to interpersonal disagreement, and because of this, they are assertorically expressible. But it may be that these attitudes place their holders in relations of disagreement because they are assertorically expressible. It is not the functional roles of the intention-like attitudes themselves that produces the disagreement. Rather, part of learning normative language is being socialised into regarding those with incompatible intention-like attitudes as mistaken.

Given this, there will be rational pressure against adopting a new normative attitude arbitrarily, the way one might change one's mind whether to go to the grocery or bookstore first. If I believe eating meat is fine, I am implicitly committed to thinking that those who believe otherwise are mistaken, and that includes a future version of me who thinks otherwise. I will be strongly disposed, then, to continue believing eating meat is fine unless presented with reasons to think my current perspective is inferior to one on which eating meat is wrong. In general, I will only revise my normative beliefs to the extent that a rival normative perspective seems superior. This means I will typically adopt a new normative belief only because some of my other normative beliefs favour doing so—and again, these normative beliefs will strike me at the time as attitudes it would be a mistake to change.

So my normative reasoning will seem to me much more like a process of discovery than invention. My normative perspective is largely fixed. Changing it without justification will seem to me like deliberately embracing error. When I change my normative perspective, I will understand myself as guided by further normative principles. These further principles will themselves strike me as not to be changed, unless that is justified on the basis of further principles, and so on. Ultimately changes to my normative perspective will feel like it is a matter of answering to principles whose status as correct principles is already in place.

So we can fear, for example, that something might be wrong, or be happy about the same. Given how our normative attitudes function, it is as though these things can

turn out well or poorly for us and for other things we care about. It might turn out that eating meat is wrong, and, since meat is tasty, Bob hopes that it is not and will be sad if it is. Even if normative beliefs are not responses to what the world is like in the way that nonnormative beliefs are, there will be an important similarity from the point of view of the normative reasoner. Whether eating meat is wrong will seem to her to be a matter to be figured out on the basis of what she already regards correct.

To summarise my answer to the massive coincidence objection: first, the expressivist is not obviously positing massive coincidence. She can identify an unified class of attitude—beliefs—with certain essential properties in common, and she simply posits that members of this class can functionally interact in the same way with other attitudes such as fear, anger, hope, and so on. If there is a massive coincidence, it is because what holds beliefs together are relatively superficial properties, and their deeper functional properties are very different. But the superficial commonalities must be accompanied by further properties, and these properties make it intelligible why we would, for example, hope or fear that *p*, for normative *p*'s.

5.3. Can We Really Be Sure about All Attitudes?

Despite its claims to neutrality and minimalism, the account here might seem to make a bold prediction. It predicts that *all* propositional attitudes can be individuated according to the final inductive definition offered. But what if there is some attitude *A*, which cannot be analysed in terms of functional relations, direct or indirect, to belief, and yet it still seems correct to us to say on occasion, 'Bob *A*'s that eating meat is wrong'? Does it follow that expressivism fails?

Probably not. The concern raised by the many attitudes problem is that expressivism is massively revisionary, declaring that it is mistaken or even unintelligible to attribute desires, hopes, fears, and so on with normative content. It is a powerful

objection that a philosophical theory is massively revisionary; it is another thing when such a theory forces us to revise at the margins. This is still a cost, but one to be weighed holistically against the theory's other costs and virtues. If expressivism can accommodate the majority of the supposedly problematic attitude-attributions, yet there remains a handful of attitude-types which it cannot, then it asks us to revise at the margins. It would be a cost, but one to be compared to the costs of metaethical alternatives.

6. Conclusion

Does this result in the systematicity we want? The view says that any utterance of the form '*A*s that *p*' refers to a mental state that stands in the right kind of functional or intentional relations to the mental state expressed by sincerely asserting the sentence 'p'. The relations in question can be direct, or indirect—that is, *A* might be functionally related to an attitude which is about an attitude which… is functionally related to the state expressed by sentence 'p'. The relations are different for different attitudes of course. But we want standing in the hope relation to *p* to be different from standing in the desire relation to *p*: and on this picture that difference amounts to a difference of functional relation to the state expressed by 'p'. Spelling out these relations in detail is a job for philosophy of mind. The only special commitment the expressivist needs to take on is that all attitudes (other than belief) that *p* can be individuated by their functional and intentional relations, direct or indirect, to the belief that *p*.

This is all to say the problem is less serious than it initially appeared—or maybe rather that the only real difficulty is whether expressivists are really entitled to normative beliefs, and once that is in place the rest comes easily. But there is an interesting upshot: so long as some part of our language has enough structure to allow for the standard array of logical embeddings, attitudinal embeddings come cheap.

Lingnan University, Hong Kong SAR, China

Works Cited

Baker, D., and J. Woods (2015) 'How Expressivists Can and Should Explain Inconsistency', *Ethics,* 125/2: 391-424.

Barrett, L. F. (2006) 'Are Emotions Natural Kinds?', *Perspectives on Psychological Science,* 1/1: 28-58.

Barrett, L. F. (2012) 'Emotions Are Real', *Emotion,* 12/3: 413-29.

Barrett, L. F. (2018) *How Emotions Are Made,* First Mariner Books Edition, New York: Mariner Books.

Beddor, B. (2020) 'A Solution to the Many Attitudes Problem', *Philosophical Studies,* 177/9: 2789-813.

Blackburn, S. (1984) *Spreading the Word,* New York: Oxford University Press.

Bourget, D. and A. Mendolovici (2019) 'Phenomenal Intentionality', in E. Zalta (ed.) *Stanford Encyclopedia of Philosophy. <https://plato.stanford.edu/archives/fall2019/entries/ phenomenal-intentionality/>.*

Bratman, M (1999) *Faces of Intention,* New York: Cambridge University Press.

Carr, J. (2015) 'Ecumenical Expressivism Ecumenicized', *Analysis,* 75/3: 442-50.

Chrisman, M (2008) 'Expressivism, Inferentialism, and Saving the Debate', *Philosophy and Phenomenological Research,* 77/2: 334-58.

Chrisman, M. (2015) *The Meaning of 'Ought': Beyond Descriptivism and Expressivism in Metaethics,* New York: Oxford University Press.

Dreier, J. (2004) 'Meta-ethics and the Problem of Creeping Minimalism', *Philosophical Perspectives* 18: 23-44.

Dunaway, B. (2010) 'Minimalist Semantics in Meta-ethical Expressivism', *Philosophical Studies,* 151/3: 351-71.

Gendron, M., D. Roberson, J. M. van der Vyer, and L. F. Barrett (2014) 'Perceptions of Emotion from Facial Expressions Are Not Culturally Universal: Evidence from a Remote Culture', *Emotion,* 14/2: 251-62.

Gibbard, A. (1990) *Wise Choices, Apt Feelings: A Theory of Normative Judgment,* First Harvard Paperback Edition 1992, Cambridge, MA: Harvard University Press.

Holton, R. (2009) *Wanting, Willing, Waiting,* New York: Oxford University Press.

Horwich, P. (1998) *Truth,* New York: Oxford University Press.

Köhler, S. (2017) 'Expressivism, Belief, and All *That', Journal of Philosophy,* 114/4: 189-207.

Köhler, S. (2018) 'Expressivism, Meaning, and All *That', Canadian Journal of Philosophy,* 48/3-4: 337-56.

Mendolovici, A. (2018) *The Phenomenal Basis of Intentionality,* New York: Oxford University Press.

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

Schroeder, M. (2013) 'Two Roles for Propositions: Cause for Divorce?', *Nous,* 47/3: 409-30.

Schroeder, T. (2004) *Three Faces of Desire,* New York: Oxford University Press.

Shiller, D. (2017) 'The Problem of Other Attitudes', *American Philosophical Quarterly,* 54/2: 141-52.

Simpson, M. (2018) 'Solving the Problem of Creeping Minimalism', *Canadian Journal of Philosophy,* 48/3-4: 510-31.

Sinclair, N. (2011) 'Moral Expressivism and Sentential Negation', *Philosophical Studies,* 152/3: 385-411.

Timmons, M. (1999) *Morality without Foundations,* New York: Oxford University Press.

Woods, J. (2017) 'The Frege-Geach Problem', in T. McPherson and D. Plunkett (eds.) *The Routledge Handbook of Metaethics,* 226-42, New York: Routledge.