ECUMENICAL EXPRESSIVISM AND THE FREGE-GEACH PROBLEM[*]

Alexander Miller
University of Otago, New Zealand

Kirk Surgener
University of Warwick, UK

1. Cognitivism and Expressivism

A traditional way of drawing the distinction between cognitivist and expressivist accounts of moral judgement characterizes cognitivists as holding that moral judgements express beliefs (and not desire-like attitudes) and expressivists as holding that moral judgements express desire-like attitudes (and not beliefs). Alternatively, the two positions could be summarised as follows:

Cognitivism: For any moral sentence M, M is conventionally used to express a belief (and not a desire-like attitude).

Expressivism: For any moral sentence M, M is conventionally used to express a desire-like attitude (and not a belief).

A central problem for expressivism, thus characterised, is the Frege-Geach Problem, the problem of accounting for the meanings of moral sentences as they appear in unasserted contexts such as the antecedents of conditionals (Geach 1960, 1965).[1] Although the problem has been tackled by leading proponents of expressivism such as Allan Gibbard (1990, 2003) and Simon Blackburn (1984, 1993, 1998) it is fair to say that the solutions offered have not been convincing (see Schroeder 2008a and Miller 2013, chapters 4 and 5). In a series of articles (2006, 2007 2008, 2009) and recent monograph (2014), Michael Ridge has developed a novel form of expressivism, Ecumenical Expressivism, according to which moral judgements express *both* beliefs and desire-like attitudes, and argued that Ecumenical Expressivism enables a relatively straightforward solution to the Frege-

---

[1] As Schroeder (2008a) rightly points out, the problem is much more general than simply dealing with the case of conditionals and concerns a family of issues surrounding compositionality in general: so the problem concerns the expressivist's capacity to preserve moral reasoning in general and not just e.g. moral modus ponens.

Geach Problem.[2] Our main aim in this paper is to challenge Ridge's claim that Ecumenical Expressivism solves the Frege-Geach Problem. We proceed as follows. In §2 we give a very brief reminder of the Frege-Geach Problem. For illustrative purposes that we shall drawn on later, we also recap the 1984 solution to the problem developed by Simon Blackburn and the main reason that Blackburn's solution fails. Following this, in §3 we explain Ridge's distinction between Ecumenical Cognitivism and Ecumenical Expressivism. In §4 we briefly outline how Ridge's Ecumenical Expressivism claims to solve the Frege-Geach Problem, before outlining, in the next four sections, a series of challenges to that solution. We set out our main conclusion and draw some broader morals in §9.[3]

2. Blackburn's Quasi-Realist Expressivism and the Frege-Geach Problem

The fundamental expressivist ideas are that we give an account of the meaning of a sentence in terms of the state of mind that it expresses and that in the case of a moral sentence such as "Murder is wrong" the relevant state of mind is a non-cognitive attitude of disapproval of murder: B!(murder).[4] These ideas, however, leave the expressivist with a problem. While it is plausible to think of the meaning of "Murder is wrong" as it appears in an asserted context such as e.g.

(1) Murder is wrong

in terms of B!(murder), it is difficult to see how this account can be extended to cover the appearance of "murder is wrong" as it appears in an unasserted context such as the antecedent of (2):

---

[2] "Non-Ecumenical Expressivism" is thus the view that moral judgements express desire-like attitudes but not beliefs.

[3] For the most part, for the purposes of evaluating Ridge's solution to the Frege-Geach Problem we focus on the simpler forms of Ecumenical Expressivism broached in his 2006 and 2008: the solution to the Frege-Geach Problem offered in Ridge 2014 is essentially the same as that offered in the earlier articles, with the additional complexities about normative perspectives, "admissible ultimate standards of practical reasoning" and "negative thinking" introduced in the 2014 Ecumenical Expressivist account playing (as far as we can see) no essential role in the attempt to defuse the Frege-Geach Problem. Likewise, we do not concern ourselves with the question as to whether expressivism is best framed as a thesis in semantics or (as Ridge now prefers) in metasemantics. As Ridge himself notes (2014: 137-38) the philosophical work that the expressivist has to carry out to deal with the Frege-Geach problem is effectively the same irrespective of whether it is couched as a view in first-order semantics or as a view in metasemantics.

[4] Ridge characterises expressivism as a form of "ideationalism", where "Ideationalism maintains that facts about the semantic contents of meaningful items in a natural language are constituted by facts about how those items are conventionally used to express states of mind" (2014: 107). For an account of the philosophical motivations for expressivism – in metaphysics, epistemology and moral psychology – see chapters 3 – 5 in Miller (2013).

(2) If murder is wrong then getting Peter to murder people is wrong,

since someone sincerely asserting (2) needn't have an attitude of disapproval towards murder (or indeed towards getting Peter to murder people) – think of how those who approve of helping the aged can still sincerely utter "If helping the aged is wrong then getting Peter to help the aged is wrong". If this extension turns out not to be possible it looks like the inference from (1) and (2) to

(3) Getting Peter to murder people is wrong

will be vitiated by a fallacy of equivocation, since "Murder is wrong" will have different meanings as it appears in (1) and in the antecedent of (2). And this is highly problematic, as the inference is an instance of Modus Ponens, a valid inference form.[5] This is the Frege-Geach Problem, and the challenge to the expressivist is therefore to give an account of the contribution made by the meaning of a moral sentence to the meaning of a more complex sentence in which it appears in terms of the state of mind it expresses when used in an asserted context, in such a way that intuitively valid inferences involving it are not impugned (by, for instance, the commission of fallacies of equivocation).

It will be useful later to contrast Ridge's attempted solution with that attempted by Blackburn in his 1984. To cut to the chase, Blackburn proposes to understand the meaning of a conditional such as (2) above in terms of a higher-order attitude of approval towards moral sensibilities that combine disapproval of murder with disapproval of getting Peter to murder people: schematically, H! [B! (Murder); B! (Getting Peter to murder people)]. If we now think of the overall state of mind of someone who accepts (1) and (2) but rejects (3) we can see that this will consist of disapproval of murder together with approval of combining disapproval of murder with disapproval of getting Peter to murder people, but will lack disapproval of getting Peter to murder people. Someone with this state of mind will be prey to a kind of incoherence: he "has a fractured sensibility which cannot itself be an object of approval" (1984: 195), and this allows us to capture the idea that the inference from (1) and (2) to (3) is valid.

This attempt at solving the Frege-Geach Problem was criticised shortly after its publication by Crispin Wright:

Anything worth calling the validity of an inference has to reside in the inconsistency of accepting its premises but denying its conclusion. Blackburn does indeed speak of the 'clash of attitudes' involved in endorsing the premises of the modus ponens example, construed as he construes it, but in failing to endorse the conclusion. But nothing worth regarding as inconsistency seems to be involved. Those who do that

---

[5] Notice that it will not do for the expressivist to simply accept that this aspect of moral discourse is in bad faith: as we noted above the problem in this area extends to most of moral reasoning. Going down this road would leave the expressivist with an account of the meaning of positive, atomic, moral statements but not much else. At this point it is unclear why developing expressivism is preferable to simply adopting an error theory.

merely fail to have every combination of attitudes of which they themselves approve. That is a *moral* failing, not a logical one (Wright 1988: 25).[6]

Blackburn's 1984 solution thus fails to capture the *logical* validity of the inference from (1) and (2) to (3). However, the key thing to note is that although it fails for the reason set out by Wright, it is nonetheless a genuine attempt to speak to the Frege-Geach worry about equivocation, since the contribution of "Murder is wrong" to the meaning of the conditional (2) is given in terms of the very same state of mind – B! (murder) - that gives its meaning in (1). This is a point we'll return to later.

3. Ecumenical Views

According to ecumenical views of moral judgement, moral judgements can be regarded as expressing *both* beliefs and desire-like attitudes: a moral sentence M is conventionally used to express *both* a belief and a desire-like attitude. This does not, however, lead to a collapse of the distinction between cognitivism and expressivism. According to Ridge, a version of this distinction survives the move towards ecumenicism:

Ecumenical cognitivism allows that moral utterances express both beliefs and desires and insists that the utterances are true if and only if one of the beliefs expressed is true. Ecumenical expressivism also allows that moral utterances express both beliefs and desires but denies that a moral utterance is *guaranteed* to be true just in case the belief(s) it expresses is (are) true (2006: 307-8, emphasis added).

And again:

So long as the belief expressed by a moral utterance *is not semantically guaranteed* to provide the truth-conditions for the utterance, the fact that the belief expressed *contingently* provides the truth-conditions for the token utterance is consistent with expressivism as characterized here (2006: 311-312, emphases added).[7]

The distinction between cognitivism and expressivism within the ecumenical framework is thus recast as follows:

Ecumenical Cognitivism: a moral judgement M expresses both a belief and a desire-like attitude, and, as a matter of semantic and conceptual necessity, M is true iff the belief expressed is true.

Ecumenical Expressivism: a moral judgement M expresses both a belief and a desire-like attitude, but it is not semantically or conceptually necessary that M is true iff the belief expressed is true.

The Ecumenical Cognitivist assigns a certain logical priority to belief: which of an agent's judgements count as moral will be determined by the type of belief with which

---

[6] See also Hale (1986) and Hale (1993). For a useful extension of the sort of objection developed by Wright and Hale, see Van Roojen (1996).

[7] See also (2008: 54, 55, 59) for further use of "guarantee", "semantic guarantee" and so on.

moral judgements necessarily co-vary: for example, a version of Ecumenical Cognitivism which took the beliefs in question to be beliefs about maximising utility would imply that the agent's moral judgements are those about the maximisation of utility. In contrast, although the Ecumenical Expressivist would regard moral judgements as expressing beliefs as well as desire-like attitudes, on this type of account logical priority would be assigned to the desire-like attitudes rather than the beliefs. For example, on the toy ("Plain Vanilla") version of Ecumenical Expressivism that Ridge sometimes uses in explaining the position:

Normative utterances express (a) a speaker's approval [disapproval] of actions in general insofar as they have a certain property, and (b) a belief which makes anaphoric reference to that property (the one in virtue of which the speaker approves [disapproves] of actions in general) (2008: 55).

Consider a utilitarian speaker ("Jeremy"). Jeremy's judgement that X is right expresses (a) an attitude of approval towards actions insofar as they maximise utility and (b) a belief that X maximises utility. Which of Jeremy's judgements count as moral judgements will be determined by the characteristics towards which he takes the moral attitude of approval: since he takes this attitude towards actions which maximise utility, his moral judgements will be those judgements which express beliefs about utility maximisation.

Note that it is the Ecumenical Cognitivist's commitment to the semantic and conceptual necessity of the biconditional relationship between moral judgement and the type of belief assigned priority in the account which leaves it susceptible to Moorean "open question" style worries. Although the Ecumenical Expressivist may well posit a biconditional relationship between moral judgements and certain sorts of belief, that this relationship holds will be a matter of first-order normative theory:

Given deflationism about truth and truth-aptness, the expressivist might hold that moral utterances are truth-apt but deny that their truth-conditions *necessarily* are provided by the beliefs they express. [T]he expressivist might argue that whether an agent's belief provides the truth-conditions for her utterance will be a substantive first-order question and not a question to be settled by metaethical theorizing (2006: 316, emphasis added).

Again

E]ven if normative utterances do express beliefs, as the Ecumenical Expressivist insists, they do not express beliefs which are such that the utterance is *semantically guaranteed* to be true just in case the belief is true (2008: 55, emphasis added).

Since the Ecumenical Expressivist does not view the relationship between the relevant type of belief and moral judgement to hold as a matter of semantic and conceptual necessity, he apparently escapes having to deal with "open question" style considerations.

And note, finally, that the Ecumenical Expressivist view leaves open the possibility of a kind of variability in what constitutes moral judgement. While Jeremy's moral judgements are keyed to utility in virtue of his attitude of approval towards utility maximising actions, Alvin's moral judgements may be keyed to a different characteristic in virtue of his attitude of approval being directed towards actions which instantiate it:

Just what the relevant property is can vary from one speaker to the next. I might approve of actions insofar as they promote happiness, while you might approve of actions insofar as they are in accordance with God's will (2008: 55).

Thus, it may be that Alvin's judgement that X is right expresses (a) an attitude of approval towards actions insofar as they accord with God's will and (b) a belief that X accords with God's will.

4. Ridge's Solution

Ridge – conscious of the problem which undermined Blackburn's attempts at solving the Frege-Geach problem – articulates a constraint which any expressivist account has to meet:

Inconsistency Constraint: the account must explain why someone who accepts the premises of a valid argument involving moral terms, but who denies the conclusion, is making a *logical* mistake. This inconsistency must be logical, rather than the pragmatic inconsistency exemplified by "Moore's paradox" style sentences, e.g. "I believe that P, but not-P" (see Ridge 2006: 313).

Since the expressivist has not – prior to solving the Frege-Geach Problem – earned the right to think of moral judgements as true or false, Ridge works with a notion of valid argument designed to avoid begging any questions by assuming that moral judgements can be regarded as having truth-values:

Validity: An argument is valid just in case any [logically] possible believer who accepts all of the premises but at one and the same time denies the conclusion would thereby be guaranteed to have inconsistent beliefs (Ridge 2006: 326, "logically " inserted).

We can see how Ecumenical Expressivism proposes to solve the Frege-Geach problem by focussing on the "Plain Vanilla" version outlined above, using our utilitarian speaker Jeremy as a representative believer. Suppose that Jeremy accepts premises (1) and (2) but rejects the conclusion (3). In virtue of accepting premise (1), Jeremy expresses the belief that murder maximises disutility; in virtue of accepting premise (2) he expresses the belief that if murder maximises disutility then getting Peter to murder people maximises disutility; in virtue of rejecting (3) he expresses the belief that getting Peter to murder people does not maximise disutility. He thus has straightforwardly inconsistent beliefs. So the argument is valid.

Nothing turns on Jeremy in particular. Suppose that Alvin accepts premises (1) and (2) but rejects the conclusion (3). In virtue of accepting premise (1), Alvin expresses the belief that murder clashes with God's will; in virtue of accepting premise (2) he expresses the belief that if murder clashes with God's will then getting Peter to murder people clashes with God's will; in virtue of rejecting (3) he expresses the belief that getting Peter to murder people does not clash with God's will. He thus has straightforwardly inconsistent beliefs. So, again, the argument is valid.

Ecumenical Expressivism thus exploits the fact that moral judgements express beliefs as well as desire-like attitudes to avoid the Frege-Geach Problem. In the

remainder of the paper, we'll outline three problems that suggest that Ecumenical Expressivism fails to provide a convincing solution to the Frege-Geach Problem.

5. First Problem: Security Against Equivocation?

What guarantees that "Murder is wrong", as it appears in the antecedent of (2), has the same meaning as it has in the initial premise (1)? Recall that the truth-conditions of the beliefs about (dis)utility expressed by Jeremy's moral judgements are not semantically or conceptually guaranteed to be the truth-conditions of those judgements: this is what makes the view a form of Ecumenical Expressivism as opposed to Ecumenical Cognitivism. So the fact that beliefs about (dis)utility are expressed by Jeremy's acceptance of (1) and acceptance of (2) cannot on its own secure the univocity of "Murder is wrong" as it appears in those premises. In order to secure the argument against equivocation, note has to be taken in addition of the role played by the desire-like attitude expressed. Ridge's idea (see e.g. 2014: 152) is that this remains constant in the states of mind expressed by the acceptance of the premises and the rejection of the conclusion and that it is the *combination* of this attitude and the relevant beliefs about e.g. (dis)utility that guarantees univocity.[8] As a first pass, we can say that the hybrid states of mind Jeremy expresses in virtue of accepting (1) and (2) and rejecting (3) are:

(i) (Belief that murder maximises disutility, B!(actions which increase disutility))

(ii) (Belief that if murder maximises disutility then getting Peter to murder people maximises disutility, B!(actions which increase disutility))

(iii) (Belief that getting Peter to murder people does not maximise disutility, B! (actions which increase disutility))

We will now argue that this fails to secure the inference against equivocation. In order to secure univocity, the contribution of the antecedent ("Murder is wrong") to the meaning of the entire conditional (2) must be given by the state of mind expressed by the antecedent as it appears in the asserted context (1). In order to see how Ridge's account fails to do this, note first that in order for the belief that murder maximises disutility and the general sentiment *B!(actions which increase disutility)* to conjointly constitute a moral judgement they have to be related in some way: Ridge says explicitly (2008: 71) that normative judgement is constituted by there being a *link* between the relevant belief and desire-like attitude, and he also (2014: 195) refers to it as a "relational state".[9] (At a minimum, presumably, the belief and desire-like attitude need to be able to interact with each other in the psychological economy of the relevant agent). Suppose that the relevant relation is R. Then, the state of mind expressed in virtue of Jeremy's acceptance of (1) is

(i*) R(belief that murder maximises disutility, *B! (actions which increase disutility)*)

---

[8] See also Schroeder (2009b: 197-8).
[9] See also Schroeder (2013: 307-8).

In other words, the complex state of mind that consists in the belief that murder maximises disutility standing in the relation R to the general sentiment *B! (actions which increase disutility)*.

Likewise, the state of mind expressed in virtue of Jeremy's acceptance of the conditional (2) is

(ii*) R(belief that if murder maximises disutility then getting Peter to murder people maximises disutility, *B!(actions which increase disutility)*)

i.e. the state of mind that consists in the belief that if murder maximises disutility then getting Peter to murder people maximises disutility standing in relation R to the general sentiment *B! (actions which increase disutility)*.

  Our key claim here is that since the state of mind contributed by "murder is wrong" to (ii*) is not (i*), Ridge fails to deal convincingly with the problem about equivocation. It is perhaps easiest to see this by reflecting on the fact that the state of mind (i*) is not a component of the state of mind (ii*) is the way in which, on Blackburn's 1984 account, the state of mind expressed by (1) is a component of the state of mind expressed by (2). Recall from §2 above that for Blackburn the state of mind expressed by (1) is

(i**) B! (murder)

while the state of mind expressed by (2) is

(ii**) H! [*B! (murder)*; B! (getting Peter to murder people)]

Here, the contribution of "murder is wrong" to the state of mind expressed by the conditional (*italicised*) is given by the very same state of mind expressed in the simple asserted context. This is not the case in Ridge's Ecumenical Expressivist account: the complex state of mind (i*) is not what "murder is wrong" contributes to (ii*). Hence Ecumenical Expressivism fails to secure univocity, and the security against equivocation required for a viable solution to the Frege-Geach problem is not provided.[10], [11]

---

[10] It appears that the most Ridge can say is that "Murder is wrong" contributes the relation R, the belief that murder maximises disutility and the attitude B!(actions which increase disutility). On its own, this isn't sufficient to guarantee univocity: it is consistent with e.g. "Murder is wrong" contributing the state of mind R(B!(disutility causing actions), belief that murder causes disutility), and since we don't know whether R is symmetric, this may well not be the same state of mind expressed in virtue of Jeremy's acceptance of (1). The most that Ridge can legitimately say here is that "Murder is wrong" contributes R, the belief that murder maximises disutility, and B! (actions which increase disutility, but – crucially – not in a way that displays them as determinants of the state of mind expressed by (i*). (The argument of this section was sparked by a suggestive comment by Neil Sinclair, and deploys a strategy similar to that used in

Sinclair (2011) against the account of sentential negation developed in Schroeder (2008b)).

[11] On Ridge's account, how does acceptance of the premises in a moral modus ponens argument commit me to acceptance of the conclusion? An answer might be that in accepting the premises I express beliefs whose acceptance commits me to the belief expressed by the conclusion. But how do these beliefs commit me to the desire-like attitude expressed in accepting the conclusion? John Eriksson notes Mark Schroeder's suggestion (2009b: 198) that the key to this is the idea – noted above – that acceptance of any moral sentence containing e.g. "wrong" will for me express the same desire-like attitude. Eriksson argues against this that while this explains why someone who accepts the premises *has* the desire-like attitude prescribed by the conclusion, it fails to explain why someone who accepts the premises is *committed* to accepting the conclusion. He writes:

> [I]t seems more reasonable to think that the kind of attitude prescribed by the conclusion is a new attitude and not an attitude one has merely in virtue of accepting the premises. For instance, it seems conceivable that an agent accepts the premises yet fails to accept the conclusion, but if someone who accepts the premises already has the desire-like attitude prescribed by the conclusion, this seems impossible (2009: 15-16).

The obvious reply to this is that someone who has the desire-like attitude expressed by the conclusion need not have the belief it expresses, so that they needn't have the belief-desire pair possession of which would constitute acceptance of the conclusion. Eriksson objects that this misses the point, since:

> [First], it should be possible to accept the premises without thereby having the attitude expressed by the conclusion. Second, the objection turns on the fact that one does not necessarily have the belief expressed in the conclusion. However, it seems possible to have the belief but, for some reason or other, fail to acquire the desire-like state of mind expressed by the conclusion. This still seems to be something that Ridge's view rules out (2009: 16, n.26).

This strikes us as weak. Without additional argument, the unsupported assertion that it should be possible to accept the premises without thereby having the *desire-like attitude* expressed in the conclusion simply begs the question against Ridge. And the possibility that Eriksson mentions in his second point is not ruled out: someone who doesn't accept the premises may on Ridge's account be able to have the belief component of the conclusion without having the desire-like attitude. Eriksson's objection to Ridge thus seems to us to fail. Whether the variant "ecumenical" position he goes on to develop as an alternative to Ridge's is itself plausible is a matter for future discussion. Likewise for the "ecumenical" position developed in Toppinen (2013). (Note that Eriksson (2009) refers to an unpublished paper by Schroeder called "Finagling Frege": the point discussed appears to have appeared in print since in Schroeder (2009b), to which we refer above).

6. Second Problem: Agnostics about First Order Nonconditional Matters

In order to outline this problem we'll work with the "Ideal Observer" version of Ecumenical Expressivism favoured in Ridge (2006). On this, an agent's judgement that e.g. X is morally required expresses (a) an attitude of approval towards actions insofar as they garner approval from a certain sort of ideal observer and (b) a belief that X would garner approval from that kind of ideal observer.

Ridge allows (2006: 334-336) that there are at least two ways in which a conditional statement can be accepted. Consider

(B) If passive euthanasia is sometimes morally required then active euthanasia is sometimes morally required.

The standard way of accepting (B) involves having a state of mind that consists of an attitude of approval towards actions insofar as they garner approval from a certain sort of ideal observer together with a belief that if passive euthanasia (PE) sometimes garners the approval of that sort of ideal observer then so does active euthanasia (AE). Ridge admits that (B) may also be accepted by an agent who has suspended judgement about all first-order moral matters (i.e. someone who neither approves nor disapproves of actions):

Here, I suggest that it is most plausible within the framework of Ecumenical Expressivism to understand such an agent as taking a stand against the approval of certain sorts of observers—those observers who would simultaneously approve of passive euthanasia but at one and the same time not also approve of active euthanasia, say. In the Ecumenical framework, this will amount to the agent's adopting a perfectly general noncognitive attitude, here an attitude of refusal—refusal to approve of an observer unless it has certain features and the belief that such features (once again we have a belief with anaphoric reference back to the content of a noncognitive attitude) preclude simultaneously approving of passive euthanasia while not also approving of active euthanasia (2006: 335).

Ridge notes a potential worry opened up by this sort of multiple realizability:

The only problem, so far as the technical details of the solution to the Frege-Geach puzzle go, would arise if it were possible for someone to accept a conditional premise in the way characteristic of someone who is agnostic on all substantive nonconditional first-order normative claims, while at one and the same time accepting a nonconditional substantive first-order premise in the more standard way. For in this sort of case, if it were possible, the belief expressed in the major premise would not "hook up" logically in the right way with the belief expressed by the conditional premise to explain the validity of the argument (2006: 335).

Call this putative "bifurcated" moral agent "Sick Boy". Suppose that he accepts (A) and (B) but rejects (C)

(A) Passive euthanasia is sometimes required

(B) If passive euthanasia is sometimes required then active euthanasia is sometimes required

(C) Active euthanasia is sometimes required

If such a "bifurcated" Sick Boy were possible this would frustrate Ridge's solution to the Frege-Geach problem: bifurcated Sick Boy would accept (A) and (B) and reject (C) but would not thereby be guaranteed to have inconsistent beliefs, so that we would have a plainly valid argument that turned out not to be valid on Ridge's conception of validity. However, Ridge argues that bifurcated Sick Boy isn't in fact possible:

[S]uch cases are not possible on the theory on offer here, properly understood. For if someone does have a normative outlook at all, as they must to accept an atomic judgment like passive euthanasia is right, then they can only count as making the relevant conditional judgment if they have the right sort of belief about that observer. Refusing to approve of certain sorts of observers can play a role in conditional (and other nonatomic) moral judgments only when someone lacks a normative outlook. *Once someone adopts a general normative stance by approving of a certain sort of observer, it is plausible to hold that this is dominant in determining their normative judgments, including their conditional judgments, and that they therefore simply do not count as judging, for example, that if passive euthanasia is right then so is active euthanasia unless they believe that the observer they take to be ideal would approve of the former only if he also approved of the latter* (2006: 335-336, emphasis added).

How plausible is Ridge's claim that there cannot be an agent who accepts nonconditional moral statements in the standard way and conditional moral statements in the manner of an agnostic about first order moral matters? It might well be true as a matter of empirical fact (or possibly even as a matter of psychological necessity) that the normative stance of the non-agnostic about first order nonconditional statements would be dominant and come into play in the agent's acceptance of conditional moral statements, but the crucial question is *whether this is so as a matter of logical necessity*: so long as bifurcated Sick Boy is logically possible, we have on Ridge's account a logically possible agent who accepts the premises of a moral modus ponens argument while rejecting the conclusion but who is not thereby guaranteed to have inconsistent beliefs.

Is bifurcated Sick Boy logically impossible? Let's think about his overall state of mind. In virtue of accepting (A) Sick Boy approves of actions insofar as they garner approval from a particular kind of ideal observer (call him I), and he believes that passive euthanasia sometimes garners approval from I. In virtue of rejecting (C) he approves of actions insofar as they garner approval from I but believes that active euthanasia does not sometimes garner approval from I. Putting these together we can say that Sick Boy approves of an observer (I) who sometimes approves of passive euthanasia without sometimes approving of active euthanasia. However, in virtue of his acceptance of the conditional (2) in the manner of an agnostic about first-order moral matters he refuses to approve of an observer unless that observer has some feature which precludes sometimes approving of passive euthanasia without sometimes approving of active euthanasia. *The most we can say about Sick Boy is that in approving of I he does something that he has a stance of refusing to do.* Plainly, this is a moral failing not unlike that of the agent who fails to have every combination of attitudes of which he himself approves. He fails to live up to his commitments. Agents who fail to live up to their commitments in this way are logically possible! Moreover, such agents commit no logical error: if there were some logical incoherence in failing to live up to one's commitments (in doing what you have a stance of refusing to do) Blackburn's solution to the Frege-Geach Problem would not have succumbed to the objection from Wright outlined in section 2 above. Since there is no logical incoherence in the idea of bifurcated Sick Boy, Ridge fails to dispatch the worry opened up by his concession that there are multiple ways in which a conditional

moral statement can be accepted.

7. Third Problem: Variability Within a Single Agent

Recall from §3 above that on Ecumenical Expressivism it is possible for different speakers to make identical moral judgments in different ways. Reverting back to "Plain Vanilla" Ecumenical Expressivism, it may be that Jeremy's judgement that x is right expresses a complex state of mind consisting of a generalised attitude of approval *H!(actions which maximize happiness)* together with the belief that x maximizes happiness, while Onora's judgement that x is right expresses a complex state of mind consisting of a generalized attitude of approval *H!(actions which comply with the Categorical Imperative)* together with the belief that x complies with the Categorical Imperative.[12]

We might ask: if we can have this sort of variability between different speakers, why not within a single speaker at a single time with respect to different types of claim. For example, say that Dee is a utilitarian vis a vis some non-conditional claims but a Kantian vis a vis some conditional claims. Then suppose that Dee accepts (a) and (b) but rejects (c) in:

(a) x is right

(b) If x is right then y is right

(c) y is right.

Then Dee will have the following hybrid states of mind:

(a*) H!(things which maximize happiness); belief that x maximizes happiness.

(b*) H!(actions which comply with the Categorical Imperative); belief that if x complies with the Categorical Imperative then y complies with the Categorical Imperative.

(c*) H!(things which maximize happiness); belief that y does not maximizes happiness.

There is no inconsistency in Dee's beliefs, supplying a counterexample to Ridge's account of validity.

Ridge must therefore argue that an agent like Dee is logically impossible. Let's call the attitudes H!(things which maximize happiness) and H!(actions which comply with the Categorical Imperative) *normative perspectives*. Our question is therefore whether there is some logical or conceptual incoherence in the idea of someone occupying variable normative perspectives in the manner of Dee. What does Ridge have to say about this?

In his 2014 book Ridge introduces the notion of a normative perspective, where

---

[12] For ease of exposition we here suppress mention of the relation which binds the belief and the attitude together in the complex state of mind: nothing turns on this here.

this is defined as the complete set of an agent's "emotionally tinged self-governing policies" (2014: 152) rather than in terms of a single generalised attitude of approval or disapproval. That an extension of this sort is required is shown by examples such as the conditional

(E) If x is right then y is wrong.

The complex state of mind expressed when Jeremy accepts this will need to contain both a generalised attitude of approval and a generalised attitude of disapproval together with beliefs keyed to the characteristics which the attitudes are directed at:

(E*) belief that if x maximizes happiness then y maximizes unhappiness; {H!(actions which maximize happiness, B!(actions which maximizes unhappiness)}

The normative perspectives that Ridge speaks of in his 2014 are simply generalized versions of the set which forms the second component of (E*).

      To return to our question: is an agent like Dee, occupying different normative perspectives vis a vis conditional and nonconditional statements, logically possible, so that equivocation in the beliefs relevant to the validity of an argument results in some valid arguments being deemed invalid? Considering a worry along these lines, Ridge writes:

> Given that an agent can at any given point in time have only one normative perspective this ensures that [there is no equivocation among] the beliefs relevant to testing the validity of the relevant arguments (2014: 152).

In Dee's case, the description of the single normative perspective that he occupies would presumably consist of the attitude *H!(things which maximize happiness)* and the attitude *H!(actions which comply with the Categorical Imperative)* together with some indication to the effect that the former kicks in when Dee is considering nonconditional statements while the latter kicks in when he is considering conditional statements. Presumably, equivocation is avoided because the contents of the beliefs involved become disjunctive. In the example above Dee's beliefs will include: the belief that x either maximizes happiness or complies with the Categorical Imperative, the belief that if x maximizes happiness or complies with the Categorical Imperative then y maximizes happiness or complies with the Categorical Imperative, and the belief that y neither maximizes happiness nor complies with the Categorical Imperative. These beliefs are inconsistent as a simple matter of logic, so that the alleged counterexample of Ridge's account of validity is avoided.

      However it turns out that this "solution" is only made possible because of a *stipulative definition* Ridge makes concerning "normative perspective":

> Another important feature of the view is that, by definition, a speaker will count as occupying at most one normative perspective at any given point in time. Whenever it seems that a speaker occupies more than one, the right thing to say is that his normative perspective is really the conjunction of what one might otherwise take to be his normative perspectives. This is simply how I am defining *normative perspective* here, as a term of art – they are by definition maximally general in this way (2014: 121).

An agent can at any given point in time have only one normative perspective because normative perspectives are just defined as the totality of the relevant sorts of emotionally tinged self governing policies (2014: 152).

It follows from this that the "solution" to the Frege-Geach offered by Ridge is merely a trivial consequence of a stipulative definition: Ridge has simply defined "normative perspective" in such a way that normative perspectives are guaranteed to have a characteristic (non-variability in a single agent at a single time), a consequence of which is that in accepting the premises but rejecting the conclusion of a moral modus ponens argument the relevant agent has inconsistent beliefs. What Ridge owes us is some *non-ad hoc, substantive* reason for thinking that no logically possible believer can occupy variable normative perspectives in this way. Given that this has not been provided, we have not been given a compelling solution to the Frege-Geach Problem.

8. Schroeder's Objection

It might be worthwhile at this point to pause briefly in order to explain how our objection to Ridge's attempted solution of the Frege-Geach problem differs from an objection that has been developed by Mark Schroeder (Schroeder 2009a).

Schroeder's objection starts out from the observation that Ridge's 2006 account of moral sentences sees them as involving a kind of sentential anaphora. "Murder is wrong", for example, is held by Ridge to express (A) a desire-like sentiment of disapproval towards action-types insofar as they possess a certain property and (B) a belief that murder possesses *that* property. The pronoun in (B) is anaphoric on the reference to the property in (A). Now consider the following:

(a) Superman flies.
(b) If Clark Kent flies then I'm a walrus. So,
(c) I'm a walrus.

This is truth-preserving but not logically valid: someone who isn't party to the substantive information that Superman and Clark Kent are the same man could rationally accept (a) and (b) and deny (c). Likewise for

(d) Superman – he flies.
(e) But Clark Kent – if he flies then I'm a walrus. So,
(f) I'm a walrus.

This is truth-preserving given the preferred interpretation of "Superman" and "Clark Kent", but for logical validity we require truth-preservingness in *any* model, not just in the preferred interpretation.

According to Schroeder the moral modus ponens argument is akin to these because seeing that the moral MPP argument is truth-preserving on Ridge's interpretation requires knowledge of the substantive assumption that moral sentences all express the same desire-like attitude. Without that assumption there is no guarantee that the belief expressed in the first premise of the moral MPP is the same as that expressed in the antecedent of the conditional second premise. So Ridge has not captured the logical

validity of moral MPP and so has failed to solve the Frege-Geach problem "on the cheap".

Schroeder's objection is subtle and deserves more careful attention than we can give it here. However, it does seem to us that Schroeder's objection is somewhat narrower than that presented in some of the influential presentations of the Frege-Geach problem in its application to Blackburn's quasi-realism, such as Hale (1986, 1993) and Wright (1988). There the objection seems to be that Blackburn cannot frame the moral MPP argument in a way that satisfies some expressivist surrogate of the notion of truth-preservingess. The moral MPP argument on Blackburn's account doesn't do this because it is no better than an argument that equivocates and which has true premises and a false conclusion – and which is therefore *a fortiori* not truth-preserving (or possessed of a surrogate thereof). We see the objection we raised against Ridge above as concerning this more general worry: the moral MPP argument on Ridge's interpretation is not *even* truth-preserving (because of its failure to deal with the worries about equivocation) and is therefore not logically valid (since being truth-preserving is a necessary – though not sufficient – condition for logical validity). Whereas Schroeder's worry is that on Ridge's account moral modus ponens arguments are truth-preserving but not truth-preserving in virtue of their form, our worry is that they are not truth-preserving at all. [13]

9. Conclusion

Overall, we can conclude that Ecumenical Expressivism does not offer a solution to the Frege-Geach problem that succeeds where the solutions offered by Non-Ecumenical Expressivism fail.

What lessons can we draw from this discussion for moral psychology in general? Ridge is committed to the Humean view that beliefs and desires are "distinct existences" (2014: 49-50). Abstracting a little from the specifics of our argument, what seems to be driving the problem for Ridge is this: to get the right kind of guarantee needed for a successful solution to the Frege-Geach problem you need a much tighter connection between the belief and desire-like elements posited than Ridge's account allows. [14] To put this into the context of the history of moral psychology, we can see now why one might be driven to posit a "besire"-friendly view, where moral judgements are taken to express

---

[13] This is not to say that Schroeder's objection to Ridge's account of formal validity is not a good one, just that it is not the most fundamental problem in the vicinity. In fact, Ridge attempts in his 2014 to extend his 2006 account of validity in a way that speaks to Schroeder's objection: see (Ridge 2014: 153-159). We remain neutral here on whether the developments introduced by Ridge succesfully deal with Schroeder's objection.

[14] Although this has not formed part of our case here, we suspect similar considerations apply to Ecumenical Cognitivism as construed by Ridge, and its attempt to secure motivational internalism – again, the framework Ridge provides doesn't allow for a tight enough connection between the cognitive and the conative to do justice to the phenomenon in question.

unitary mental states with both desire-like and belief-like features.[15] Whatever the deficiencies of such a position at least the view earns a robust connection between desire-like and belief-like features through commitment to a non-Humean metaphysics of mental states. What we are suggesting is that Ridge cannot have his cake and eat it: without a more radical departure in our theory of motivation than he countenances a viable solution to the Frege-Geach problem will elude him. Alternatively, one could retain a commitment to a Humean theory of motivation but then the view will have no substantial advantage over other, non-ecumenical, versions of expressivism that allow for ethical statements to communicate descriptive information.[16] Thus the terrain of moral psychology is much more tightly constrained than in Ridge's vision.

This final consideration allows us to note that Ecumenical Expressivism's inability to succeed where Non-Ecumenical Expressivism fails should perhaps have been obvious from the start. In *Spreading The Word*, Blackburn wrote:

We can see that it does not matter at all if an utterance is descriptive as well as expressive, provided that its distinctive meaning is expressive. It is the *extra import* making the term evaluative as well as descriptive, which must be given an expressive role. It is only if that involves an extra truth-condition that expressivism about values is impugned (Blackburn 1984: 169-70).

In effect, Blackburn is here countenancing the type of Ecumenical Expressivist view favoured by Ridge. It seems, then, that either Ridge has a simple solution to the Frege-Geach Problem that Blackburn somehow missed despite countenancing the possibility of the view or what Ridge takes be to be a simple solution to the Frege-Geach Problem is in fact no solution at all. The problems outlined above suggest that the latter is the case

References

Altham, J. 1984: "The Legacy of Emotivism", in G. Macdonald and C. Wright (eds.) *Fact, Science and Morality* (Oxford: Basil Blackwell).
Blackburn, S. 1984: *Spreading the Word* (Oxford: Oxford University Press).
Blackburn, S. 1993: *Essays in Quasi-Realism* (Oxford: Oxford University Press).
Blackburn, S. 1998: *Ruling Passions* (Oxford: Oxford University Press).
Eriksson, J. 2009: "Homage to Hare: Ecumenicism and the Frege-Geach Problem", *Ethics* 120, 8-35.
Geach, P.T. 1960: "Ascriptivism", *Philosophical Review* 69, 221-225.
Geach, P.T. 1965: "Assertion", *Philosophical Review* 74, 449-465.
Gibbard, A. 1990: *Wise Choices, Apt Feelings* (Cambridge MA: Harvard University Press).
Gibbard, A. 2003: *Thinking How to Live* (Cambridge MA: Harvard University Press).
Hale, B. 1986: "The Compleat Projectivist", *Philosophical Quarterly* 34, 65-84.
Hale, B. 1993: "Can There Be a Logic of Attitudes?", in J. Haldane and C. Wright (eds.) *Reality, Representation and Projection* (Oxford: Oxford University Press).
Miller, A. 2013: *Contemporary Metaethics: An Introduction* (Cambridge: Polity Press)
Ridge, M. 2006: "Ecumenical Expressivism: Finessing Frege", *Ethics* 116, 302-36.

---

[15] See for instance Altham (1984). For discussion of the deficiencies of this kind of view, see Smith (1994).

[16] For example, if you know that I morally approve of all and only actions that maximize the number of green things in existence, you will be able to infer from my calling an action right that I believe it will maximize the green things in existence. For a brief overview of views in this ballpark see van Roojen (2018).

Ridge, M. 2007: "Epistemology for Ecumenical Expressivists", *Proceedings of the Aristotelian Society Supplementary Volume* LXXXI, 83-108.

Ridge, M. 2008: "Ecumenical Expressivism: The Best of Both Worlds?", in R. Shafer-Landau (ed.) *Oxford Studies in Metaethics* vol. 2 (Oxford: Oxford University Press), 51-76.

Ridge, M. 2009: "The Truth in Ecumenical Expressivism", in D. Sobel and S. Wall (eds.) *Reasons for Action* (Cambridge: Cambridge University Press).

Ridge, M. 2014: *Impassioned Belief* (Oxford: Oxford University Press).

Schroeder, M. 2008a: "What is the Frege-Geach Problem?", *Philosophy Compass* 3/4, 703-720.

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

Schroeder, M. 2009a: "Hybrid Expressivism: Virtues and Vices", *Ethics* 119, 257-309.

Schroeder, M. 2009b: *Noncognitivism in Ethics* (New York: Routledge).

Schroeder, M. 2013: "Tempered Expressivism", *Oxford Studies in Metaethics* 8, 283-313.

Sinclair, N. 2011: "Moral Expressivism and Sentential Negation", *Philosophical Studies* 152, 385-411.

Smith, M. 1994: *The Moral Problem* (Oxford: Blackwell).

Toppinen, T. 2013: "Believing in Expressivism", *Oxford Studies in Metaethics* 8, 252-282.

Van Roojen, M 1996: "Expressivism and Irrationality", *Philosophical Review* 105, 311-335.

Van Roojen, M. 2018: "Moral Cognitivism vs Non-Cognitivism" *Stanford Encylopedia of Philosophy*.

Wright, C. 1988: "Realism, Anti-Realism, Irrealism, Quasi-Realism", reprinted in *Saving The Differences* (Camb, MA: Harvard University Press 2003).