

Thick Concepts, Non-Cognitivism, and Wittgenstein's Rule-Following Considerations¹

Adam M. Croom

Department of Philosophy
University of Pennsylvania
433 Cohen Hall, Philadelphia PA 19104
E-mail: adam.m.croom@gmail.com

Abstract:

Non-cognitivists claim that thick concepts can be disentangled into distinct descriptive and evaluative components and that since thick concepts have descriptive shape they can be mastered independently of evaluation. In *Non-Cognitivism and Rule-Following*, John McDowell uses Wittgenstein's rule-following considerations to show that such a non-cognitivist view is untenable. In this paper I do several things. I describe the non-cognitivist position in its various forms and explain its driving motivations. I then explain McDowell's argument against non-cognitivism and the Wittgensteinian considerations upon which it relies, because this has been sufficiently misunderstood by critics and rarely articulated by commentators. After clarifying McDowell's argument against non-cognitivism, I extend the analysis to show that commentators of McDowell have failed to appreciate his argument and that critical responses have been weak. I argue against three challenges posed to McDowell, and show that the case of thick concepts should lead us to reject non-cognitivism.

1. Cognitivist and Non-Cognitivist Claims

Cognitivists claim that moral statements express beliefs that are truth-evaluable. Cognitivists are normally moral realists, but they needn't be. For instance, the error theorist John Mackie (1977) claims that, although moral statements express truth-evaluable beliefs, moral statements are all systematically false because the world as it genuinely is does not contain moral features. So Mackie is a cognitivist yet a moral anti-realist. However, cognitivism is most attractive to moral realists who claim that moral statements express truth-evaluable beliefs and that at least some of these beliefs are true by virtue of representing moral features that genuinely obtain in the world. In contrast, non-cognitivists claim that there are no moral features that are genuine. Further, they claim that moral statements do not express truth-evaluable beliefs, but rather work primarily to express something non-cognitive, such as attitudes of approval or disapproval, or desire.

¹ This paper was awarded the University of Pennsylvania's Elizabeth F. Flower Prize for best philosophical paper of the year (2009-2010). Special thanks go out to Elisabeth Camp for helpful comments on earlier drafts of this paper.

2. Non-Cognitivist Accounts and Thick Concepts

Let's review some non-cognitivist accounts of moral statements, starting with A.J. Ayer's emotivism. Emotivism holds that 'in so far as statements of value... are not scientific, they are not in the literal sense significant, but are simply expressions of emotion which can be neither true or false' (Ayer 1952: 102-103). Since Ayer believes that only naturalistic scientific statements represent the world as it genuinely is, and since moral statements are not naturalistic scientific statements, Ayer concludes that moral statements are not 'significant propositions' suitable for truth-value. However, as Ayer himself understood, this entails that genuine disagreement on moral issues is impossible. If moral concepts are only 'pseudo-concepts' as the emotivist claims, then the propositions in which they occur are not genuine propositions at all. And if moral statements are not genuine propositions but mere emotive ejaculations, moral statements are not truth-evaluable, and resultantly, unsuitable for genuine agreement or disagreement (Ayer 1952: 107-108).

Most non-cognitivists have been unwilling to accept that disagreement over moral statements is impossible. Indeed, accounting for moral disagreement is one thing we expect from a successful account of moral statements (Stevenson 1937: 16). Resultantly, one classical non-cognitivist move has been to suggest that the content of a moral concept is not *exclusively* emotive or evaluative; moral concepts also contain sufficient descriptive content such that we can genuinely disagree over the propositions in which they occur. R.M. Hare, for instance, suggests that the term 'good' is composed of both an evaluative and descriptive meaning (1970: 118-119). Hare further claims that:

although with 'good' the evaluative meaning is primary, there are other words in which the evaluative meaning is secondary to the descriptive. Such words are [e.g.] 'tidy' and 'industrious' (1970: 121).

The concepts 'tidy' and 'industrious', which Hare points out, are called 'thick concepts'. Thick concepts are characterized as 'hold[ing] together a property and an attitude... or, as it is also sometimes put, description and evaluation' (Dancy 1996: 263). Typical examples include: 'shameful', 'lewd' (Gibbard 1992: 278-279), 'treacherous', 'rude', 'cruel', 'honest' (Blomberg 2007: 63), 'courage', 'delicate' and aesthetic concepts (Burton 1992: 28-29) and concepts of moral virtues (McDowell 1981: 144) and vices (Dancy 1996: 263). The thick moral concept 'courageous', for instance, is typically conceived as including both description and evaluation. 'Courageous' is descriptive in that it's normally applied to the same types of descriptive items, e.g. to acts of strength and determination in the face of conflict,² and it's by virtue of this descriptive component that thick moral concepts are considered *genuine* concepts. Yet 'courageous' is evaluative in that it's normally indicative of an evaluative pro-attitude, and it's by virtue of this evaluative component that thick moral concepts are considered *moral* concepts. That is, if thick moral concepts lacked description then according to the non-cognitivist we could not *genuinely disagree* over the moral statements in which they occur, and if thick moral concepts lacked evaluation there would be no

2 What's important here is the point that 'courageous' is descriptive in that it normally applies to the same types of descriptive items; the particular descriptive example I chose – e.g. strength and determination in the face of conflict – is inessential to my general point and it may be substituted for another description that the reader finds more apt. But the point here remains that 'courageous' is descriptive in that it normally applies to *those* same types of descriptive items.

moral statements over which to disagree. So the thick moral concept ‘courageous’ presumably involves both description and evaluation, e.g., a pro-attitude towards acts of strength and determination in the face of conflict.

3. Non-Cognitivist Motivations

We saw that Ayer’s early non-cognitivism was motivated by a naturalistic conception of what counts as a ‘significant proposition’. Since Ayer believed that moral statements are not naturalistic scientific statements and that only the latter genuinely represent the world, Ayer concluded that moral statements are not significant truth-apt propositions. It’s because Ayer assumes that there are no moral features in the world that he denies that moral statements are propositions representing how the world is.

It’s unsurprising, then, that non-cognitivism is typically motivated by a naturalism that takes for granted what Bernard Williams calls the ‘absolute conception’ of the world. As Williams explains, this is a conception of the world ‘consisting of nonperspectival materials available to any adequate investigator, of whatever constitution’ (1985: 139-140) which serves to distinguish ‘the world as it is independent of our experience’ from ‘the world as it seems to us’ (1985: 139-140). The appeal of the absolute conception is presumably due to its promise to explain ‘how the more local representations of the world [as it appears to us] can come about’ and how ‘to relate them to each other, and to the world as it is independently of them’ (Williams 1978: 245-246).

In order to articulate ‘the world as it is independent of our experience’ in a way that is distinct from our articulation of ‘the world as it seems to us’, we’ll need to utilize two distinct types of concepts. ‘Response-dependent concepts’, as Mark Johnston defines them, are ‘those concepts which exhibit a conceptual dependence or interdependence with concepts of our responses’ (1989: 145). As concepts that are conceptually interdependent with concepts of our responses, response-dependent concepts are those suitable to articulate the features of ‘the world as it seems to us’, e.g. concepts of color and other secondary-quality features, evaluative features, etc. (Johnston 1989: 146). In contrast, ‘response-independent concepts’ are those that are not dependent or interdependent in this way (Johnston 1989: 146). As concepts that are conceptually independent of concepts of our responses, response-independent concepts are those suitable to articulate features of ‘the world as it is independent of our experience’, e.g. concepts of science, mathematics, etc. Indeed, it seems that the best ‘place to begin... explicating the primary/secondary quality distinction is with the response-independent, response-dependent distinction’ (Johnston 1989: 147).

Now, someone accepting the absolute conception of the world, as McDowell points out, may come to distrust the genuineness of secondary-quality features. Assuming that genuine features must be represented in concepts of naturalistic primary-qualities, and that our experiences of secondary-qualities represent the world as containing features other than those of primary-quality and so must be further explained down into primary-quality concepts, one might resultantly believe that an experience of secondary-quality ‘fails to be a transparent mode of access to something that is [genuinely] there [in the world] anyway’, independent of us (McDowell 1981: 142). This in turn suggests that by maximally purifying from our conception of the world those secondary-quality features the concepts of which are response-dependent, our inquiries should ideally ‘converge on an answer, where the best explanation of the convergence

involves the idea that the answer represents how things [genuinely] are' (McDowell 1981: 136).

However, it's implausible that the convergence of judgments found in the area of moral discourse will replicate that of the natural sciences (Williams 1985: 136). One might be inclined to reject moral features as genuine as one had rejected secondary-quality features. For instance, secondary-quality features such as color are dependent upon our subjective constitution in that an investigator of very different constitution may un-illusorily encounter a world completely lacking in color features. And as concepts of color are interdependent with concepts of our responses, so in a similar way concepts of morality are interdependent with concepts of our responses insofar as they concern how we ought to respond with our actions. Moreover, as McDowell argues, moral features can be seen as dependent upon us because it's by virtue of 'our affective and attitudinative propensities... that we can be brought to care in appropriate ways about the things we learn to see as collected together by the classifications' of 'courageous', etc. (1981: 142). So one might argue that since our best naturalistic conception of the world is seen as a result of purifying secondary-quality features from it, on similar grounds we must reject moral features. On this view, moral features have no place in our account of the world as it genuinely or objectively is (McDowell 1981: 154).

The absolute conception is considered an 'objective' account of the world precisely because it factors out features the concepts of which are response-dependent, leaving in its account only those features that are considered fundamentally primary by virtue of describing the world as it is independently of us. On this view, only those features that can be described in response-independent concepts will be considered genuine, because those are the only features that are represented as obtaining in the world as it is anyway, independently of us. This view further suggests that features that are described in response-dependent concepts are not genuine, since these features aren't represented as obtaining in the world as it is anyway. Resultantly, the absolute conception in its purest form denies the genuineness of secondary-quality and evaluative features. Accordingly, if beliefs are to track genuine features, they are to track those features articulated in descriptive naturalistic concepts, not moral or evaluative ones. And it is precisely the acceptance of such a view that leads non-cognitivists to assume that there are no moral features in the world for beliefs to truth-aptly track.

4. Rejecting the Disentangling Manoeuvre

As McDowell explains, non-cognitivism holds that:

when we ascribe value to something, what is actually happening can be disentangled into two components. Competence with an evaluative concept involves, first, a sensitivity to an aspect of the world as it really is... and second, a propensity to a certain [non-cognitive] attitude... from which items in the world seem to be endowed with the value in question... [so] in making value judgments, [an agent] register[s] the presence in objects of some property they authentically have, but enrich their conception of this property with the reflection of an attitude (1981: 143-144).

The non-cognitivist Alan Gibbard provides an example of how an 'evaluative concept' might be applied, which will be helpful here. Gibbard, in describing how a foreign

tribe – the Kumi – apply their thick moral concept ‘gopa’, explains their application procedure as consisting of two stages:

[first] they observe an act they *know* to be the killing of an outgroup member in the face of danger. They conclude ‘This act is *gopa*’. Then they further conclude, ‘Let us glory in this act!’ (1992: 268).

According to Gibbard, the Kumi first cognitively track descriptive features (the killing of an out-group member in the face of danger) and then respond to these features with an evaluation (a positive evaluation such as ‘how glorious!’). As this view suggests, the descriptive content of the Kumi’s cognition is distinct from their evaluative non-cognitive pro-attitude since the latter is only a response to the former.

Non-cognitivists claim that it’s the descriptive content of the thick concept that determines its correct application, because it’s by virtue of its description alone that the concept is applied to the same type of descriptive things. In other words, non-cognitivists claim that the concept has descriptive *shape*; the concept is shaped by description. Non-cognitivists further claim that the evaluative component is distinct from the descriptive and so does not determine the concept’s shape (Williams 1985: 141). They deny that evaluation is itself ‘in some way percipient’ or capable of ‘expanding our sensitivity to how things are’ (McDowell 1981: 143). Rather, non-cognitivists think that the evaluation piggybacks on the description to express our non-cognitive attitude towards the action we are applying the thick concept to.

Now, McDowell denies that thick concepts can be disentangled into distinct descriptive and evaluative components. Here’s his central argument against the ‘disentangling manoeuvre’:

If the disentangling manoeuvre is always possible, that implies that the extension of the associated [thick] term [e.g. ‘courage’]... could be mastered independently of the special concerns which, in the [moral] community, would show themselves in admiration or emulation of actions seen as falling under the [thick] concept [of courage]. That is: one could know which actions the [thick] term [‘courage’] would be applied to... without even embarking on an attempt to make sense of their admiration. That [attempt to make sense of their admiration] would be an attempt to comprehend their perspective [from inside the moral community]; whereas, according to the [non-cognitivist] position I am considering, the genuine feature to which the [thick] term is applied should be graspable without benefit of understanding the special perspective [from inside. That is, it should be graspable from the outside, independent of the moral community’s special perspective] (McDowell 1981: 144).

McDowell denies the non-cognitivist claim that thick concepts can be disentangled into distinct descriptive and evaluative components because this ‘disentangling manoeuvre’ would require that one could master the extension of the thick concept independently of understanding why we evaluate the items falling under its extension as we do. That is, in order to master the thick concept ‘courageous’, the non-cognitivist claims that we needn’t understand the practice of moral evaluation at all. Rather, they claim that it’s sufficient to cognitively track purely descriptive features, e.g. strength and determination in the face of conflict. Presumably, evaluations are simply post-cognitive responses and don’t contribute to the shape of the thick concept, and so aren’t required for one to have mastery over its extension.

But McDowell argues that this non-cognitivist view is mistaken. In order to understand a thick concept such as ‘courageous’ and gain mastery over its extension, one must be able to understand why certain actions, and not others, are considered the ones apt for falling under the extension of that thick concept. In other words, mastery over the extension of a thick concept requires *evaluating* candidates for its extension and determining which ones are *apt*. For instance, not just *any* act of strength and determination in the face of conflict will count as courageous. One might face the conflict of trying to read peacefully in a room of crying babies (e.g. at a hospital) and so evince the determination to strongly slap each baby to shut them up. But there’s nothing courageous about that. Or consider a racist that evinces strength and determination in derogating racial groups that he despises. That needn’t be ‘courageous’ either. Clearly, then, only *certain* acts of strength and determination in the face of conflict are aptly ‘courageous’, and it’s an evaluative outlook that reveals that the baby-slapper and racist’s actions, although satisfying the description, are not *apt* for the concept ‘courageous’.³

Moreover, it’s no use for the non-cognitivist to suggest that we should include into the meaning of ‘courageous’ e.g. an *appropriate amount* of un-foolishness, un-derogatoriness etc. because to do this is simply to pack an *evaluative* requirement into the very definition of ‘courageous’. That is, an evaluative outlook is still required in order to determine what e.g. *counts as* an ‘appropriate amount’ of un-foolishness for an act to still count as a ‘courageous’ one. To suggest that we pack an evaluative requirement into the definition of a thick concept would be to allow evaluation to shape the thick concept, and this is something the non-cognitivist is strictly committed *against*. So such a suggestion remains unavailable to the non-cognitivist.

Clearly only certain descriptive features are apt for thick concepts such as ‘courageous’, and an understanding of why certain acts and not others are apt involves more than simply tracking some pre-evaluative set of descriptive features. In order to properly evaluate actions e.g. as courageous requires the development of a certain sort of sensitivity and sensibility: to share in a perspective that views actions as pro or con, as worthy of praise or worthy of contempt. Through our development and enculturation into a social practice of evaluating actions and appreciating certain features in those actions as salient, we learn to collect certain actions, but not others, together under the extension of particular moral concepts. As one grows into the moral community and acquires a morally sensitive outlook, one learns how to pick out which features aptly belong under the extension of a thick moral concept, and how to respond to the actions grouped together by this concept in certain ways. Mastery over moral concepts requires more than cognitively picking out natural descriptive features in the world, it requires a moral upbringing, moral eyes, and a moral point or purpose, and so a ‘person from another culture who failed to see the evaluative point of a thick [moral] con-

3 Consider also the thick aesthetic concept ‘delicate’. Non-cognitivists typically assume that in order to master this aesthetic concept it is sufficient that one cognitively track purely descriptive features, e.g. objects with the descriptive properties of being small, pale, and fragile. But as Stephan Burton has rightly argued, “not just any small size, pale colors, and fragility will do the trick. Only some small, pale, fragile things are delicate; the vast majority are merely bland” (1992: 30). In other words, whereas it’s typically apt to ascribe the aesthetic concept ‘delicate’ to flowers or to a young female lover, it would typically be considered *inapt* to ascribe that same term to, for instance, a small piece of rotting flesh that was nonetheless small, pale, and fragile. Thus, only *certain* small, pale, and fragile objects are aptly ‘delicate’, and it’s from within an evaluative framework that one comes to understand that the small, pale, and fragile piece of rotting flesh is not *apt* for the concept ‘delicate’.

cept would not be able to predict local use of it on the basis of descriptive similarities alone' (Dancy 1996: 263).

To recap: the non-cognitivist claimed that thick concepts contain distinct descriptive and evaluative components that can in principle be disentangled, and that since the latter component is simply a response to the former one, the extension of a thick moral concept can in principle be mastered by someone outside the viewpoint of the community of moral evaluators. Yet contrary to this claim, we saw that for thick concepts, evaluation partly determines description since we must rely on an evaluative outlook in order to pick out those features that aptly belong under the extension of the thick concept. For instance, only certain acts of strength and determination in the face of conflict are aptly 'courageous', and how these 'certain' acts are to be picked out is by reference to the outlook of the community of moral evaluators. Resultantly, one cannot master the extension of a thick concept independently of evaluation. And since the extension of thick moral concepts is determined by both description and evaluation, McDowell argues that the non-cognitivist is wrong to suppose that 'corresponding to any value concept, one can always isolate a genuine feature of the world... a feature that is there anyway, independently of anyone's value experience being as it is – to be that which competent users of the concept are to be regarded as responding when they use it' (1981: 144). Description, considered by the non-cognitivist as something distinct and separable from evaluation, is insufficient for those outside the community of moral evaluators to master the extension of thick moral concepts because thick concepts have both descriptive and evaluative shape. Therefore, contra non-cognitivism, evaluation is not something that can simply be 'disentangled' away from the thick concept.

5. The Objective Conception of Rules: Psychologically Grasping Independently Fixed Rules

It's commonly understood that in order to correctly apply a concept, moral or otherwise, to the same types of things, one must use the concept in accord with the rule for its application: for 'judgments or utterances to be intelligible as applications of a single concept to different objects, [they] must belong to a practice of going on doing the same thing' (McDowell 1981: 145). This 'going on to do the same thing' is typically conceived by the non-cognitivist as fixed by rules that determines a concept's correct application. Such rules are fixed 'independently of the responses and reactions a propensity to which one acquires when one learns the practice [involving the rule] itself' (McDowell 1981: 146) because, presumably, if rules were *not* independent of human responses then we would be unable to account for the *normativity* of rules, i.e. that one's responses can be *wrong*. By accounting for the fixedness or objectivity of rules on the grounds that rules are independent of human responses, the non-cognitivist assumes that rules can only be objectively fixed within the objective account of the world offered by the absolute conception. Accounting for the fixedness of rules outside of the absolute conception's framework strikes the non-cognitivist as simply implausible. As Gibbard says:

I think I get the idea of a non-objectivist model, where we see... judgments as a cultural artifact. But where in this model is there room for truth and falsehood? There is only a way of living (1992: 269).

However, an account of the fixedness of rules does not yet give us an account of how we grasp and act in accord with rules. So corresponding to their account of the fixedness of rules, non-cognitivists must further provide an account of how we grasp and act in accord with these rules. Unsurprisingly, the non-cognitivist explains our ability to grasp and act in accord with rules by positing some 'special psychic mechanism that ties discussion to action' (Gibbard 1992: 278). Presumably, it's because we possess the appropriate psychological machinery that we can grasp rules and act in accord with them. So the non-cognitivist offers a two-component account of rule-following involving (1) that there are rules fixed independently of human responses, and (2) that humans possess psychological machinery by which we can grasp these independently fixed rules and act in accord with them. If the non-cognitivist lacks (1), then there are no rules for humans to psychologically grasp and act in accord with. If the non-cognitivist lacks (2), then there are no means for human to grasp rules and act in accord with them.

So the non-cognitivist account of rule-following is successful just in case (1) and (2) are both accounted for. McDowell's strategy against the non-cognitivist account of rule-following then, is to use Wittgenstein's rule-following considerations to discredit (2), thereby rendering (1) suspect. McDowell clearly suggests that this is his strategy:

the idea that the rules of a practice mark out rails traceable independently of the reactions of participants is suspect... and insistence that wherever there is going on in the same way there must be rules that can be conceived as marking out such independently traceable rails involves a misconception... We can begin working up to this conclusion by coming to appreciate the emptiness... of the psychological component of the picture: that is, the idea that grasp of a rule is a matter of having one's mental wheels engaged with an independently traceable rail (1981: 146).

McDowell later continues:

The [non-cognitivist] picture [we are considering here] has two interlocking components: the idea of the psychological mechanism correlates with the idea that the tracks we follow are objectively there to be followed, in a way that transcends the reactions and responses of participants in our practices. If the first component is suspect, the second component should be suspect to. And it is (1981: 150).

Unfortunately, other commentators attempting to analyze McDowell's Wittgensteinian argument against non-cognitivism have typically failed to properly discuss the crucial role that (2) plays, and as a result, fail to present McDowell's argument thoroughly and adequately. Gerald Lang, for instance, only discusses (2) in passing (2001: 191) and Olle Blomberg (2007) doesn't discuss (2) at all. I think that Lang and Blomberg pose unconvincing arguments against McDowell because they fail to appreciate (and adequately articulate) this aspect of McDowell's Wittgensteinian argument. In this section I'll discuss the crucial argument against (2), which is a neglected yet central aspect of McDowell's argument. In §7 I'll explain how the rule-following considerations cause problems for non-cognitivist accounts of thick concepts in general, and in §8 and §10 I'll return to offer particular arguments against Lang and Blomberg, respectively.

One might conceive of an initial argument against the non-cognitivist's (2) along the following lines. Imagine that a child, Jones, has correctly solved a finite set of addition problems involving numbers < 57 . After sampling Jones's finite success with such problems, we claim that 'Jones understands the plus rule'. To account for the 'continuity' of Jones's behavior in these cases, we posit the rule 'plus' as one that is fixed independently of the responses of Jones and other problem solvers (Wittgenstein 1953: §218) and we posit a psychological mechanism in Jones by virtue of which this rule is grasped (Wittgenstein 1953: §146). But, as this argument might go, how does positing a psychic grasp of the 'plus' rule account for the 'continuity' of Jones's behavior? For consider also the following rules:

- (i) quus: denoted by ' \oplus ' where $x \oplus y = x + y$, if $x, y < 57$, but $= 5$ otherwise (Kripke 1982: 9)
- (ii) guus: denoted by ' \oplus^1 ' where $x \oplus^1 y = x + y$, if $x, y < 57$, but $= 6$ otherwise
- (iii) buus: denoted by ' \oplus^2 ' where $x \oplus^2 y = x + y$, if $x, y < 57$, but $= 7$ otherwise
- (iiii) tuus: denoted by ' \oplus^3 ' where $x \oplus^3 y = x + y$, if $x, y < 57$, but $= 8$ otherwise

Since, by hypothesis, Jones has so far only solved problems involving numbers < 57 , one might just as legitimately posit Jones as grasping any one of these other rules, or infinitely many others, to 'explain' his finite behavior. Yet if any of an infinite number of rules can be used to explain his finite behavior, then it becomes implausible that his behavior is satisfyingly explained by positing his psychic grasp of the plus rule in particular. That is, positing such a mediating mental state based on Jones's previous behavior gets us no further in understanding how Jones's behavior will continue. Or, imagine that in a new case Jones is given a problem involving numbers > 57 , e.g. '67 + 92'. He answers '7', and *insists* that this is how the rule he had learned was to continue (e.g. he says confidently, 'I know it! Seriously, the rule continues *like this*...'). Although Jones's response here strikes us as odd, his response is still compatible with the examples from which he learned to solve problems, which all involved numbers < 57 . Again, this shows 'that his behavior hitherto was not guided by the psychological conformation we were picturing as guiding it' and that 'the pictured state, then, always transcends any grounds there may be for postulating it' (McDowell 1981: 147). As a result, the 'postulation of the mediating state is an idle intervening step; it does nothing to underwrite the confidence of our expectation [of an agent's behavior]' (McDowell 1981: 148). So the positing of such a psychic 'grasp' does nothing to ground the continuity of Jones's (or any agent's) behavior.⁴

However, the problem with the non-cognitivist account of rule-following is not *just* that some finite behavior can be interpreted by infinitely many rules, although one might see that as a problem (Kripke 1982: 8-27). For if that's all the problem consists in then the non-cognitivist can still claim that there *really are* infinitely many independently fixed rules out in reality, and all that needs to be provided is a convincing story for *how* one psychically grasps the *right* rule. That is, the non-cognitivist might still think of 'grasping a rule' as coming to possess the right mediating mental state – the mental state the wheels of which engage with the right independently traceable

⁴ In his work *Wittgenstein: On Rules and Private Language* (1982) Saul Kripke concedes the skeptical point that there are thus *no facts* that we can follow in accord with a rule; however, Kripke attempts to explain that there are still conditions under which we are *warranted in asserting* of others that they are following a rule. For reasons why Kripke's solution to the so-called 'rule following paradox' fails, see my paper *Wittgenstein, Kripke, and the Rule Following Paradox* (2010).

rule-road – and so think that it's the possession of a mediating mental state that's essential to grasping a rule.

To see why this view fails, let's review a variety of ways it might take shape. Account 1: when we grasp a rule we acquire a mediating mental state the content of which *encodes* a descriptive procedure for how we are to act in accord with the rule (Wittgenstein 1953: §197; Kripke 1982: 10-22). But what exactly is encoded as the content of what we've grasped? One might suggest that the content consists in a descriptive list of procedures. On this view, grasping the plus rule consists in grasping that '1 + 1 = 2', '1 + 2 = 3', '1 + 3 = 4', and so on. But it's surely not the case that in grasping a rule, its entire application somehow appears before the agent's mind (Wittgenstein 1953: §139; Kripke 1982: 22), because the entire application is potentially infinite and includes cases the finite agent has never yet considered. So as it stands, this view is implausible.

Maybe, then, one could adjust Account 1 to Account 2: when we grasp a rule we acquire a mental state the content of which encodes (a) an *abbreviated* list of procedures – e.g. '1 + 1 = 2', '1 + 2 = 3' – and (b) the additional procedure, expressed by *and so on*, indicating that one is to use the examples displayed in (a) as samples to continue on in the same way (Kripke 1982: 10-16). Yet to tell someone to 'continue the same way' is just to tell them to 'continue following the rule', so 'and so on' is not helpful because the very question is what *counts as continuing the same way*. In other words, if I don't already know how to follow the rule for plus then I won't know how to continue from '1 + 2 = 3' in the same way in accord with the plus rule, and so to tell me to continue on with the plus rule in the same way with the phrase 'and so on' is of no help to me. Yet the non-cognitivist might try to adjust (b) of Account 2 into a new Account 3 in order to avoid the circular explanation that an agent can only grasp a rule if she has already understood it in the first place. Account 3: the 'and so on' in Account 2 does not simply mean 'continue on in the same way', but rather contains a further rule providing procedures for continuing on from the first rule (Kripke 1982: 17). For example, this further rule might contain the procedure: 'and so on' is to be continued as '1 + 3 = 4', '1 + 4 = 5', *and so on*. But this move is also unhelpful, because it returns us right back to the problems of Accounts 1 and 2: either (1) the entire application of this further rule is to appear before the agent's mind, or (2) this further rule contains an abbreviated portion of its application along with the procedure, expressed by *and so on*, that one is to use these abbreviated portions as samples to continue on in the same way. And we have already seen why these options are inadequate.

Since accounts 1-3 failed, one might change their approach to Account 4: when we grasp a rule we acquire a mental state that *causally disposes* us to act in accord with a rule (Kripke 1982: 22-37). On this view, the connection between grasping a rule and subsequently applying it is a causal one, so there's no worry of interpreting what is encoded in the rule we grasp. Since our grasp of a rule causes us to act in accord with it, the rule that 'someone means is to be *read off* from his dispositions' (Kripke 1982: 29). But this account cannot work because if we claim that our grasp of a rule causes us to act in accord with it, then there's no possibility of our grasping a rule and *failing* to act in accord with it. Here we are either causally disposed to act in way A and so act in accord with A, or we are causally disposed to act in way B and so act in accord with B. But in either case, since the rule that 'someone means is to be *read off* from his dispositions', there is no *failing* to act in accord with a rule, e.g. failing to act in accord with A. There is just acting in accord with some *other* 'rule', e.g. acting in accord with

B. So dispositional Account 4, as merely descriptive of acts, fails to account for the *normativity* of rules, i.e. that one's responses can be *wrong*. Moreover, since Account 4 fails to provide normative criteria for correct applications of a rule that are independent of an agent's causally disposed responses, it betrays the very reason non-cognitivist appealed to such an objective account of rule-following in the first place; i.e. to maintain objectivity by avoiding rules that are dependent upon agent's responses.

I've now shown that four non-cognitivist accounts of rule-following fail. In Account 4, the connection between grasping a rule and subsequently applying it was too strict. Having rules causally determine our actions does not allow for an account of their normativity. In Account 3 the connection between grasping a rule and subsequently applying it was too loose. Encoding rules with further rules made fixing on an action impossible. And Accounts 1-2 required that we understand a rule before we can grasp it, which puts the cart before the horse. The failure of these accounts shows us that the connection between grasping a rule and subsequently applying it is not plausibly established via the mediation of that rule somehow appearing before our minds. Furthermore, even if a rule *did* appear before one's mind, we can easily imagine that what appears before the mind of an agent that correctly applies a rule can also appear before the mind of an agent that does *not* correctly apply it (Wittgenstein 1953: §152). So what comes before one's mind is not what *grounds* one's grasp of a rule, even if what comes before one's mind is often *associated* with it. Thus, the non-cognitivist has not adequately supported their claim (2) that humans possess psychological machinery with which to grasp independently fixed rules and to act in accord with them. And this in turn renders suspect their claim (1) that there are such independently fixed rules, the following of which is to be explained by a psychic grasp of them. The non-cognitivist's two-component view of rule-following, by divorcing rules from the 'responses and reactions a propensity to which one acquires when one learns the practice' involving the rule itself, suggests a fanciful picture of what a rule is and what it takes to act in accord with it. Upon closer inspection we find that 'there is no such thing here as, so to say, a wheel that he is to catch hold of, the right machine which, once chosen, will carry him on automatically' (Wittgenstein 1970: §304). So we see that the non-cognitivist's view of rule-following not only fails to explain how we can grasp and act in accord with rules, but also fails to suggest a realistic picture of how this is to be achieved. For if we consider how students actually learn rules in concrete cases, we see that all that's involved is a certain (usually extensive) training (Wittgenstein 1953: §198); training into and active involvement in a communal practice where students learn to develop a knack, technical skill, or ingrained sensibility to act and react as others do within their practice. We do not in real cases expect a student to acquire mastery of a rule 'in a flash' via a psychic grasp of all the possible instances of a rule. Rather, we expect the student to progressively become better at mastering the rule as their skill develops and as their active involvement, in the practice where that rule functions, matures and is fine-tuned. As McDowell rightly points out:

the structure of the space of reasons is not constituted in splendid isolation from anything merely human. The demands of reason are essentially such that a human upbringing can open a human being's eyes to them (1996: 92).

6. Our Rules, Culture, and Nature

The ‘objective’ conception of rules appealed to non-cognitivists because they thought it only makes sense to say that one correctly acts in accord with a rule if there are independently fixed rules that one is capable of psychically grasping. Accordingly, they thought that it only makes sense to say that one correctly applies a concept, moral or otherwise, if there are independently fixed rules for the application of that concept that one is capable of psychically grasping. In other words, the non-cognitivist assumes the objective conception of rule-following in an attempt to account for the normativity of concept application; to explain how it is that an application of a concept, moral or otherwise, can go wrong. The problem with ‘non-objective’ accounts of rules, as Gibbard expressed earlier, is that they seem unable to account for the fixedness, and so normativity, of rules. That is, one might worry that by rejecting the objective conception of rules, the Wittgensteinian analysis has led us to what Michael Dummett has called ‘full-blooded conventionalism’, the view that a given statement’s necessity and truth ‘consists always in our having expressly decided to treat that very statement as unassailable’ (1959: 329, 337, 348; McDowell 1981: 150-152). The worry is that because there are no independently fixed rules for us to grasp, our treating any statement as true or false must ultimately result from a decision of how to treat that statement. But if the truth-value of a statement is a result of our decisions, then the truth-value of any statement is as flexible as our decisions are, and in that case we can no longer account for how our decisions about the truth-value of statements can be wrong. The general worry is that non-objective conceptions of rules must also be non-normative ones.

One of Wittgenstein’s philosophical contributions was in showing that this worry is misplaced. By bringing rules out of individual minds and into the public space of culture, he avoids the worrisome claim that it’s the decisions of individuals that determines the truth-value of statements. Indeed, the very point of dismantling Accounts 1-4 in §5 was to show that our rule-following ability is not suitably grounded in the mind (and hence the decisions) of individuals at all (Wittgenstein 1974: II, §33, 70). As Wittgenstein argues in *Philosophical Investigations*:

‘obeying a rule’ is a practice. And to *think* one is obeying a rule is not to obey a rule. Hence it is not possible to obey a rule ‘privately’: otherwise thinking one was obeying a rule would be the same thing as obeying it (§202).

Gibbard and others that worry that non-objective conceptions of rules must also be non-normative ones have missed this point. Rules are not ‘in the mind’ at all, but are rather diachronically stable patterns of activity that have been shaped by the shared sensitivities and responsive actions of the agent’s involved in the communal practice in which that rule has the particular function that it does (see Wittgenstein 1953: §198-199). Like a well-trodden path through the wilderness a rule receives its significance in part by the history of its use, the constitution of its users, and the purpose for which they used it, while further hinting at the direction for how it should continue.

A rule is a constituent path within a community-wide practice, and because of its public nature, what constitutes following in accord with a rule is not some individual’s decision. Insofar as rules are not already ‘out there’ independently fixed from human response, our ‘going on in the same way’ in accord with a rule is in fact contingent insofar as it depends upon facts regarding our natural and socio-cultural constitution and

tendencies. But it would be confused to think that this entails that our ‘going on in the same way’ in accord with a rule is decidedly *arbitrary*. I suggest that it’s just a contingent fact that we are born into a world equipped with needs (e.g. to eat and mate, etc.), purposes to achieve (e.g. to acquire food and mates, etc.), and instinctive sensitivities (e.g. to see things as edible or mate-able, etc.) and for, presumably evolutionary, reasons, nature has it that those of us who are alive today share in our basic needs, purposes, and sentient architecture such that we can work to attain these together.⁵ It’s this common rootedness in nature, need, and purpose that brings our sensitivities and responsive actions together to form the practices that we have, and we don’t decide to accept the procedures of our practices anymore than we decide to accept these practices themselves:

We do not decide to accept them or reject them at all, any more than we decide to be human beings as opposed to trees. To ask whether our human practices or forms of life themselves are “correct” or “justified” is to ask whether we are “correct” or “justified” in being the sorts of things we are (Stroud 1965: 518).

Furthermore, insofar as we are to achieve the ends of our practices, the steps we take are not arbitrary. As Wittgenstein puts it, ‘the rules aren’t arbitrary if the purpose of the game is to be achieved’ (1974: §140; Stroud 1965: 515). And since we share in our ways of living it’s also evident that we share in the rules for living those ways. So to ask if an agent is following in accord with a rule, then, only makes sense when that agent is considered within the wider context of her community of fellow practitioners. And it’s with respect to these wider communal practices that an individual’s actions are to be considered as correct or incorrect. That’s why it’s not sensible to ask how an agent, divorced from the context of practice, can follow in accord with a rule that derives its sense from within that practice. A ‘person goes by [an expression of a rule, such as] a signpost only insofar as there is an established usage, a custom’ (Wittgenstein 1953: §198).

7. Implications of McDowell’s Wittgensteinian Argument against Non-Cognitivism

Given that McDowell’s article *Non-Cognitivism and Rule-Following* relies extensively on Wittgenstein’s work on rule-following, it’s unfortunate that commentators of McDowell have not provided an adequate picture of Wittgenstein’s work as a background to McDowell’s argument against non-cognitivism. This is what I have provided in §5-6. Let’s now review how McDowell’s Wittgensteinian rule-following analysis dismantles non-cognitivism.

First, the non-cognitivist thought that in order to apply a concept correctly, there must be some objective rule that is independent of human sensitivities and responses that determines that concept’s correct application. Further, the non-cognitivist claims that we can only act in accord with the rule for that concept’s correct application by

⁵ For example, it might be the case that our ability to enter into complex practices with one another is grounded upon and made possible by ‘The psychic unity of humankind – that is, a universal and uniform human nature – [which] is necessarily imposed to the extent and along those dimensions that our psychologies are collections of complex adaptations [... it may be that] selection, interacting with sexual recombination, tends to impose at the genetic level near uniformity in the functional design of our complex neurocomputational machinery’ (Tooby & Cosmides 2005: 39) and, presumably, it’s such a uniformity in aspects of our nature that makes social coordination in judgments and actions possible. Note that this point is not attributed to Wittgenstein.

cognitively grasping it. In §5-6 we considered four different accounts of how this 'grasping' might be achieved, but we found them all inadequate. So thus far the non-cognitivist has failed to provide a convincing account of how it is that we can follow in accord with a rule for the correct application of our concepts, moral or otherwise. Moreover, it was precisely this failed non-cognitivist assumption that rules for the correct application of concepts are fixed 'independently of the responses and reactions a propensity to which one acquires when one learns the practice itself' (McDowell 1981: 146) that motivated the non-cognitivist to suppose that 'any respectable evaluative concept must correspond to a classification intelligible from outside the evaluative outlook within which the concept functions' (McDowell 1981: 153). That is to say, it's because the non-cognitivist supposes that the correct use of concepts in general requires that one psychically grasp an independently fixed rule for their correct application, that the non-cognitivist also supposes that the correct use of evaluative concepts in particular requires that one psychically grasp an independently fixed rule for their correct application (as the latter is a subset of the former). And it was this erroneous assumption that led the non-cognitivist to suppose that the extension of a thick moral concept could be mastered independently of an involved understanding of the practice of moral evaluation itself. However, as the Wittgensteinian analysis has shown, it's misguided to construe rules as fixed independently of human responses. Resultantly, there need be no supposition that the correct use of concepts in general requires that one grasp an independently fixed rule for their correct application, and so there need be no supposition that the correct use of evaluative concepts in particular requires that one grasp an independently fixed rule for their correct application either. So the Wittgensteinian analysis does not suppose that the extension of a thick moral concept could be mastered independently of an involved understanding of the practice of moral evaluation. Indeed, it's suggestive of the idea that it only makes sense to say of an agent that she can correctly apply a moral concept when that agent is considered within the wider context of her community of fellow moral evaluators. So it's clear that the Wittgensteinian analysis that McDowell offers not only avoids the problems that the non-cognitivist faced with respect to mastery of thick concepts in §4 and with respect to rule-following in §5, but also explains how these problems in §4 and §5 are to be corrected.

Since thick concepts are not independent of evaluation, it also follows that thick concepts don't have purely descriptive shape. Resultantly, mastery over thick concepts is also guided by an evaluative sensitivity. But notice that this casts doubt on the non-cognitivist's view that (a) cognitive-states with descriptive content are distinct from (b) non-cognitive affective or evaluative states, and that (b) are simply post-cognitive responses that are incapable of disclosing features of the world (Williams 1985: 141). Indeed, since mastery over thick concepts requires utilizing evaluative and descriptive sensitivity in order to pick out those features that aptly belong under that thick concept's extension (see §4), we have reason to be suspicious of the non-cognitivist's strict separation of (a) from (b) since it claimed that only (a) is capable of determining a concept's extension. In §4 we saw why this non-cognitivist claim was incorrect. So McDowell's argument not only attacks the non-cognitivist account of thick concepts in particular, but their presupposed bifurcated conception of mind more generally (McDowell 1981: 143, 154-156).

Notice also that the absolute conception was originally appealing because it was thought to be explanatorily powerful (Williams 1978: 245-246; McDowell 1981: 142;

Lang 2001: 194). The absolute conception was supposed to explain how the world as it appears to each of us relates to each other and to the world as it is independently of us (Williams 1978: 245-246). However, as we saw in §5, the absolute conception of rules as fixed objectively and ‘independently of the responses and reactions a propensity to which one acquires when one learns the practice itself’ (McDowell 1981: 146) was not only unable to account for the vast range of thick concepts we actually use in articulating features of people’s actions and the world as conceived by us (see §2) but also rendered suspect how one ‘grasps’ and applies the concepts of *any* features of the world at all. By imposing the restriction that what is ‘objective’ is fundamentally separated from our responses, the absolute conception made it impossible for the world to be something graspable by us at all. The purity it sought was otherworldly, and so it seems reasonable that we ‘should accept sometimes that there may be nothing better to do than explicitly appeal to a hoped-for community of human response’ (McDowell 1981: 153). The demand for more, that there *must* be a sense of objectivity or genuineness beyond that of universal agreement in human responses, actions, and judgments, is presumably rooted in nothing more than the metaphysical assumption that something like the absolute conception *must* be right. But that’s an assumption we’ve seen is problematic. Importantly, the point here is not that there’s something fundamentally wrong with naturalistic accounts of the world. Rather, the analysis here simply shows that the absolute conception has not lived up to its expectations of explanatory power and that a naturalistic account of the world construed in terms of the absolute conception will invariably remain an inadequate one.

8. Challenge One: Gerald Lang’s Criticism

McDowell’s argument against non-cognitivism has seen numerous challenges. The first that I’ll discuss is one offered by Lang. According to Lang:

McDowell’s... argument takes the form of a *modus tollens*... [but his argument is] flawed, because the conditional statement ‘If the outsider could perform the disentangling manoeuvre, then he could predict the extension of the concept’, is false. The ability to perform the disentangling manoeuvre does not entail mastery of the extension of the concept... [and so] differences between insiders [of the community of moral evaluators] and outsiders do not support McDowell’s arguments against, specifically, the disentangling manoeuvre (2001: 203).

Now, this passage suggests that Lang has unfortunately misunderstood McDowell’s argument. Because, contra Lang, McDowell doesn’t claim that the ability to perform the disentangling manoeuvre – the ability to distinctly separate description from evaluation – *entails mastery* of the extension of a thick concept. That’s not McDowell’s point at all. Rather, McDowell is claiming that if the disentangling manoeuvre is possible, then ‘outsiders’ should be able *in principle* to master the extension of a thick concept *without needing to hold or understand the perspective from inside the community of evaluators and evaluative practice*. As I’ve explained in the preceding sections, it’s clear that McDowell’s argument is that thick concepts have both descriptive and evaluative shape, and that as a result an ‘outsider’ tracking purely descriptive features independently of an evaluative outlook (of those inside the community of evaluative practitioners) *cannot master* the extension of a thick concept. The outsider lacks a sense for ‘aptness’; e.g. a sense for what counts as an action appropriately belonging to the extension of a thick moral concept in certain circumstances. So Lang is mis-

guided in claiming that McDowell's argument against the disentangling manoeuvre is not supported by the differences between (a) 'outsiders' that track purely descriptive features independently of an evaluative outlook, and (b) 'insiders' whose evaluative outlook enables them to pick out those features that aptly belong under the extension of a thick concept. As I've argued at great length by now, it's precisely *this* difference of being outside the evaluative practice, and so being restricted to tracking purely descriptive features independently of evaluation, that prevents 'outsiders' from determining which features are apt for a thick concept and so prevents them from developing a mastery over a thick concept's extension.

9. Challenge Two: The Sufficiency of Imaginative Identification

Another challenge to McDowell that appears often is what I'll call 'the sufficiency of imaginative identification' challenge. J. E. J. Altham (1986) posed this type of challenge as follows:

it does not seem that, in order to grasp the extension of the [thick] term, one must actually share the evaluative perspective of those who use it. It would be enough to have a merely imaginative identification with their perspective... [moreover] if a merely imaginative identification suffices, then the thought arises that once it has been achieved, and the term grasped, a neutral equivalent *can* be introduced (278-279).

Olle Blomberg (2007) also writes:

it is not clear what the difference is between (a) accepting the evaluation embedded in a thick concept ("sharing values"), and (b) grasping the thick concept's "evaluative point" imaginatively without accepting it, where (b) cannot amount to acquiring a purely descriptive equivalent of the concept. Hare's attribution [that one must *accept* the evaluation embedded in the community where the thick concept functions in order to master that concept's extension] is not entirely unwarranted until Entanglers [McDowellians] provide a clear account of this difference (72-73).

First, we shouldn't get carried away with the idea of 'imaginative' identification; it's still to be identification nonetheless. And we already saw in §4 that e.g. only certain acts of strength and determination in the face of conflict are aptly 'courageous', and that how these 'certain' acts are to be picked out involves more than simply tracking some pre-evaluative set of descriptive features. If Jones is to master the extension of the thick concept 'courageous' through an 'imaginative identification' with the perspective from which those actions are appropriately collected together, then the imaginings of Jones must be *appropriately constrained* such that they are in line with the perspective that other moral evaluators hold *non-imaginatively*. That is, the imaginative case is parasitic on the genuine case and so not just *any* imaginings will enable Jones to collect just those actions that are apt for the extension of 'courageous'. Only *certain* imaginative identifications will be appropriate, and as we saw in §4-5, determining what *counts as appropriate* here is dependent upon an evaluative outlook. Furthermore, mastery over a thick concept consists in 'be[ing] able to predict applications and withholdings of it in new cases' (McDowell 1981: 144) which requires that one is capable of utilizing this evaluative outlook in *new* cases in order to pick out e.g. just those actions that are aptly 'courageous' in some new case (one will hardly ever be interested in only identifying 'courage' in old cases). So, on the one hand, it was by vir-

tue of utilizing an evaluative outlook that one identified *old* cases of ‘courageous’ acts, and on the other hand, it’s by virtue of utilizing an evaluative outlook that one can pick out just those acts that are aptly ‘courageous’ in *new* cases too. So insofar as identification of evaluative features in both old and new cases requires an evaluative outlook, *re-identification* of evaluative features *as the same* evaluative features requires an evaluative outlook too. So even if I wanted to construct a ‘neutral equivalent’ of an evaluative term and apply it in new cases, it’s still required that I utilize an evaluative outlook in order to identify those features in the object to which I intend to apply the neutral term as possessing features that are *aptly similar* to those of the evaluative term which I am trying to replace. Without utilizing an evaluative outlook I may not be using the ‘neutral term’ as an *apt equivalent* of the evaluative term at all. Thus, it’s unclear what this ‘imaginative identification’, which Altham (1986) suggests, actually consist in, and how one in this ‘imaginative’ state is to determine what actions *count as appropriate* ones for the extension of a thick concept without already depending upon an evaluative outlook to do this.

However, there is something right in raising the challenge of imaginative identification. Its importance is that it brings out a question about what constitutes the difference ‘between (a) accepting the evaluation embedded in a thick concept... and (b) grasping the thick concept’s “evaluative point” imaginatively without accepting it’ (Blomberg 2007: 72-73). The difference between (a) and (b) will be best explicated by an example. Imagine Jones, a young boy that, throughout some period during his upbringing, was taught to evaluate e.g. a person’s actions of strength and determination in the face of conflict. Maybe he grew up in a tribe where such actions are helpful to one’s community at large, or maybe he was born into city life where such actions help one procure jobs and income. Regardless, Jones’s training consists in being educated into a social world that makes sense of evaluating people on the basis of their actions in the face of conflict. It’s likely that Jones was taught that such actions are worthy of praise and emulation, but it’s possible that he was taught otherwise. Either way, Jones’s moral education is still such that he has acquired a sensitivity and sensibility to see that a person’s actions are such that *some evaluation or other* is called for. One could say: regardless of the evaluative *direction* Jones’s attitude is pointing (e.g. pro or con), it’s first required that Jones has become sensitive to an evaluative *point* (e.g. *that* a person’s actions are such that some evaluation or other is called for). And as long as Jones acquires sensitivity to the evaluative practice such that he can identify certain relevant features of an action as salient (e.g. strength and determination in the face of conflict), and so understand that some evaluative attitude or other is called for by that action, Jones is *capable of conceiving* that one might evaluate these same features with an attitude *opposite* of his. This is an option that, although not likely to convince Jones, is at least intelligible to him. And so *this* is how (b) grasping the thick concept’s evaluative point imaginatively without accepting it, is to be distinguished from (a) accepting the evaluation embedded in a thick concept. In other words, the difference is that in (b) one must simply acquire sensitivity to an evaluative practice such that one can identify certain relevant features of an action as salient (e.g. strength and determination in the face of conflict) and so understand *that* an evaluative attitude is called for by that action, while in (a) one has further fixed on *which* direction their evaluation points. And if one cannot make sense of some situation having an evaluative point at all, then that is all the same as saying that one cannot grasp the thick concept’s evaluative point imaginatively.

Although Jones is capable of conceiving of the possibility that one might evaluate these same features with an attitude opposite of his, Jones, having himself learned the evaluative point for particular practical purposes, will of course find one evaluative direction most natural since this is the evaluative direction in line with the purpose for which the evaluative point was taught. For example, Jones might have been taught to evaluate a person's act of strength and determination in the face of conflict positively, because such acts are those that best enable one to survive in the kind of environment one happens to be living in. Given Jones's bio-cultural environment, and thus the sort of evaluative training he received in order to productively engage in those environments, Jones will naturally find one evaluative direction as his default. A 'neutral' evaluative direction (of which it's hard, but maybe not impossible to find genuine cases) would be an exception case because to find no evaluative direction as natural to hold would suggest that there was no strong initial motivation for which the evaluative point was learned, and in that case, the relevant thick concept would have been without much purpose anyway. However, I doubt that there are genuine cases of 'neutral' evaluation that are still genuinely evaluative; they are more likely to be cases of multi-directional evaluation of which the weighing of directions makes a single direction less compelling.

Finally, it's important to mention (since I'll make use of this point later) that what I've said so far in no way precludes the possibility that one can *change* the direction of their evaluative attitude, or even that there must be *one* attitude connected to a particular thick concept. It's unlikely that an attitude connected to a thick concept will be only *loosely* connected, varying unpredictably, if the purpose for which one learned the evaluative point is sufficiently strong or widespread in their community. However, there will be cases when one, after having learnt an evaluative point, can change their evaluative direction. For instance, our young boy Jones might have been taught an evaluative point for a purpose that is marginal to the community in which he lives. Or maybe he has just learnt the evaluative point and has not encountered many real-life instances in which he has exercised his evaluative direction. But it's possible that Jones will hereafter change his evaluative direction given that he's sufficiently educated to do so.⁶ That is, since Jones had originally learned the evaluative point (*that* a person's actions are such that some evaluation or other is called for) for particular practical purposes, he found one evaluative direction most natural since that was the evaluative direction in line with the purpose for which the evaluative point was taught. But this is compatible with the possibility that Jones may learn the same evaluative point for a *different* purpose, and so in turn come to find another evaluative direction most natural since this is now the evaluative direction in line with the new purpose for which the evaluative point was taught. If the purpose of the latter e.g. is more central to the community in which Jones is involved or taught to Jones to a greater extent, there's nothing in what I've said that suggests that a change in Jones's evaluative direction would be impossible. And given the complexity of our lives in general and our social ties in particular, and the different directions in which they pull at us, there may

6 This is an important point to make, and I think it's especially helpful when applied to, for instance, racial discrimination. That is to say, this point makes it clear that, even if e.g. one was raised as a child to evaluatively pick out certain racial features of others for certain practical purposes, we're still capable of *changing* or *neutralizing* our *evaluative direction* with respect to racial features. See my papers *Racial Epithets: What We Say and Mean by Them* (2008) and *Slurs: A Philosophical Analysis* (m.s.) for an analysis of how we might come to characterize and evaluate people on the basis of racial features.

be several different crosscutting purposes for which we learn an evaluative point. So although evaluative concepts require an evaluative *outlook* for their mastery, this in no way entails that an evaluative concept is restricted to *one* evaluative *direction* (see Blackburn 1992: 294-299; see also Dancy 1996: 266-272). In short, the problem with the sufficiency of imaginative identification challenge that critics of McDowell have posed, is that it assumes that in order to master the extension of a thick concept one must, not only grasp its evaluative point, but also hold *a particular* evaluative direction. But as I've explained in this section, this needn't be the case at all.

10. Challenge Three: Alan Gibbard and Simon Blackburn's Non-Cognitivist Accounts

Let's now turn to more recent developments of non-cognitivism, starting with Gibbard's norm-expressivism, to see if they fare better than the other classical accounts we've considered. According to Gibbard:

when the licensed user of the [thick] term calls an act [e.g.] lewd, he means that L-censoriousness [i.e. 'the feeling of outraged shock and censure that goes with finding something lewd'] toward its agent is warranted for passing beyond these limits. The content of these limits is not built into the meaning... Rather, he presupposes that there are such limits... [these limits refer to] standards of warrant for feelings of L-censoriousness toward lewd acts (1992: 280-281).

According to Gibbard, moral statements express non-cognitive attitudes, such as acceptance, towards a system of norms warranting feelings in response to certain actions. But notice that Gibbard's account diverges from classical accounts of non-cognitivism in that he does not claim (like Hare had) that description and evaluation are both *lexical* components in the meaning of the thick concept. For Gibbard, e.g., the non-cognitive attitude of accepting a system of norms warranting feelings of shock 'is not built into the meaning' of the thick concept 'lewd'. One just expresses their acceptance of a system of norms warranting feelings of shock towards the object that they are calling 'lewd'. Gibbard's also differs from classical accounts in that he doesn't claim that description alone guides the application of thick concepts; sometimes our feelings do: 'In applying a [thick] term like 'lewd'... people are guided alike by the descriptive features of the case and by their feelings about... [what] various descriptive kinds of acts warrant' (Gibbard 1992: 282, 271-272). Although Gibbard concedes that description and evaluation may not be independent (1992: 272), he clearly thinks that description and evaluation can still be *disentangled* (1992: 268). This is evident given that Gibbard considers evaluative feelings as *warranted responses* to descriptive kinds; when we are not guided directly by description, we are guided by those feelings that '*descriptive kinds of acts warrant*' (Gibbard 1992: 282, my emphasis). Gibbard's non-cognitivism may differ from classical versions in interesting ways, but he's done nothing to distance himself from accepting the disentangling manoeuvre and so he fares no better against McDowell's arguments than classical non-cognitivists. Finally, Gibbard claims that 'Norm acceptance, I speculated, is the work of a special psychic mechanism that ties discussion to action' (1992: 278). That is, Gibbard claims that the normative rules that one accepts, and that one e.g. expresses acceptance of in moral statements, is grasped by some mediating psychic mechanism in the mind. But I've already argued at great length in §5 that this idea is implausible.

Let's now turn to Simon Blackburn's non-cognitivist account. Like Gibbard, Blackburn doesn't claim that description and evaluation are both lexical components in the meaning of the thick concept. Indeed, Blackburn is suspicious that there even *are* thick concepts (1992: 285). According to Blackburn:

If a person's industry is entirely misdirected... we do not deny that it is [still] industry... people can be too prudent and too discerning. But that does not *stop* them from being industrious, prudent and discerning (1992: 286).

The idea is that although one can be industrious, which is thought to involve a pro-attitude, one can be *too* industrious, which is thought to involve a con-attitude. Blackburn regards this as showing that it 'is actually extremely difficult to say *which* attitude [e.g. pro or con], if any, is fixed as part of the literal meaning of most of these terms... [and so] I suggest that there is no stable connection between any single attitude and such a term' (1992: 294). And because there's no stable connection between a single attitude and a thick term, e.g. between the thick concept 'greedy' and the con-attitude of moral disapprobation, Blackburn argues that 'the point remains that an individual element of the mix – moral disapprobation, for example – can be removed without semantic rupture' (1992: 297). That is, Blackburn claims that the con-attitude of moral disapprobation can be removed from the thick concept 'greedy' without changing that concept's meaning, and as a result, one could master the extension of that thick concept independently of an involved understanding of the practice of moral evaluation. In another example, Blackburn claims that 'discussion, for instance of whether Pavarotti is fat↓' [where 'fat' is the description and '↓' is the tone expressing a con-attitude] is nothing new from discussion of whether to feel repelled or not at his weight' (1992: 297). It's clear from these examples that Blackburn thinks that the extension of a thick concept is determined by descriptive features alone (e.g. weight), and that an evaluative attitude (e.g. as carried by the tone) can be disentangled from the description, and even removed, without alteration to that concept's shape. So on Blackburn's view, applying a thick concept to Pavarotti is similar to applying a descriptive concept to Pavarotti but with a tone in your voice that, considering the context of your speech act, licenses the hearer to suppose that you are expressing an attitude towards Pavarotti on the basis of the applied description (Blackburn 1992: 289, 297). And in general, to utter an evaluative statement is similar to uttering a descriptive statement with a tone in your voice that, considering the context of your speech act, licenses the hearer to suppose that you are expressing an attitude on the basis of the description. Blackburn thinks this analysis of his has undermined the McDowellian picture of thick terms.

There are several points to make here. First, Blackburn hasn't really offered an independent *argument* against McDowell on behalf of the legitimacy of the disentangling manoeuvre. Blackburn simply presupposes its legitimacy in the examples he uses to support non-cognitivism. In the Pavarotti example above, for instance, he simply assumes that the extension of the thick concept 'fat' is determined by descriptive features alone (e.g. weight), and that an evaluative attitude (e.g. as carried by the tone) can be disentangled from the description, and even removed, without alteration to that concept's shape. But as I argued in §4, an evaluative outlook is required in order to determine the extension of a thick concept, so even the thick concept 'fat' cannot be determined by descriptive features alone. For instance, during Arnold Schwarzenegger's bodybuilding career he had a competition weight of 240-250 pounds, but as evidenced by the fact that he won Mr. Olympia seven times, it's clearly not appropriate to apply

the thick concept 'fat' to him on the basis of his weight alone. And surely there are voluptuous women, such as the supermodel Tyra Banks, to which the thick concept 'fat' would be an inapt application in spite of the fact that they are heavier than less voluptuous women that couldn't make it as supermodels. So the legitimacy of the disentangling manoeuvre is not something that one can simply assume.⁷

Blackburn attempted to motivate his non-cognitivism by pointing out that although one can be 'industrious', one can be *too* 'industrious'. This was supposed to show that there is *no* attitude, e.g. pro or con, that has a stable connection to the thick concept 'industrious'. Presumably then, the evaluative attitude is an inessential component to the thick concept 'industrious'. But Blackburn's move here is suspect. First, notice that Blackburn's criticism of McDowell runs parallel to the criticism posed by the challenge of imaginative identification. Blackburn claims that the McDowellian picture of thick terms implies that ethical activity are 'exercises in talking past each other' on the grounds that 'if you don't respond to [a thick term such as] lewdness as I do, then... your amalgamated concept of lewdness is not mine, and we are left in incommunicable solitude' (1992: 299). The reason why Blackburn thinks that McDowell is committed to the claim that a correct application of a thick concept to an object requires that one *accept a particular* evaluative direction for that object runs parallel to the reason why, in §8, Blomberg claimed that 'Hare's attribution [that one must *accept* the evaluation embedded in the community where the thick concept functions in order to master that concept's extension] is not entirely unwarranted'. Given what I showed in §8 in response to Blomberg and Altham, we can see that the problem with Blackburn's criticism rests on his mistakenly conflating an evaluative point with an evaluative direction. Blackburn assumes that if there is no singularly stable evaluative attitude connected to a thick concept, then there is no 'thickness' to that concept at all and that we should therefore be able to master that concept without an evaluative outlook. But it doesn't follow that because it's not required that a thick concept have a singularly stable *evaluative direction*, that a thick concept can therefore be mastered independently of a sensitivity to its *evaluative point*. Indeed, in §8 I explained how, although one will find a particular evaluative direction natural to a thick concept by virtue of that evaluative direction being in line with the purpose for which the evaluative point was taught, there needn't be a *singularly stable* evaluative direction for a thick concept at all.

In §8 I pointed out that moral education has a socio-cultural function, and when one acquires mastery over a thick moral concept one acquires a sensitivity to an evaluative point for a particular purpose or purposes. However, evaluative directions can change and one can be pulled in different directions, because the purpose or purposes for which one learns an evaluative point needn't be singular. In fact, they are presumably as complicated as our social networks and lives are. Of course for each thick concept there *will* be an evaluative direction that one finds most natural (i.e. the thick concept's 'default direction'). And this is clearly the case since even Blackburn relied upon a thick concept's natural evaluative direction in order to suggest that thick concepts *have no* evaluative direction. That is, Blackburn initially attempted to motivate his non-cognitivism by pointing out that although one can be 'industrious', one can be *too*

7 Indeed, it seems that one would be in violation of an appropriateness condition on the term 'fat' by applying it to paradigm bodybuilders and supermodels, and if a speaker continued to misapply such a concept in this way, that speaker could provide warrant for their linguistic community to identify them as incompetent or unworthy of engaging in sensible communication.

‘industrious’. What’s important to notice is that Blackburn clearly could not arouse our intuition that ‘industrious’ is to be negatively evaluated without *augmenting* ‘industrious’ with ‘too’, which is defined *evaluatively* as ‘an excessive extent or degree; beyond what is desirable, fitting, or right’ (Collins English Dictionary 2003). That is, our default understanding of ‘Jones is industrious’ or ‘Jones is courageous’ is with a positive evaluative direction, and augmenting such statements with ‘too’ alters our understanding on semantic grounds because saying of X that it is *too* X is to say that X is ‘excessive [in] extent or degree; beyond what is desirable, fitting, or right’. It is to express semantically that something that we may have had a pro-attitude towards is excessive to the extent that it is no longer *apt* for a pro-attitude. So if attitudes are really as fluid and inessential to thick concepts as Blackburn suggests, then it’s at least suspicious why Blackburn found it necessary to augment the thick concept ‘industrious’ with ‘too’ to make his point. Strangely, Blackburn relies on the natural evaluative directions of thick concepts in attempting to prove that such concepts have no natural evaluative directions at all. I find this move on Blackburn’s part suspect. Resultantly, it’s clear that Blackburn has not only failed to convincingly criticize the evaluative nature of thick concepts (of which he thinks there are none), but has also failed (along with other non-cognitivists) to defend the legitimacy of his use of the disentangling manoeuvre against McDowell’s arguments.

11. Thick Moral Concepts and the Moral of Thick Concepts: Concluding Remarks

More than fifty years after Ayer’s *Language, Truth and Logic*, non-cognitivists are still motivated by a naturalism that takes for granted the external viewpoint of the absolute conception (Blomberg 2007: 64). In line with this external viewpoint, the non-cognitivist thought that in order to apply a concept correctly, one must psychically grasp a rule for its correct application that is fixed independently of human responses. We saw how McDowell used Wittgenstein’s rule-following considerations to show that this view is fancifully false, and that this view imposes on us the demand that we transcend the bounds of our cognitive powers (McDowell 1981: 153). However, as demanding as the non-cognitivist position is, it has itself been unable to account for even the simplest cases of moral reasoning (see especially Schueler 1988; van Roojen 1996; Unwin 1999; Schroeder 2008). Take an example of moral reasoning from Geach (1965: 463):

If doing a thing is bad, getting you little brother to do it is bad.

Tormenting the cat is bad.

Ergo, getting your little brother to torment the cat is bad.

The problem here, which is known as the Frege-Geach problem, is that:

The whole nerve of the reasoning is that “bad” should mean exactly the same at all four occurrences – should not, for example, shift from an evaluative to a descriptive... use. But in the major premise the speaker... is certainly not uttering acts of condemnation (Geach 1965: 463-464).

That is, we derive the conclusion via *modus ponens* from the first and second premise without equivocation. However, while in the second premise ‘bad’ occurs as part of the assertion ‘tormenting the cat is bad’, and so e.g. is presumably expressive of a non-cognitive attitude towards the tormenting of cats, the first premise is a conditional statement where ‘bad’ occurs *unasserted* as part of the antecedent and consequent of

the conditional, and so no non-cognitive attitude towards the tormenting of cats is presumably expressed by that conditional statement. The trouble for the non-cognitivist here is how to account for the unequivocal meaning of the moral term 'good' that occurs in such embedded and un-embedded contexts. The problem for the non-cognitivist in general is to explain how, in spite of their claim that thick moral concepts have a different kind of meaning than normal descriptive concepts, both thick moral concepts and normal descriptive concepts function similarly across linguistic constructions and are subject to reasoning with the same logical tools. Unlike non-cognitivists, cognitivists claim that statements involving moral concepts express truth-evaluable beliefs, so they face no problem accounting for how we can reason with moral concepts.

I bring up the Frege-Geach problem here to solidify a final point: non-cognitivism lacks explanatory power across the board. Just from what I've pointed out in this paper, we see that non-cognitivism is unable to live up to its claims for how one can reason with thick concepts, how one can master the extension of a thick concept independently of evaluation (see §4), how one can grasp and act in accord with rules in general and for the correct application of thick concepts in particular (see §5), and how one can have an 'imaginative identification' with an evaluative perspective that is not itself evaluative (see §9). Yet I have accounted for all of these points from the broadly cognitivist (or, to put it more accurately yet less elegantly, 'anti-non-cognitivist') position I share with McDowell (McDowell 1981: 154). If the strongest motivation one has to accept non-cognitivism is a prior appeal to a metaphysical picture of an absolute conception of the world with mental mechanisms that somehow hook onto this world, then it's about time we got rid of this picture (see §7). Indeed, my goal in this paper was not simply to show why the case of thick concepts should lead us to reject non-cognitivism, but why it should lead us to question these deeper commitments upon which non-cognitivism is typically based.

References

- Altham, J.E.J. 1986. 'The Legacy of Emotivism'. In: *Fact, Science and Morality*, G. MacDonald and C. Wright (eds.), Oxford: Blackwell.
- Ayer, A.J. 1952. *Language, Truth and Logic*, New York: Dover Publications.
- Blackburn, S. 1992. 'Through Thick and Thin', *Proceedings of the Aristotelian Society*, Supplementary Volumes 66: 285-299.
- Blomberg, O. 2007. 'Disentangling the Thick Concept Argument', *Sats – Nordic Journal of Philosophy*, 8(2): 63-78.
- Burton, S. 1992. "'Thick' Concepts Revised", *Analysis*, 52(1): 28-32.
- Croom, A. 2008. 'Racial Epithets: What We Say and Mean by Them', *Dialogue: Journal of Phi Sigma Tau – International Honor Society for Philosophy*, 51(1): 34-45.
- Croom, A. 2010. 'Wittgenstein, Kripke, and the Rule Following Paradox', *Dialogue: Journal of Phi Sigma Tau – International Honor Society for Philosophy*, 52 (3): 103-109.
- Croom, A. manuscript. 'Slurs: A Philosophical Analysis'.
- Dancy, J. 1996. 'In Defense of Thick Concepts', *Midwest Studies in Philosophy*, XX: Moral Concepts: 263-279.

- Dummett, M. 1959. 'Wittgenstein's Philosophy of Mathematics', *The Philosophical Review*, 68(3): 324-348.
- Geach, P.T. 1965. 'Assertion', *The Philosophical Review*, 74(4): 449-465.
- Gibbard, A. 1992. 'Thick Concepts and Warrant for Feelings', *Proceedings of the Aristotelian Society*, Supplementary Volumes 66: 267-284.
- Hare, R.M. 1970. *The Language of Morals*, London: Oxford University Press.
- Johnston, M. 1989. 'Dispositional Theories of Value', *Proceedings of the Aristotelian Society*, Supplementary Volumes 63: 139-174.
- Kripke, S. 1982. *Wittgenstein: On Rules and Private Language*, Cambridge: Harvard University Press.
- Lang, G. 2001. 'The Rule-Following Considerations and Metaethics: Some False Moves', *European Journal of Philosophy*, 9(2): 190-209.
- Mackie, J. 1977. *Ethics: Inventing Right and Wrong*, Harmondsworth: Penguin.
- McDowell, J. 1981. 'Non-Cognitivism and Rule-Following'. In: *Wittgenstein: To Follow A Rule*, S. Holtzman and C. Leich (eds.), London: Routledge and Kegan Paul.
- McDowell, J. 1996. *Mind and World*, Cambridge: Harvard University Press.
- Schroeder, M. 2008. 'What is the Frege-Geach Problem?', *Philosophy Compass*, 3/4: 703-720.
- Schueler, G.F. 1988. 'Modus Ponens and Moral Realism', *Ethics*, 98(3): 492-500.
- Stevenson, C. 1937. 'The Emotive Meaning of Ethical Terms', *Mind*, New Series, 46(181): 14-31.
- Stroud, B. 1965. 'Wittgenstein and Logical Necessity', *The Philosophical Review*, 74(4): 504-518.
- 'Too'. Def. 2. 2003. *Collins English Dictionary*. 6th ed.
- Tooby, J. & Cosmides, L. 2005. 'Conceptual Foundations of Evolutionary Psychology'. In: *The Handbook of Evolutionary Psychology*, D.M. Buss (ed.), Hoboken: Wiley.
- Unwin, N. 1999. 'Quasi-Realism, Negation, and the Frege-Geach Problem', *The Philosophical Quarterly*, 49(196): 337-352.
- van Roojen, M. 1996. 'Expressivism and Irrationality', *The Philosophical Review*, 105(3): 311-335.
- Williams, B. 1978. *Descartes: The Project of Pure Enquiry*, Harmondsworth: Penguin.
- Williams, B. 1985. *Ethics and the Limits of Philosophy*, Cambridge: Harvard University Press.
- Wittgenstein, L. 1953. *Philosophical Investigations*, G.E.M. Anscombe and R. Rhees (eds.), G.E.M. Anscombe (trans.), Oxford: Blackwell.
- Wittgenstein, L. 1970. *Zettle*, G.E.M. Anscombe (trans.), Berkeley: University of California Press.
- Wittgenstein, L. 1974. *Philosophical Grammar*, R. Rhees (ed.), A. Kenny (trans.), Berkeley: University of California Press.