



# Perspectival representation and fallacies in metaethics

Max Kölbel<sup>a,b</sup>

<sup>a</sup>Philosophy Department, University of Vienna, Vienna, Austria; <sup>b</sup>Philosophy, ICREA (Catalan Institution for Research and Advanced Studies), Barcelona, Spain

## ABSTRACT

The prevailing theoretical framework for theorising about representation construes all representation as involving objective representational contents. This classic framework has tended to drive philosophers either to claim that evaluative judgements are representations and therefore objective, or else to claim that evaluative judgements are not really representations, because they are not objective. However, a more general, already well-explored framework is available, which will allow theorists to treat evaluative judgements as full-fledged representations (thus doing justice to their representational aspects) while leaving open whether they are objective. Such a more general conception of representational content is exemplified, e.g. by Lewis's 'centred contents' and Gibbard's framework of 'contents of judgement', thus it is not new. I shall start in §1 by introducing the more general framework of perspectival contents and then illustrate in §2 how awareness of it can help expose the fallaciousness of certain widely used forms of argumentation in metaethics.

**ARTICLE HISTORY** Received 31 December 2017; Accepted 10 January 2018

**KEYWORDS** Perspectival content; centred content; content; representational content; representation; Frege-Geach Problem; metaethics; David Lewis; Allan Gibbard; Nicholas Unwin; Mark Schroeder

## 1. Perspectival representational contents

### 1.1. Representation

All representation minimally involves representing things as being some way. If things are as represented, the representation is accurate or correct. For example, a representation might represent things in general as being such that there are cats, or it might represent a particular cat as being large, amusing or hungry. We

**CONTACT** Max Kölbel  [max.koelbel@univie.ac.at](mailto:max.koelbel@univie.ac.at), [max.kolbel@icrea.cat](mailto:max.kolbel@icrea.cat)

© 2018 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.  
This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way.

will here be interested in certain types of representation, namely representations that are mental states, such as beliefs, or linguistic acts, such as assertions.

Uncontroversial examples of representational mental states or linguistic acts would be, for example:

- (B1) My belief that there are cats.
- (B2) My belief that Apolo is a cat.
- (B3) My belief that my neighbours have a cat.
- (B4) My belief that Apolo is my neighbours' cat.
- (U1) My utterance of the sentence 'There are cats.'
- (U2) My utterance of the sentence 'Apolo is a cat.'
- (U3) My utterance of the sentence 'My neighbours have a cat.'
- (U4) My utterance of the sentence 'Apolo is my neighbours' cat.'

For one reason or another, the following mental states and linguistic acts are not uncontroversial examples of representations:

- (B5) My judgement that Apolo is amusing.
- (B6) My judgement that it is bad to leave Apolo alone in the flat.
- (U5) My utterance of the sentence 'Apolo is amusing.'
- (U6) My utterance of the sentence 'Leaving Apolo alone in the flat is bad.'

I think I could have equally characterized the mental states (B5) and (B6) as my *beliefs* that Apolo is amusing and that locking him in is bad respectively. I chose the term 'judgement' in order to have a term that is more neutral and is not associated with certain assumptions that are usually made about beliefs in the context of non-cognitivism, namely that beliefs are descriptive, representational, truth-apt, objective, etc.

Why should these examples be controversial? I suggested above that representations represent things as being some way, and that they are accurate or correct just if things are as represented. If that is all there is to representation, then (B5), (B6), U5) and (U6) should qualify as representations, for each of them seems to represent things in some way, and each of them will be accurate or correct if and only if things are as they have been represented.

However, the trouble starts when we specify how things are represented by a given representation (and thereby also specify what would be required for the representation to be correct). It seems that B2 represents Apolo as a cat, and that B2 will therefore be correct just if Apolo is indeed a cat. Similarly, we might say, B5 represents Apolo as amusing, and it should be regarded as correct just if Apolo is indeed amusing. However, at this point, some will say that whether B5 is correct depends not just on how things are with the object of representation,

Apolo, but also on how things are with me, the owner of B5. To see this, compare a further mental state, and corresponding utterance:

(B7) Your judgement that Apolo is amusing.

(U7) Your utterance of 'Apolo is amusing.'

Let us suppose that at the time of B5, my dispositions to be amused are such that Apolo is (at the time of B5) liable to amuse me, but that your dispositions (at the time of B7) are such that he will not amuse you. It would seem, then, that B5 is correct, while B7 is not. Nevertheless, both B5 and B7 can be described as representing the same object, Apolo, in the same way, namely as amusing. So, if we accept this description of the representational content of B5 and B7 (they both represent Apolo as amusing), then clearly the correctness of B5 and B7 does not depend only on how things are with Apolo, the object of representation. It seems to depend also on how things are with the owner of the mental state in question, for it is the difference between me and you (the respective owners of B5 and B7) that accounts for the difference in correctness.

Some philosophers think that we cannot, therefore, construe B5 and B7 as representing things in the same way, i.e. as having the same representational content.<sup>1</sup> On this kind of view, representational contents must be construed in such a way that if two representations (that occur in the same world) have the same content, then it cannot be that one is correct, but the other is not. Representational content, on this conception, determines correctness. Or perhaps I should precisify: given a state of the world, representational content determines representational correctness.

## 1.2. *The classic conception*

This sort of conception is in the tradition of Frege, who stipulated that *thoughts* (his analogue of contents) were to determine a truth-value, and did not allow that anything short of determining a truth-value could qualify as a thought (Frege 1918). On this conception, representational contents are by definition objective: if anyone believes such a content correctly at any time, then any belief with the same content, no matter by whom or at what time, will also be correct (assuming that truth is necessary and sufficient for correctness). Let's call this the 'classic conception', and let's call representational contents construed to meet this constraint 'classic contents'.

On the classic conception, then, if we are to say that B5 is correct while B7 is not, we *have to* construe them as having distinct representational contents, i.e. as not representing things in the same way. For example, on the classic conception, one might say that B5 and U5 represent Apolo as amusing *me*, while B7 and U7 represent Apolo as amusing *you*, thus assimilating B5 and B7 to B3 (or U3): my belief that my neighbours have a cat does not have the same content as B8 (U8):

(B8) Your belief that your neighbours have a cat.

((U8) Your utterance of 'My neighbours have a cat.')

even though you would use the same words to express B8 (as in U8), or to ascribe it to yourself ('I think my neighbours have a cat.'). In the case of B3 and B8, we are less tempted to say that they share the same content, for we would not use the same words to specify how they represent things as being. This was different in the case of B5 and B7: both can be correctly described (by me, you or anyone else) as representing Apolo as amusing. On the classic conception, however, this description is not complete and masks a difference in representational content and B5 and B7 should be treated as having contents, respectively, like those of B9 and B10:

(B9) My belief that Apolo amuses me.<sup>2</sup>

(U9) My utterance of 'Apolo amuses me.'

(B10) Your belief that Apolo amuses you.

(U10) Your utterance of 'Apolo amuses me.'

### **1.3. *Perspectival content***

The classic conception, however, is not the only option. There is no compelling reason for stipulating that representational content must determine correctness (i.e. is objective by definition). Alternative, perspectival conceptions drop this requirement, so that we can, after all, maintain that the way B5 and B7 represent things as being is the same. Both B5 and B7 represent Apolo as amusing, but they do so, as it were, from different perspectives. On this conception, then, whether a representation with a certain representational content is correct can depend on who represents and at what time. We might say: whether it is correct to represent things a certain way, whether it is correct to believe a certain representational content, depends on the identity of the thinker and the time of thinking. In other words, this conception operates with perspectival contents, contents the correctness of believing which depends on a perspective or standpoint. For our purposes, we can think of perspectives as ordered triples of a thinker, a time and a possible world. On this conception, while a content (and the world) alone do not determine a truth-value, a content does so relative to a perspective.

I proposed initially that representation minimally involves representing things as being a certain way and that representations are correct iff things are the way they are represented as being. The classic conception adds a further objectivity requirement, namely that the ways things are represented as being can only be specified in a certain way, namely such that if two representations share the same representational content, they cannot differ with respect to correctness. Construing representational contents as perspectival involves maintaining the essential feature of representation but dropping the classic

constraint on representational content. As a result, the perspectival approach will allow us to classify B5 and U5 as representations.

B6 and U6 – the judgement and claim that leaving Apolo alone in the flat is (morally) bad – are subject to similar worries as B5 and U5, in so far as there are philosophers who argue that the correctness of moral judgements may, in addition to depending on facts concerning the objects of moral evaluation, depend on features of the judge, such as her system of moral norms, or that of her community. Those who argue this will have options that are analogous to the ones just discussed for B5 and U5: adopt a classic or a perspectival conception.<sup>3</sup>

Besides allowing B5/U5 and B7/U7 to be uncontroversially classified as representations, *and* as representations with the same content, the perspectival approach also permits more *differentiation* in some cases of uncontroversial representations. Thus, we can distinguish between the content of B3 and the content of B11, while arguably the classic conception does not:

(B11) My belief that MK's neighbours have a cat.

It would also allow us to treat, for example, B9 and B10 as having the *same* content: they both represent the owner of the belief as being amused by Apolo. Of course this will be another case where correctness depends not only on how things are represented as being, but also on who does the representing.<sup>4</sup>

#### 1.4. *What is at stake?*

It is not obvious what, if anything, is at stake in the choice between classic and perspectival conceptions of representational content. In my view, the situation is best described as follows.

Focus on our minimal starting point: representations represent things as being some way and they are correct iff they represent things as they are. However, there are many different respects of similarity that one could usefully focus on when thinking about how representations represent things as being and the conditions under which they would be correct. In some respects, B3 and B8 are similar: both of them represent the thinker's neighbours as having a cat. Both of them are representations that are correct just if the thinker who has the representation has neighbours who have a cat. A perspectival conception of content allows modeling this similarity as sameness of content. The proponents of the classic conception may well be interested in this kind of similarity too, they do not, however, want to capture this respect of similarity in terms of sameness of content. Instead, they try to abstract from a different respect of similarity, namely that obtaining between B3 and B11: both represent the same object (MK's neighbours) as being the same way. The respect of similarity they are interested in entails sameness of correctness status: any two representations similar in this respect are also similar in respect of correctness. This means that the classic conception needs to capture the earlier mentioned similarity between B3 and B8 in a way that does not require sameness of content.<sup>5</sup> As long

as each conception is able to offer alternative ways of modelling similarities that are captured by the other, there seems to be nothing substantial at issue between the two approaches.

I conceded that similarities captured as sameness of content on perspectival conceptions can be described in other ways on classic conceptions. However, are perspectival conceptions conversely able to model everything classic conceptions can model? Yes, they can. Let us focus for a moment on a well-known, simple and elegant implementation of the classic conception: the conception of contents as functions from possible worlds to truth-values (or, equivalently, sets of possible worlds). Compare this classic notion of content with a corresponding perspectival notion of contents as functions from 'centred worlds' to truth-values, where centred worlds are simply ordered pairs  $\langle w, c \rangle$  consisting of a world  $w$  and a center  $c$  (where centres are in turn construed as ordered pairs of thinkers and times). David Lewis (1979) has shown that *this* perspectival notion is a generalization of *this* classic notion, in the following sense: Amongst all the perspectival contents some are what has been called 'boring'<sup>6</sup> or 'portable'<sup>7</sup>: the centre-component of each centred world is irrelevant for the value of the function, because if the value of the function is truth for some centred world  $\langle w, c \rangle$ , then it is truth also for every centred world  $\langle w, c' \rangle$  with a different centre, but the same world. Or if we think of contents as sets: a set  $P$  of centered worlds is *portable* just if: for all centers  $c$  and worlds  $w, \langle c, w \rangle \in P$  only if for every center  $c', \langle c', w \rangle \in P$ . Now, for each classic content  $C$ , there is exactly one portable perspectival counterpart  $P$ , such that  $w \in C$  iff for some (or all) centres  $c, \langle c, w \rangle \in P$ . Since there is exactly one counterpart portable centred content for each classic possible-worlds content, and counterparts share all the theoretical properties that do any work in the classic conception,<sup>8</sup> the framework of centred contents can do everything the classic framework can do.

Lewis's point seems to hold not only for unstructured contents, but for any conceptions of classic and perspectival structured content which are such that the only difference between them is that the classic conception imposes an extra condition of objectivity, and thereby rules out all except the portable ones amongst the perspectival contents.<sup>9</sup> For we will then again be able to find exactly one counterpart perspectival content for each classic content, which shares all the theoretically relevant properties. Thus, whatever else we think a content needs to do in addition to determining conditions under which it can be correctly believed: a perspectival version is just a generalization and can do everything its classic counterpart can do.

While this result is good news for the perspectival conceptions, it by no means decides which conception we should adopt. A theorist who is only interested in using her representational contents to explain phenomena that can be explained by classic contents, and who reserves other resources for other phenomena (or is not interested in them) can of course rationally stick to the classic conception.

Even if I am right, and the choice between the two frameworks is a matter of stipulation and a pragmatic choice, it will make a rhetorical and presentational difference whether we do or do not allow non-objective representations. Those convinced that evaluative thought is non-objective, in the sense that correctness depends on one's perspective, standpoint etc., may have a harder time selling their view if they are forced to say that evaluative thought does not qualify as representational (as they would be on the classic view).

The mere rhetorical or presentational significance of the choice, however, can mislead if the category of representation (as classically delimited) is erroneously imbued with further significance. For example, if it is assumed that everything that shares certain patterns of behaviour with paradigm representations should also (on pain of explanatory disunification) be a representation. This would be a fallacy, and demonstrably so, if there is an alternative notion of representation available that would restore explanatory uniformity. The second part of my paper attempts to diagnose such a fallacy.

## **2. Fallacies in metaethics**

### **2.1. Introduction**

I want to illustrate that our choice of notion of representational content – classic or perspectival – can mislead us into substantial philosophical mistakes. I believe that mainstream thinking in metaethics in the last half-century or so has been misguided by the narrow classic picture, because it makes certain metaethical views (non-cognitive ones) appear less unified. The basic theme is the one already mentioned: if we regard all representation as involving objective content, then our conceptual framework, our terminology, portrays anything that is otherwise a bona fide judgement, but lacks such a content, as something categorically different: as not a representation or as lacking proper representational content.

I suggested above that the matter is merely terminological: if we stipulate that representation must involve classic contents, then instead of arguing whether the contents of certain representations are objective, we'll be arguing whether certain mental states are or are not representations. That approach may be rhetorically unfair (to those who have to argue that certain bona fide judgements are not genuine representations), but as long as we do not imbue our choices of conceptual and terminological framework with substantial, perhaps metaphysical or explanatory, significance, we are completely free to choose as we wish. However, it is sometimes hard to remember later on that certain bits of framework that one has grown fond of were in fact the result of a stipulation that one could revoke at any time. The familiarity of notions employed in a theory (e.g. classic content) may also count as an explanatory advantage, but it

is at most an historically contingent advantage that can vanish without residue as rival notions become more familiar.

This is all highly abstract. In order to fill this abstract possibility with concrete life, I will now show how awareness of the perspectival content framework can expose some well-known forms of reasoning in metaethics as fallacious. The fallacious reasoning I have in mind is frequently exhibited in the family of objections against metaethical non-cognitivism and expressivism that is associated with the label 'Frege-Geach point'.

I shall first outline a certain non-cognitivist/expressivist outlook with an emphasis on those aspects that are important for my point. Here I will distinguish non-cognitivism: a view about the nature of the mental states one might call 'moral judgements'; from expressivism: a view specifically about moral language. Then I shall outline some conclusions drawn from the Frege-Geach point that exhibit a fallacy to which I think our stipulation of classic representational contents can lead. In order to demonstrate that I am not merely taking down a straw man, I shall then show how even the most sophisticated and sympathetic engagement with expressivism, namely Mark Schroeder's, is guilty of a version of this fallacy. Finally, I shall outline a rudimentary theory of 'judgements', understood as those representational mental states that have perspectival contents.

## 2.2. *Non-cognitivism*

Metaethical non-cognitivism is a view about the nature of moral judgements, namely the view that moral judgements are non-cognitive mental states. Paradigm examples of cognitive mental states are beliefs, while paradigm examples of non-cognitive mental states are desires or preferences. Thus non-cognitivism is a view about a certain class of mental states, namely the class of moral judgements. It is often pointed out that the term 'judgement' is ambiguous between an act-sense and a state sense. It is here used in the state-sense: it denotes a state of mind, i.e. an opinion or conviction, rather than a mental act of judging.

Expressivism, as I shall explain below, is a view about moral language which, even though usually motivated by non-cognitivism, is a clearly separable idea. Unfortunately, it is not usually clearly separated. For some reason, the canonical or textbook formula widely used to characterize non-cognitivism, emotivism and expressivism is as 'the view that moral judgements express non-cognitive attitudes'. This formula ('The Formula') does not make sense on either of the two senses of 'judgement' I just distinguished. For neither does anyone intend to characterize a view according to which one mental *state* expresses another, nor does anyone intend to characterize a view according to which one mental *act or episode* expresses another. What the many users of The Formula clearly intend is the view that some overt act or behaviour, such as that of *making a moral statement*, expresses a non-cognitive attitude. Thus, the only way to make



sense of 'judgement' in The Formula is to take it in a third sense, namely as denoting an overt act of making a claim or statement. Understood in this way, The Formula does contain a decent enough characterization of what I shall call 'expressivism' below, but this third sense of 'judgement' is highly unusual and therefore potentially confusing. I shall stick to 'judgement' in the state-sense and stay away completely from The Formula. Instead, I shall do as I did above and characterize non-cognitivism as the view that moral judgements *are* non-cognitive states, while characterizing expressivism as the view that moral statements, i.e. utterances of moral sentences, *express* non-cognitive states.

A quick reminder of the typical motivations for non-cognitivism will help us get the right perspective on this family of views. Let me distinguish metaphysical, epistemological and psychological motivations. *Metaphysical*: if moral judgements were cognitive states, i.e. beliefs, then they would aim to represent how things are objectively. Accordingly, any correct moral judgements would be correct because they represents some aspect of objective reality correctly. Therefore cognitivism, if combined with the view that some moral judgements are correct, requires there to be aspects of reality that make moral judgements true: moral facts. This is regarded as problematic by philosophers with a naturalist outlook who also think that there is no naturalistic reduction of moral values.

*Epistemological*: relatedly, if cognitivism is true and some moral judgements do indeed represent moral facts correctly, then how do we acquire knowledge of such facts, given that there seems to be no suitable empirical or a priori form of justification available.

*Psychological*: if we accept a Humean picture of the mind, then beliefs or cognitive states (the outputs of a Humean faculty of reason) cannot by themselves motivate thinkers to act. Beliefs only serve to inform our desires and preferences, but without any desires or preferences they would be completely motivationally inert. Combined with motivational internalism about moral judgements, the Humean picture therefore motivates non-cognitivism: if moral judgements in conjunction with certain beliefs, do motivate us in the absence of any further motivating mental states (desires, preferences), then they must themselves be on the non-cognitive side of Hume's divide.

Based on these typical motivations, let me make some observations that will be relevant later in §2.9.<sup>10</sup> The non-cognitivist's claim is that moral judgements are non-cognitive mental states. Given the psychological motivation, these non-cognitive mental states must fall on the motivating side in a Humean picture of the mind. Thus, the idea is that judgements in general divide into cognitive (motivationally inert) and non-cognitive (intrinsically motivating) ones, and that moral judgements in particular (perhaps alongside other types of value-judgement) fall on the motivating side. The important point to retain here is that non-cognitivists are committed to a specific claim of discontinuity between moral judgements (and perhaps some other value judgements) and

other judgements, concerning their motivational features: moral judgements fall on the motivational side of Hume's divide between reason and passion.

Similarly, the epistemological motivation commits the non-cognitivist to the view that not all judgements admit of straightforward empirical justification, and that only those that do should be regarded as cognitive. Ayer, for example, points out that moral intuition would not qualify as an appropriate source of empirical justification, because there is no way to settle the matter when two different thinkers' intuitions conflict. The metaphysical motivation also relies on the assumption that cognitive states (beliefs), when successful, correspond to objective aspects of reality, i.e. facts. Thus, the epistemological and metaphysical motivation seem to require that cognitive states correspond to, and (when true) are made true by, objective aspects of reality, for which there is an empirical method that could settle conflicts effectively.

To summarise, the non-cognitivist seems committed to three kinds of difference between moral and cognitive judgements: (a) moral judgements do not, like cognitive judgements, represent aspects of objective reality, (b) the method used to justify moral judgements cannot effectively settle inter-subjective disagreements, as empirical methods in science (supposedly) can, and (c) moral judgements are not, like beliefs, motivationally inert. As long as the non-cognitivist retains her specific claim of discontinuity between moral and other types of judgement, she is free to view the wider category of judgements as otherwise highly continuous.

### **2.3. Expressivism**

As already mentioned briefly, expressivism is a thesis about moral language that naturally complements non-cognitivism. It seems undeniable that in some good sense declarative sentences, or their utterances, express judgements. Thus, the sentence 'Gambling is a lucrative business.' serves to express the judgement that gambling is a lucrative business, and the sentence 'Gambling is bad.' serves to express the judgement that gambling is bad. What does it mean to say that a sentence serves to express a mental state, or perhaps more accurately, that we can utter a sentence and thereby express a mental state? There are many things one could say here. Minimally, the claim that a sentence, or utterances of it, express a mental state seems to be a characterization of an aspect of the sentence's conventional meaning. We might say further, and still fairly generally, that the sentence, by virtue of its conventional meaning, is suitable for performing a certain speech-act that requires, as a condition on the speech-act's sincerity, that the sentence-user have the mental state. This much should be fairly uncontroversial.<sup>11</sup>

Expressivism, now, seems to be built on this fairly uncontroversial assumption that sentences are meant (in some sense) for the expression of mental states. It seems sensible to say that moral sentences like 'Gambling is bad.' serve to

express moral judgements in precisely the sense just mentioned. If moral judgements are non-cognitive mental states, as the non-cognitivist claims, then moral sentences express non-cognitive mental states. This is the expressivist's thesis. It follows from non-cognitivism on the simple assumption that moral sentences serve to express moral judgements.

Looked at in the way I just did, expressivism is not a particularly exciting addition to non-cognitivism. However, it immediately looks more interesting once we place it in the context of some surrounding assumptions about language. Let us consider the following assumptions:

A1 Declarative sentences express beliefs.

A2 Declarative sentences are assertoric (i.e. they serve, in virtue of their meaning, for the performance of assertions).

A3 Assertions aim at objective truth (they have objective contents).

The expressivist has to deny A1, because expressivism is the thesis that some declarative sentences, namely moral declarative sentences, do not express beliefs, for they express non-cognitive states. A2 and A3 can be seen as background assumptions that motivate A1. If the expressivist does not deny A2, then at least she has to deny A3.<sup>12</sup>

Thus, the interest of expressivism lies in its denial of A1, which is the result of either denying A2 or A3. This is regarded by many as a theoretical cost, because a certain account of meaning and conversation seems at first sight to depend on these assumptions being true. The expressivist's denial of A1 and either A2 or A3 threatens to throw over board some valuable and successful theories of meaning and conversation. The family of objections surrounding Geach's Frege-point and Unwin's negation problem spell out this worry in more detail, as we shall see in a moment.

But the expressivist need not give up any of the advances that may seem to be encapsulated in A1–A3. The expressivist could, for example, simply accept A1\*, A2 and A3\*:

A1\* Declarative sentences express judgements.

A2 Declarative sentences are assertoric (i.e. they serve, in virtue of their meaning, for the performance of assertions).

A3\* Assertions aim at correctness (they have perspectival contents).

These alternative assumptions about declaratives allow the preservation of everything worth preserving in established theories of meaning and conversation. To see this, let us now examine Geach's Frege point and related objections in a little more detail.

#### **2.4. Geach's Frege point against expressivism**

Geach (1960, 1965) targeted expressivist claims like the following:

(E1) The meaning of the moral predicate 'bad' consists in the fact that it can be used to condemn, or express disapproval of, something.

Thus, Geach's target is a proposal about moral *predicates* rather than the expressivist claim about the meaning of moral *sentences* that I have been discussing. However corresponding claims about moral sentences can presumably be derived from E1:

(E2) The meaning of the moral declarative sentence 'Gambling is bad.' consists in the fact that it can be used to condemn gambling or to express disapproval of gambling.

Geach objected (i) that there were many occurrences of the predicate 'bad' in which it cannot be used to condemn anything – for example when occurring embedded in the following sentences:

(G1) Gambling is not bad.

(G2) If gambling is bad, then seducing others to gamble is bad.

(G3) Is gambling bad?

On these occurrences, the expressivist's meaning analysis cannot be correct: neither of these sentences can be used to condemn gambling or express disapproval of it. Moreover, Geach pointed out (ii) that we can't say that the meaning of 'bad' changes from unembedded uses to embedded uses as in G1–3. For if we say that, then (for example) the valid argument from:

G4 Seducing others to gamble is not bad.

and G2 to G1 as conclusion comes out as *not* of valid *modus tollens* form, but rather as a mere equivocating *modus tollens* look-alike. However, the argument is clearly valid, and it seems to be a perfect instance of *modus tollens*.

Geach's diagnosis was that the expressivist's thesis E1 does not take into account the difference between mere predication and assertion. When 'bad' is predicated of gambling and the resulting content is asserted, we get a sentence that asserts badness of gambling. However, 'bad' is also predicated of gambling in G2 and G3, but without asserting the predication we do not get an assertion of badness (Geach 1960, 223; 1965, 464).

In effect, Geach's point is that, in order to challenge the established account of the function of 'bad' and 'gambling is bad', it is not sufficient for the expressivist only to offer a speech-act analysis of the predicate 'bad' in simple declarative sentences, or to say what mental state such sentences express. For the challenge to be effective, the expressivist's alternative account needs to say something about the unchanging contribution the predicate 'bad' makes on all its occurrences. For this is what the established account does: it is a systematic account that tells us how the propositional content of sentences containing 'bad' result from the meaning of 'bad' and that of other constituent expressions. The communicative function of sentences (e.g. making assertions, expressing belief)

then results from the account of the propositional content together with the type of sentence it is (declarative, interrogative etc.).

To take the simplest case, the established account will say that G1 has as its content the negation of the proposition that gambling is bad, where the negation of a proposition  $p$  is simply the proposition that is true in all those cases in which  $p$  is not true. Thus, the expressivist needs to say likewise what contribution the negation is making on her account.

Geach's point is well taken and essentially correct. The fallacy I am going to diagnose occurs en route to the further conclusion, drawn by many theorists, that the expressivist cannot meet Geach's challenge.

### 2.5. Unwin's negation problem

Unwin's negation problem (Unwin 1999) extends Geach's challenge for expressivism to non-cognitivism. Geach is asking for an account of the meaning of moral expressions that works for all possible embeddings of these expressions. Unwin's negation problem concerns the expressivist's and non-cognitivist's account of the non-cognitive mental states expressed by complex moral sentences. Unwin focusses on negation as the simplest case of embedding, but the problem can be raised equally for other forms of embedding and the corresponding mental states expressed. The problem is this: the expressivist claims that to judge that gambling is bad is to disapprove of gambling. For someone to fail to judge that gambling is bad, correspondingly, should then be construed as failing to disapprove of gambling. However, what is it, then, to judge that gambling is not bad? It couldn't be disapproval of not gambling, because merely thinking that gambling is not bad does not commit one to thinking that not gambling is bad. So it looks like the expressivist has to introduce a further non-cognitive attitude, say *toleration*, in order to account for the judgement that gambling is not bad: it is the same as tolerating gambling:

(N1) Nick judges that gambling is bad. = Nick disapproves of gambling.

(N2) Nick does not judge that gambling is bad. = Nick does not disapprove of gambling.

(N3) Nick judges that gambling is not bad. = ? (proposal: = Nick tolerates gambling.)

(N4) Nick judges that not gambling is bad. = Nick disapproves of not gambling.

Thus, the non-cognitivist, in order to deal with all four states of mind N1–N4 needs to introduce a further non-cognitive attitude, *toleration*.

However, what should be the systematic contribution of negation in thought that allows us to predict that thinking that gambling is bad is disapproval of gambling while thinking that gambling is not bad is the attitude of tolerating gambling? Moreover, similar questions can be raised about other forms of complex thought or embedding, thereby forcing the non-cognitivist-expressivist to resort to further, unsystematic postulations of additional attitudes.<sup>13</sup>

## 2.6. Response: distinguishing force and content

How can the expressivist respond to Geach's challenge? How can the non-cognitivist respond to Unwin's negation problem? Geach's own diagnosis seems sensible. Established theories distinguish the compositional contribution a predicate makes to the content of sentences (predicating a property of something) from the ultimate, and composite, effect that the predicate has when it is used assertorically and without being embedded (attributing the property). If we claimed that 'bad', in virtue of its meaning, serves to attribute the property of badness, we would get into the same kind of trouble as the expressivist who claims that 'bad' serves to condemn, or to express disapproval. 'Bad', when simply predicated of something and then asserted, can be plausibly claimed to attribute badness. But when embedded in negation, a conditional, a disjunction, etc. it will often generate a content asserting which will clearly not amount to attributing badness (see G1, G2). And even a simple predication of 'bad' does not serve for the attribution of badness unless the result is asserted (see G3).

The problem exposed by Geach, therefore, is simply that the expressivist, who claims that 'Gambling is bad' serves to condemn or express disapproval (E2), must separate the content-determining compositional function of 'bad', which it has in all contexts in which it is used with the same sense,<sup>14</sup> from the effect 'bad' has only when used unembedded and asserted. In other words, the expressivist should drop or modify E1, and treat E2 as a derivative fact that results from the compositional meanings of all the syntactic elements in 'Gambling is bad', including its assertoric force indicator (mood, word order). The expressivist should say that on the one hand there is the general effect of the assertoric mood of the sentence, which it has wherever it occurs. On the other hand there is the content one can assert by using the sentence 'Gambling is bad' (which is negated in G1, which occurs as the antecedent of G2 and which it shares with G3). This content is determined by the content-determining semantic features of the sentence's syntactic elements. Instead of E1, the expressivist could claim, say, that the content-determining function of 'bad' is to predicate the *feature* of badness (where 'feature' is supposed to be more general and less committal than 'property'). That would be just as (un)informative as the line according to which 'bad' predicates the property of badness. Talk of 'predicating features' would need to be further explained, just as talk of predicating properties does. This further explanation will need to underpin E2 and the expressivist's general claim that moral sentences serve to express non-cognitive attitudes (just as Geach's talk of predicating properties would ultimately underpin the general view that moral sentences serve to express beliefs).

An answer to Unwin's negation problem (and its generalization to other forms of complex moral thought) would already be contained in such an answer to Geach's challenge. For it will involve separating out in all moral judgements the content from an attitude of acceptance. For example, disapproval of gambling

will be factored into a judgeable content, that gambling is bad, and an attitude of acceptance. Negated, conditional, disjunctive, etc. moral judgements will be analysed as functions of simpler judgeable contents, and even atomic judgeable contents will be constructed from simpler elements that can re-occur in other contents ('concepts').

This kind of answer to Geach and Unwin has had the benefit of being forcefully and clearly presented by Mark Schroeder in a series of papers (e.g. 2008a, 2008b) and in his book *Being For* (2008c). In Schroeder's work, the attitude of disapproval of gambling expressed by 'Gambling is bad.' is factored into a general moral attitude of acceptance or endorsement (called 'being for') and a judgeable content, explicated as 'blaming for gambling', something that one can merely consider without being for it – just as one can predicate without attributing. That the content is *blaming*, rather than, say, praising, caning, thanking or whatever is the compositional contribution made by the expression 'bad', and this is the contribution it always makes, even in the various embedded contexts (e.g. G1–3). This takes care, very nicely, of Geach's challenge and Unwin's negation problem.

Despite being clearly indicated by Geach, this sort of approach curiously has not been particularly prominent in the literature about expressivism (Gibbard 1990, 2003 are exceptions). This is probably because expressivists (and their opponents) were for a long time stuck on E1, the idea that moral predicates are force-indicators, and did not genuinely consider the possibility of factoring the contributions made by the constituents of moral sentences in the more straightforward way suggested by Geach and expounded by Schroeder.<sup>15</sup> Historically, this may well be due to the influence of authors like Hare (1952, 1970) and Blackburn (1984, 1988), who both started with the idea that the linguistic function of moral predicates like 'bad' was fundamentally different from that of descriptive predicates, as they found it in authors like Ogden and Richards (1946), Ayer (1936/1946) or Stevenson (1937).

## 2.7. Fallacies 1

I now want to point out two tempting but fallacious lines of reasoning into which the classic conception outlined in part 1 can mislead us. I leave it to the reader to determine the extent to which these fallacies have in fact occurred. It is enough for my purposes, to establish that these forms of reasoning are indeed fallacious.

Geach's challenge draws attention to the fact that when we give an account of the meaning of an expression, the account should (i) cover all possible occurrences of the expression. In other words, the account of meaning should be *compositional*: we should be able to derive the meaning of complex expressions from the meanings of their constituent parts. Moreover, this compositional account should (ii) correctly predict/explain which sequences of (utterances of) sentences constitute formally valid arguments.

Standard frameworks for modelling meaning deliver (i) and (ii) and involve a classic conception of content. The compositional meaning involves assignments of meanings to (some) simple expressions and some compositional rules that specify how the meaning of compound expressions are determined by the meanings of constituent parts. In this way, the theory assigns to each sentence (or sentence in context) a classic content. Such contents are, or at least determine, functions from the possible state of the world to truth-values.

The compositional rules and semantic assignments predict certain logical properties and relations: the contents of some sentences (in context) are true at every possible state of the world (*logical necessity*). Some sentences are such that if their contents (in a given context) are true at a certain state of the world, then the content of another sentence (at that context) must also be true at that state of the world: *logical consequence*. If a certain pattern of constructing a series of sentences  $\langle S_1, S_2, \dots, S_n \rangle$  always yields sequences of contents  $\langle C_1, C_2, \dots, C_n \rangle$  such that  $C_1, C_2, \dots$  is true at a state of the world, then  $C_n$  must also be true at that state of the world, then we say that that pattern is a *valid argument form*.

As a matter of historical fact, such compositional semantic theories have been first articulated for formal languages; and in so far as theorists have been interested in natural languages, they have been interested primarily in fact-stating discourse, especially in scientific discourse (consider, e.g. Russell's and Frege's motivations). Given this historical background, it is only natural, that the notions of content they used were classic: whether it is correct to believe such a content depends only on the state of the world (not on who believes or when). Suppose, now, that we regard these theories as successful, at least within their narrow range of application.

Now consider the following argument:

F1 The compositional and logical patterns of behaviour of paradigmatic beliefs and their linguistic expressions are correctly modelled by assigning them classic contents.

F2 Non-cognitive mental states do not have classic contents.

F3 The compositional and logical patterns of behaviour of moral judgements and their linguistic expressions are the same as those of paradigmatic beliefs and their linguistic expressions.

C1 So, moral judgements are not non-cognitive mental states.

This argument is not yet deductively valid: it ignores the possibility of a non-classic explanation for the compositional and logical patterns of behaviour of moral judgements and their expressions (despite the similarity of the explananda). Something like this gap in the argument was recognized by Susan Hurley (1984) and she articulated an additional constraint: ordinary modus ponens arguments and moral modus ponens arguments clearly have the same valid form, they are valid for the same reason. So it is upon the expressivist to ensure that her



explanation preserves this uniformity of explanation. Adapting this lesson to our purposes, we might add a further premiss in Hurley's spirit:

F4 The compositional and logical patterns of behaviour that beliefs and moral judgements and their linguistic expressions share, must receive a uniform explanation.

The argument from premisses F1–F4 to conclusion C1 does look tempting: we already have a correct theory for descriptive beliefs and their expression (F1), but the shared behaviour of moral judgements and descriptive beliefs should receive a uniform explanation (F3 and F4), however, the correct theory of beliefs cannot be used for non-cognitive states (F2). So, since we already have the correct treatment for descriptive beliefs, a treatment that does not work for non-cognitive states, we can't say that moral judgements are non-cognitive states (not without violating F4).

But there is a mistake: even if the classic account mentioned in F1 is correct, it may still be a special case of a more general account, an account which could offer a uniform explanation for the shared pattern of behaviour but still be compatible with moral judgements not being cognitive.

To see this, focus on a further silent premiss, which is needed to establish C1:

F5 There is no uniform explanation of the shared behaviour that is compatible with non-cognitivism (and with the correctness of the classic explanation).

This silent premiss is clearly unwarranted, and it can be shown to be false by looking at part 1 above. There are uniform non-classic explanations of the relevant patterns of behaviour of both cognitive and non-cognitive representations. Classic contents are special cases in the perspectival content framework, and the explanations needed for the shared patterns of behaviour do not rely on anything not true also of these special cases.

For concrete illustration, suppose the classic theory mentioned in F1 operated with classic contents that are functions from worlds to truth-values, and that the definition of logical consequence used to explain cases of valid inference was this: A content  $C_n$  is a consequence of contents  $C_1, C_2, \dots$ , =<sub>def</sub> for all worlds  $w$ : if the  $C_1, C_2, \dots$  all have the value true at  $w$ , then  $C_n$  also has the value true at  $w$ .

Suppose further that we now generalize to a perspectival conception on which contents are functions from centred worlds to truth-values. Then each of the earlier classic contents has a unique objective counterpart in our new framework (see §1.4 above): for each original classic content  $C$ , its perspectival counterpart  $C^*$  fulfils the following condition: for all worlds  $w$  and centres  $c$ ,  $C^*(\langle w, c \rangle) = \text{truth}$  iff  $C(w) = \text{truth}$ .

We can then adopt the following generalized definition of logical consequence: A content  $C_n$  is a consequence of contents  $C_1, C_2, \dots$ , =<sub>def</sub> for all centred worlds  $\langle w, c \rangle$ : if the  $C_1, C_2, \dots$  all have the value true at  $\langle w, c \rangle$ , then  $C_n$  also has the value true at  $\langle w, c \rangle$ .

Clearly, any explanations of cases of logical consequence in the classic framework can be translated into equally good explanations in the extended

perspectival framework. But the extended framework can accommodate beliefs and utterances with non-objective contents. The explanations offered for behaviour shared between moral judgements and descriptive beliefs will be uniform.

It is easy to overlook the silent and false assumption F5, if one forgets that nothing compels us to use a classic conception of representational content.<sup>16</sup>

## 2.8. Schroeder's fallacy

Mark Schroeder has done a lot to improve our understanding of the Frege-Geach point: as already mentioned, he has moved the focus away from the less promising (but once prominent) expressivist approaches that start from E1. Moreover, he has articulated a proposal that implements this advice by separating a general non-cognitive acceptance attitude 'being for' from the specific contribution made by 'bad' and other moral predicates to the contents that can be the objects of the general attitude of being for (see Schroeder 2008c. Compare also Silk 2014). However, Schroeder argues that his own solution to the Frege-Geach problem does not ultimately work. As we shall see, his reasoning to this conclusion also contains a fallacy that is quite tempting to those brought up on the classic view.

Schroeder argues that his own solution works only if we consider purely normative (moral) thought and language by itself. But in fact, we form complex thoughts (and sentences) involving both moral elements and descriptive ones. We can think, for example, that either Ed didn't steal the money, or he ought to be punished; or that if carbon dioxide emissions cause global warming, then emitting carbon dioxide is bad.<sup>17</sup> The following two sentences seem to serve for expressing such thoughts:

(M1) 'Either Ed didn't steal the money, or he ought to be punished.'

(M2) 'If CO<sub>2</sub> emissions cause global warming, then emitting CO<sub>2</sub> is bad.'

Dealing with such mixed thoughts and claims, he thinks, would require that the expressivist 'allow that all sentences express the same general kind of attitude' (2008a, 597). I concur. But Schroeder continues as follows:

For cognitivists, this is easy to do, for the attitude is simply belief. But for expressivists, this turns out to be tricky. For descriptive sentences, according to expressivists, express beliefs. So if they are to express states of being for, then it must turn out that beliefs are really states of being for – that belief needs to be analyzed in terms of a non-cognitive attitude. (2008a, 597. See also 2008c, 91–92, for a very similar passage.)

Schroeder is arguing that the following two options are exhaustive:

I all declarative sentences express belief.

II all declarative sentences express 'being for' (Schroeder's general non-cognitive attitude).

Option I is not acceptable to the expressivist. Option II seems radical and surprising, as Schroeder himself observes. It would involve, for example, as Schroeder is proposing, that the descriptive sentence 'Grass is green.' serves to express the non-cognitive attitude of 'being for' directed at the content of 'proceeding as if grass is green.'<sup>18</sup> However Schroeder spends the remainder of his book exploring how the expressivist could best defend this view, i.e. that all (declarative) sentences express non-cognitive states of being for, and that analogously beliefs are also really non-cognitive states. The final verdict he reaches, after detailed discussion, is negative: ultimately expressivism fails.

If I have done my job well, readers of this paper will have easily spotted Schroeder's fallacy. Of course there is a third option that Schroeder has not considered:

III All sentences express a general attitude of which belief and non-cognitive attitudes are species.

There is a general kind of attitude – call it 'judgement', for example – which falls into two mutually exclusive subspecies: cognitive and non-cognitive judgements. The answer is simple for those who have stepped back a little bit from the classic view, and who have present in their minds the notion of a perspectival content, which is a generalization of that of a classic content. The next and final section will flesh out this view.

I believe that this illustrates once more that the decision to adopt a classic notion of content, while not pernicious in itself, can make it hard to see the uniformity of a certain class of mental states: judgements.<sup>19</sup> Even though the class of judgements contains both beliefs and non-cognitive states, it is nevertheless importantly uniform in just the way needed: all judgements are states with perspectival representational contents.

## ***2.9. The positive account: judgements as states with perspectival contents***

I have claimed that non-cognitivists can perfectly coherently regard the class of judgements as uniform in all respects besides those that make up the cognitive/non-cognitive divide on which their defining thesis is based (see §2.2 above). What, then, might be those respects of uniformity amongst all judgements? Let me sketch a reasonable view (but one that does depart from some central elements of the Humean picture of Human psychology).

All judgements differ from other mental states or mental dispositions in that they involve the employment of concepts and that they are subject to rational norms. Another way of saying this, of course, is to say that they have contents, for contents serve, amongst other things, to model the conceptual and rational relations amongst different judgements (or other contentful mental states). This differentiates judgements from certain non-contentful, non-rational mental

states, such as certain moods or feelings, for example the feeling of giddiness or the feeling of anger (at least on a non-cognitive contrual of those feelings).

Judgements differ from other conceptual and contentful states in the kind of rational norm that governs them. Judgements are governed by norms of correctness: their contents determine what conditions need to hold for a judgement with that content to be correct. Thus, a judgement to the effect that the water is boiling is correct iff the contextually relevant water is indeed boiling. My judgement that the taste of the coffee is pleasant will be correct iff I will normally be pleased by the contextually relevant coffee. This norm of correctness, in turn, gives rise to general norms of rationality: the requirement to avoid judgements with inconsistent contents and the requirement to avoid judgements not obtained with correctness-conducive methods.

Desires or hopes may also have content, but in so far as they are subject to norms of rationality or correctness, these norms play a different role. For example, suppose I have an instrumental desire to acquire Dr Gordon's Hair Elixir, which is based on the belief that applying Dr Gordon's Hair Elixir to one's scalp will boost hair growth and the more fundamental desire to have fuller hair. Then my instrumental desire for the elixir might be incorrect or irrational because my belief in the elixir's effectiveness, on which the desire is partly based, might be incorrect or irrational. Thus, while other contentful mental states may be subject to certain derivative constraints of correctness or rationality, only judgements are non-derivatively subject to such constraints.<sup>20</sup>

This story, while reasonable, diverges from Hume's at a key point: Hume thought that only beliefs, or the mental states produced by the faculty of reason, were subject to (underived) rational norms, while the passions, i.e. non-cognitive states, were not subject to underived rational constraints at all. Thus, in Hume's picture, the division between the intrinsically motivating states and motivationally inert cognitive states coincides with the division between the mental states that are subject to underived rational norms and those that are not. My 'reasonable view' thus emphatically allows for rationally constrained judgements that are nevertheless on the motivational side of Hume's divide.<sup>21</sup>

I believe that all these general features of judgements can be modelled perfectly within a perspectival content framework. Let us begin with the distinction between beliefs and non-cognitive attitudes, that is so dear to the non-cognitivist's heart. The most obvious proposal is to distinguish those perspectival contents that comply with the classic constraint: i.e. if accepting such a content is correct for one thinker at one time, then it is correct to accept it for any thinker at any time. Call this subclass of perspectival contents 'objective'. Then one good starting point would be to say that judgements with objective contents are beliefs, while non-cognitive judgements have non-objective contents.

Why should some of the judgements with non-objective contents be intrinsically motivating, as is required by the psychological motivation for non-cognitivism (§2.2)? This is because the correctness of accepting a non-objective

content will depend not just on the aspects of reality represented, but also on features of the judge. In the case of judgements of pleasantness, this is fairly straightforward: in order for my judgement that the taste of that coffee is pleasant to be correct, the coffee needs to be such that my dispositions are such that under certain normal conditions I experience pleasure when drinking the coffee. In so far as pleasure is an end in itself, such a judgement will constitute a motive, *ceteris paribus*, to prefer the coffee in question to other options that do not produce pleasure.

Hume's thought that motivational states are rationally unconstrained can be partially preserved: those features of thinkers at times on which correctness may depend when contents are non-objective will often be features that are not subject to thinkers' rational control. In this sense, the reasonable picture can retain a non-rational aspect of intrinsically motivating judgements.

The moral case is more complicated: what are to be the requirements on the judge if she is to be correct in her judgement that gambling is bad? There are many options. Some may base moral judgements on certain general individual preferences (such as a desire to promote pleasure, happiness, justice or what have you). Others may have a more social view of moral norms and tie the individual's interests somehow to the values prevalent in that individual's society. This is not an issue to be decided here: all we need to show is that there are plausible proposals on each of the outlooks that might combine with a non-cognitivist view (like the ones just mentioned). Perhaps what is required in an individual, in order for that individual's judgement that gambling is bad to be correct, is that gambling have a tendency to prevent the realization of that person's moral goals. Then that would explain why judging gambling to be bad is *ipso facto* a reason to avoid gambling. For if the judgement is correct, then the avoidance of gambling will avoid something that prevents the realization of that judge's moral goals.

Why should beliefs, i.e. judgements with objective contents, be epistemologically less problematic than judgements with non-objective contents? The epistemological problem that Ayer saw with an empirical account of the justification for moral claims, construed as irreducible was this: if moral judgements were based on 'intellectual intuition', then we would get conflicting intuitions and have no way of settling the conflict (1946, 108–109). Any adequate method for justifying judgements with objective content will not suffer from this defect: since adequate justification tracks correctness, correctness for one is the same as correctness for anyone, anything that is an adequate method for me is also an adequate method for you – as long as the content is objective. Thus while that does not guarantee that we have an effective method for every objective question, it at least guarantees that once we have exchanged our justifications and agreed on what are adequate methods, we will share the same view. In the case of non-objective contents, this is not guaranteed. Here, what tracks correctness for you may not track it for me. Thus the reasonable view makes

room for the kind of epistemological asymmetry that traditionally motivates non-cognitivism.

Finally, the metaphysical motivation for non-cognitivism (given non-reductivism): why should correct judgements with non-objective contents be less metaphysically committing than judgements with objective contents? Assume, minimally, that reality consists in those aspects of the world that we all share and to which all our judgements are answerable. Then it seems clear that those judgements that have objective contents, if correct, represent those aspects of reality that all thinkers are answerable to, while judgements with non-objective contents, even if correct, may not describe an aspect of reality that everyone faces or is answerable to.<sup>22</sup>

I hope that this concluding sketch has put some flesh on the bones of the abstract idea of a mental state that is defined as a state with a perspectival content – a judgement –, and on how it can play the role needed by the non-cognitivist expressivist route I have recommended.<sup>23</sup>

## Notes

1. For example Frege (1918), Wright (1992, 91–93, 2005, 170), and Moore (1997) imposes the same constraint on contents, but construes representation more liberally, allowing that a representation could be what he calls ‘perspectival’.
2. Treating B5 and B9 as having the same representational content has its problems, of course. But I do not want to enter this discussion here. My main purpose here is to characterize the classic and the perspectival conceptions of representational content as alternatives. The second part of the paper will then identify some substantial conclusions that may be derived from the choice between the approaches. I have offered arguments in favour of the perspectival approach to content in many places, including Kölbel (2013, 2014, 2015a).
3. And arrive, respectively, either at a view similar to Harman’s (1975), Dreier’s (1990) or Björnsson and Finlay’s (2010) indexical relativism or at one similar to the invariant moral relativism outlined in Kölbel (2015b).
4. Wanting to distinguish the contents of B3 and B11, as well as treating B9 and B10 as having the same content, correspond to the classic motivations for introducing perspectival contents (also often called ‘*de se* contents’). Two loci classici are Perry (1979) and Lewis (1979).
5. For example: while B3 and B8 attribute possession of a cat to different people, the people in question are the believer’s neighbours in each case.
6. See Egan (2007).
7. See Kölbel (2013).
8. The primary role of representational contents is that they must determine necessary and sufficient conditions for when it is correct to have a representation with a given content. Classic contents and their centred counterparts determine the same such conditions.
9. See Dreier (1999) for a nice recipe for transforming unstructured perspectival contents into structured ones.
10. I will not be assuming that every non-cognitivist must have all these motivations at once. All I am saying is that typically non-cognitivists have one of these motivations. There are also some important metaphysical motivations that I

have not mentioned, such as the argument that if there were moral facts and properties, it would be hard to explain why they should supervene on non-moral facts and properties (this argument, again, is directed against non-reductivist forms of cognitivism only).

11. In order to accept this much, one does not need, for example, to follow the programme of Green and Bar-On, who claim that linguistic meaning is in general based on such expressive properties (see, e.g. Bar-On 2013; Green 2007).
12. See Kölbel (2013) for some proposals for the dynamics of conversation when assertion is not construed as aiming at objective truth and the contents of assertion are not objective (not 'portable').
13. See Blackburn (1988); who postulates an attitude of toleration in response to Hale (1986, 1993); who objects. See also Schroeder (2008a, 2008b, 2010) for a good account of how Unwin's challenge generalizes to other forms of embedding.
14. I.e. not, e.g. when used in the slang sense, in which it predicates something like attractiveness.
15. Even though Kölbel (1997) distinguishes the 'force-indicator approach', which keeps E1, from the 'content-indicator approach', which drops E1 and pursues the solution here proposed. It emerges that the latter is more elegant and that the former, when executed viably, leads to something structurally similar to the latter. Still, the focus is on the former, thus exemplifying the prevailing tendency to regard force-indicator approaches as the more typical forms of expressivism. Similarly, Kölbel (2002), ch. 6.8, explicitly classifies Gibbard (1990) as not genuinely pursuing an expressivist agenda, precisely because of disanalogies with the E1 approach.
16. I suggested in part 1 that the choice between classic and perspectival conceptions is a pragmatic decision that can have no substantial consequences. So how would one deal with this problem on a classic conception? To start with, presumably, one would have to deny F2 and find some suitable classic contents even for non-cognitive states.
17. I am assuming that the question of whether global warming is caused by CO<sub>2</sub> emissions is a descriptive one. Those who accept a non-cognitivist account of causation should choose a different example.
18. I believe that Schroeder saddles the expressivist with an (additional) unnecessary problem at this point: the construction he chooses for descriptive declarative sentences  $F(a)$  as expressing being for proceeding as if  $F(a)$  is not symmetrical to the construction he chooses in cases where  $F$  is a moral predicate: here he will say that the sentence expresses being for  $F$ -ing  $a$  (where  $F$ -ing = blaming for if  $F$  = 'bad'). The asymmetry at the level of attitudes expresses is mirrored at the level of the proposed expressivist logical forms for descriptive and moral sentences. I believe that many of the problems he subsequently detects for the expressivist are the result of this unfortunate and unmotivated choice.
19. I myself have been, to some extent, a victim of this effect: In Kölbel (1997), I considered roughly the expressivist theory here proposed (under the label 'content-indicator approach'). I suggested then that expressivists pursuing this agenda would need to develop a more general novel theory of content, and that that theory would probably need to be an inferential role theory, i.e. a theory that does not rely on compositionally specified truth-conditions. I now know better: as we saw in the first part of this paper the theory of perspectival contents can perfectly well take one of the habitual truth-conditions-specifying forms: it just needs to add a further parameter of variation of truth-value, such as a location or centre.



20. Many readers will ask (and will have asked throughout) how my *correctness* is related to *truth*. I can't treat this question adequately here, so let me just gesture in the direction of a possible answer: truth is a special case of correctness, namely the correctness of a representation with an objective content. This would work well with a view of truth as objective truth, on which a non-cognitivist would then be entitled to deny that moral judgements are truth-apt. There also seems to be a (minimal) truth concept that we apply across the board to objective and non-objective contents. Moreover, there also seem to be truth-like concepts that we employ in formal semantics. These different concepts of truth need to be distinguished. For more on this, see Kölbel (2008a, 2008b).
21. I am aware that it may seem terminologically odd to call certain contentful judgements 'non-cognitive'. For often the label 'cognitive' is used precisely to mark out the features of contentfulness and being subject to rational constraints. I am nevertheless sticking to the term because the position discussed here is non-cognitivism. That is why in §2.2 above I spent some time investigating the significance of the non-cognitivist thesis by reviewing the typical motivations for non-cognitivism.
22. I believe that this 'reasonable view' is broadly in line with Gibbard's view articulated in *Wise Choices, Apt Feelings* (1990).
23. I am very grateful to the editors, Matt Bedke and Stefan Sciaraffa, of this special issue for their very useful critical feedback, which has helped improve this paper. Thanks also to a very critical audience at the 2017 *Representation & Evaluation* conference in Vancouver, including especially Matthew Chrisman, Jamie Dreier, Anandi Hattiangadi, Gurpreet Rattan, Stefan Sciaraffa, Mark Timmons. Thanks also to audiences at Barcelona and Salzburg, where versions of this paper were presented and discussed in June and November 2017.

## Disclosure statement

No potential conflict of interest was reported by the author.

## Notes on contributor

**Max Kölbel** is Professor of Analytic Philosophy and Philosophy of Language at the University of Vienna and ICREA Research Professor at Institució Catalana de Recerca i Estudis Avançats. He works in the philosophy of language, epistemology, metaphysics and metaethics.

## References

- Ayer, Alfred J. 1946. *Language, Truth and Logic*. 2nd ed (1st ed 1936). London: Victor Gollancz.
- Bar-On, Dorit. 2013. "Origins of Meaning: Must We Go Gricean?" *Mind & Language* 28: 342–375.
- Björnsson, Gunnar, and Stephen Finlay. 2010. "Metaethical Contextualism Defended." *Ethics* 121: 7–36.
- Blackburn, Simon. 1984. *Spreading the Word*. Oxford: Clarendon Press.
- Blackburn, Simon. 1988. "Attitudes and Contents." *Ethics* 98: 501–517.



- Dreier, Jamie. 1999. "Transforming Expressivism." *Noûs* 33: 558–572.
- Dreier, James. 1990. "Internalism and Speaker Relativism." *Ethics* 101: 6–26.
- Egan, Andy. 2007. "Epistemic Modals, Relativism and Assertion." *Philosophical Studies* 133: 1–22.
- Frege, Gottlob. 1918. "Der Gedanke. Eine Logische Untersuchung." *Beiträge zur Philosophie des deutschen Idealismus* 1: 58–77. English translation as "The Thought: A Logical Inquiry", *Mind* 65 (1956), pp. 287–311.
- Geach, Peter. 1960. "Ascriptivism." *Philosophical Review* 69: 221–225. Reprinted in his *Logic Matters* (1972), pp. 250–54, Oxford: Blackwell.
- Geach, Peter. 1965. "Assertion." *Philosophical Review* 74: 449–465. Reprinted in his *Logic Matters* (1972), pp. 254–69, Oxford: Blackwell.
- Gibbard, Allan. 1990. *Wise Choices, Apt Feelings*. Oxford: Oxford University Press.
- Gibbard, Allan. 2003. *Thinking How to Live*. Boston: Harvard University Press.
- Green, Mitchell. 2007. *Self-Expression*. Oxford: Oxford University Press.
- Hale, Bob. 1986. "'The Compleat Projectivist' (Critical Study of Blackburn 1984)." *Philosophical Quarterly* 36: 65–84.
- Hale, Bob. 1993. "Can There Be a Logic of Attitudes." In *Reality, Representation and Projection*, edited by Crispin Wright and John Haldane, 337–363. Oxford: Oxford University Press.
- Hare, Richard M. 1952. *The Language of Morals*. Oxford: Oxford University Press.
- Hare, Richard M. 1970. "Meaning and Speech Acts." *Philosophical Review* 79: 3–24.
- Harman, Gilbert. 1975. "Moral Relativism Defended." *The Philosophical Review* 84: 3–22.
- Hurley, Susan. 1984. "Frege, the Proliferation of Force, and Non-Cognitivism." *Mind* 93: 570–576.
- Kölbel, Max. 1997. "Expressivism and the Syntactic Uniformity of Declarative Sentences." *Crítica* 29: 3–51.
- Kölbel, Max. 2002. *Truth without Objectivity*. London: Routledge.
- Kölbel, Max. 2008a. "'True' as Ambiguous." *Philosophy and Phenomenological Research* 77: 359–384.
- Kölbel, Max. 2008b. "Truth in Semantics." *Midwest Studies in Philosophy* 32: 242–257.
- Kölbel, Max. 2013. "The Conversational Role of Centered Contents." *Inquiry* 56: 97–121.
- Kölbel, Max. 2014. "Agreement and Communication." *Erkenntnis* 79: 101–120.
- Kölbel, Max. 2015a. "Relativism 1: Representational Content." *Philosophy Compass* 10 (1): 38–51.
- Kölbel, Max. 2015b. "Moral Relativism." In *Routledge Encyclopedia of Philosophy*, Taylor and Francis. doi:10.4324/9780415249126-L099-2. Accessed 29 January 2018. <https://www.rep.routledge.com/articles/thematic/moral-relativism/v-2>
- Lewis, David. 1979. "Attitudes *De Dicto* and *De Se*." *Philosophical Review* 88: 513–543. Reprinted in David Lewis, *Philosophical Papers* vol. 1, Oxford: Oxford University Press, pp. 133–59.
- Moore, Adrian. 1997. *Points of View*. Oxford: Oxford University Press.
- Ogden, C. K., and I. A. Richards. 1946. *The Meaning of Meaning*. 8th ed. (1st ed. 1923). New York, NY: Harcourt, Brace & World.
- Perry, John. 1979. "The Essential Indexical." *Noûs* 13: 3–21.
- Schroeder, Mark. 2008a. "How Expressivists Can and Should Solve Their Problem with Negation." *Noûs* 42: 573–599.
- Schroeder, Mark. 2008b. "What is the Frege–Geach Problem." *Philosophy Compass* 3 (4): 703–720.
- Schroeder, Mark. 2008c. *Being for: Evaluating the Semantic Program of Expressivism*. Oxford: Oxford University Press.
- Schroeder, Mark. 2010. *Noncognitivism in Ethics*. London: Routledge.

- Silk, Alex. 2014. "How to Be an Ethical Expressivist." *Philosophy and Phenomenological Research* 91: 47–81.
- Stevenson, Charles L. 1937. "The Emotive Meaning of Ethical Terms." *Mind* 46: 14–31.
- Unwin, Nicholas. 1999. "Quasi-Realism, Negation, and the Frege-Geach Problem." *Philosophical Quarterly* 50: 337–353.
- Wright, Crispin. 1992. *Truth and Objectivity*. Cambridge, MA: Harvard University Press.
- Wright, Crispin. 2005. "Relativism about Truth Itself: Haphazard Thoughts about the Very Idea." In *Relative Truth*, edited by Max Kölbel and Manuel García-Carpintero, 157–185. Oxford: Oxford University Press.