Conceptual Role Semantics and the Reference of Moral Concepts

Neil Sinclair

Abstract: This paper examines the prospects for a conceptual or functional role theory of moral concepts. It is argued that such an account is well-placed to explain both the irreducibility and practicality of moral concepts. Several versions of conceptual role semantics for moral concepts are distinguished, depending on whether the concept-constitutive conceptual roles are (i) wide or narrow (ii) normative or non-normative and (iii) purely doxastic or conative. It is argued that the most plausible version of conceptual role semantics for moral concepts involves only 'narrow' conceptual roles, where these include connections to motivational, desire-like, states. In the penultimate section it is argued, contrary to what Wedgwood, Enoch and others have claimed, that such an account of moral concepts cannot plausibly be combined with the claim that moral concepts refer to robust properties.

What is a moral concept? An answer to this question should cohere with our best general theories of the nature of concepts. In this paper I take one such theory—conceptual or functional role semantics—and consider its application to the moral case.[1] I proceed as follows. In Section 1, I provide definitions of key terms. In Section 2, I motivate a conceptual or functional role approach to moral concepts. In Section 3, I elucidate some of the decisions a defender of this theory must face, resulting in several distinct versions of the theory. In Section 4, I argue that all but two of these versions are implausible. In Section 5, I consider one of the remaining versions —elucidated by Wedgwood and tentatively endorsed by Enoch and Eklund—and ask whether it can be plausibly combined with the further claim that moral concepts refer to robust properties. I argue that it cannot. I conclude that the only plausible version of conceptual role semantics for moral terms is a version of metaethical non-cognitivism: according to this view, moral concepts are individuated in terms of functional ties to motivating states.

1. Definitions

This paper is concerned with whether conceptual role semantics can provide a plausible account of moral concepts and, if so, how this might affect existing debates in metaethics, specifically debates between cognitivism and non-cognitivism. The underlined terms need further explication, which will occupy the rest of this section.

2

1.1. Cognitivism and Non-cognitivism

This debate concerns the status of moral states of mind, such as my judgement that murder is wrong. According to cognitivists, moral judgements are maps of a moral reality: they are beliefs about the way of the world, morally speaking. Hence moral predicative concepts, such as **good** and **evil**, are taken to refer (or putatively refer) to robust properties in the world.[2] Non-cognitivists deny that moral judgements are moral beliefs, claiming instead that they are distinctive insofar as they involve a particular type of motivational state, akin to a desire or preference.[3] Hence, non-cognitivists also deny that moral predicative concepts refer to worldly properties; instead, they are distinguished by their connection to motivational states. Neither cognitivism nor non-cognitivism *per se* is committed to the existence (or non-existence) of a moral reality which our judgements might correctly describe: theirs is a psychological rather than metaphysical dispute. Even so, the most popular form of cognitivism *is* realist in this way. For this paper the concern is with the psychological dispute, not the related metaphysical one.

1.2. Concepts

Roughly, concepts are parts of thoughts. Precisely, a concept is an element of the semantically evaluable content of a propositional attitude that contributes to that content in a systematic way. Relatedly, where a propositional attitude has such a content we can say that that attitude involves or deploys its constituent concepts.[4] For example, a belief that grass is green is a propositional attitude with the propositional content *grass is green*, which is semantically evaluable as true or false.[5] The concepts **grass**, **green** and **is** (more precisely, the **'is' of predication** or is_p) are elements of this content (and thus are involved in the mental state) and contribute to it and other contents in systematic ways. For example, the presence of is_p in an unembedded position in this content indicates that the content concerns the attribution of a (putative) property to a subject.

Many concepts are associated with Fregean senses or modes of presentation (the exceptions are concepts such as conjunction and, perhaps, proper names).[6] In other words, many concepts specify ways in which objects, properties or space-time regions can be represented in thought. Imagine, for instance, that the lawn before me is mine, green, the prettiest in the neighbourhood and the cause of my uncle's heart attack. To think of the lawn under these different descriptions is to think of it using different concepts. Being thus at the level of Fregean senses, concepts are individuated by considerations of cognitive significance. Thus two concepts are distinct iff. two complete propositional contents that differ only insofar as they involve replacing one concept for the other can differ in informativeness.[7] For instance, today it is news to me that my lawn is on the route of a proposed bypass, whereas it is not news to me that my lawn is green. Hence, **on the route of a proposed bypass** and **green** are distinct concepts.

It is worth noting that concepts can be shared by speakers of different languages and that not all significant linguistic expressions express concepts. Greetings such

as 'Hiyo!' and exclamations such as 'What Ho!' are significant, yet do not express concepts. The reason is that such utterances do not express attitudes with semantically evaluable contents, and concepts are parts of such contents. It follows that a complete theory of concepts cannot be the sole basis for a complete theory of linguistic meaning.[8] Note finally that the preceding explication of concepts is not intended solely to reflect common usage of the word 'concept'; it is partly stipulative, intending to delineate a notion appropriate to the goals of philosophical and psychological theory.[9]

1.3.   Moral Concepts

Concepts are parts of thoughts and moral concepts are parts of moral thoughts. More particularly, they are the parts of moral thoughts that make those thoughts distinctively moral or, more simply, they are distinctively moral concepts. That these remarks are unhelpful is to be expected, for one of the issues at stake in the current context is the correct account of moral concepts. One way of avoiding this paradox of inquiry is to use exemplars. If anything is a moral concept, then the concept expressed by sincere assertoric use of the English expression 'morally right' is. When Theresa May, say, claims that giving harsh judicial sentences to rioters is the morally right thing to do, there is little doubt that she is expressing a moral thought. What makes this thought moral is the concept of moral rightness. What is required is an account of *that* sort of concept and its kin (although the resulting account may force us to rethink where we draw the kinship boundaries). Hence, by moral concepts I shall mean concepts of the type (whatever it is) expressed by Theresa May's sincere use of the expression 'morally right'.

1.4.   Conceptual Role Semantics

As understood here, conceptual role semantics (henceforth 'CRS') is a general account of the intentional content of concepts and thoughts, that is, a theory of mental representation. Related views are sometimes presented as (foundational) accounts of linguistic meaning or more generally as theories of 'representations', a term intended to cover both types of case.[10] However, the focus in this paper will be with CRS as an account of the content of concepts and thoughts and in particular as an account of the content of moral concepts and thoughts. I leave open the question as to how this might relate to an account of linguistic meaning.[11]

In its most basic form, CRS claims:

> The contents of concepts and thoughts are determined by their functional role in a person's psychology. (Harman 1982: 242)

Where 'functional role' includes a state's actual and hypothetical

> … role in thinking, problem solving, deliberation and the like—and, in general, in mediating between sensory inputs and behavioural outputs. (Block 1987: 160)

Here, 'deliberation' includes both theoretical and practical reasoning: the former being the process of deriving or inferring mental states from others, where this ends in belief, and the latter being the process of deriving or inferring mental states from others, where this ends in intention or action.[12]

The general idea behind CRS is as follows. A particular intentional state will play a particular role in an agent's psychology, both in the actual case and in hypothetical cases where it is grouped with other, now absent, mental states.[13] It may have typical causes in perception and/or typically lead, via processes of theoretical or practical reasoning, to further mental states. It may make rational, provide evidence for, or itself be evidentially supported by (the contents of) other mental states. It may affect the agent's behaviour in predictable ways, the precise effects depending on the states it is coupled with. CRS claims that some proper set of these relations to perception, other mental states and action—that is, the state's 'conceptual role'—determines the intentional content of the state. For concepts, content is also determined by conceptual role, where the conceptual role of a concept (as opposed to the conceptual role of an intentional state) is determined by the systematic contribution it makes to the conceptual roles of the intentional states in which it occurs.[14] According to this view, two concepts are identical iff. replacing one for the other in a propositional attitude results in an attitude with an unchanged conceptual role. According to CRS, therefore, possessing a concept is a matter of possessing mental states that conform to certain conceptual roles: mastering the concept involves mastering these roles, that is, at a minimum, having a general (perhaps defeasible) disposition to follow them.[15]

2. Motivating CRS for Moral Concepts

This characterisation of CRS will need to be refined later, in light of the several choices which defenders of the view face (Section 3). But it is already sufficient to see that the prospects for a CRS account of moral concepts look bright. This is for two general types of reason. First, because many of the general arguments in favour of CRS apply to the moral case. Second, because there are particular features of moral concepts that make them amenable to a CRS-based account. I take these types of arguments in turn.

2.1. General Arguments for CRS

2.1.1. The 'There-Must-be-More-to-Content-Than-Reference' Argument

Reference (in the case of concepts) and truth-conditions (in the case of propositional contents) are too crude for differentiating content. For example, the concepts **Hesperus** and **Phosphorus** have the same reference but distinct content, as shown by usual tests of cognitive significance: one can believe that Hesperus is bright without believing that Phosphorus is bright. Likewise, *Hesperus is a planet* and *Phosphorus is a planet* have the same truth-condition—they are true iff. the thing

we now call 'Venus' is a planet—but differ in cognitive significance. Hence, reference and truth-conditions are insufficiently fine-grained to provide a complete account of the differences between the members of these pairs of concepts and thoughts.[16] If connections with objects and conditions in the world cannot explain the difference between these pairs of concepts and thoughts, it seems that connections within the head must. The obvious account is, therefore, that the members of these pairs are distinct because they have different conceptual roles 'inside the head'. For example, **Hesperus** contributes to conceptual role by tying intentional states which contain it to thoughts of the evening, whereas **Phosphorus** contributes by providing connections to thoughts of the morning. Likewise, believing that *Hesperus is a planet* ties to the belief that the morning star is a planet, whereas believing that Phosphorus is a planet does not (ceteris paribus). Hence, the difference in conceptual role provides for the difference in content.

This argument extends to the moral case. The concepts **German Chancellor from 1933 to 1945** and **most evil individual of the 20th century** refer to the same person but differ in cognitive significance. The best explanation of this is their differing conceptual roles inside the head: thoughts containing the former are tied to thoughts of Germany (for example), whereas thoughts containing the latter are tied to thoughts of moral reprehensibility and (perhaps) certain emotions. Whatever the precise differences in conceptual role, it seems that some such difference is the best explanation of the difference between these co-referring concepts. Hence, moral concepts are, at least in part, distinguished by their conceptual role.

2.1.2.   The Argument from Understanding

This argument can be stated as follows.[17] First, a theory of concepts is a theory of what it is to possess concepts, that is, to grasp or understand them. If we want a theory of **grass,** for example, we need to examine thinkers who understand **grass** and ask what distinguishes them from those that do not. As Peacocke (1992: 5) puts it: '[t]here can be nothing more to the nature of a concept than is determined by a correct account of the capacity of a thinker who has mastered the concept to have propositional attitudes to contents containing the concept'. To deny this claim (Peacocke calls it the 'Principle of Dependence') is to hold the implausible view that two agents with exactly the same propositional-attitude-forming capacities (that is, who instantiate the same possession conditions) might nevertheless be employing different concepts.[18] Second, understanding or grasping a concept consists in having an appropriate functional organisation, that is, mastering the relevant conceptual role. For example, one will not count as possessing **grass** unless one is generally disposed to accept that if an ostended object is grass, then it is a plant (ceteris paribus). As Millar (1994: 74) puts it 'possession conditions … reflect the actual belief-forming and belief-adjusting practices of those who possess the relevant concepts'. Hence, the argument proceeds: a theory of concepts is a theory of concept-possession; concept-possession involves elements of the agents'

thoughts instantiating certain conceptual roles; hence, a theory of concepts is a theory of conceptual roles.

Again, this argument applies to the moral case. If we want a theory of **evil,** for instance, we need to examine thinkers who possess this concept. Further, one will not count as possessing it unless one is disposed (ceteris paribus) to accept certain inferences—that evil is not good, for example—and also (perhaps) to have certain attitudes (to disapprove of evil, for example). Thus, this argument supports CRS for moral concepts.

2.1.3. The Argument from Charity

A final general argument concerns the principle of charity. In the present context, this is the claim that in interpreting the thoughts of another 'we cannot rationally attribute irrationality without limit'.[19] We cannot, for example, suppose that an agent possesses the concept **grass** and has no disposition whatsoever to make rational inferences from propositional attitudes whose contents contain it (of course, such dispositions may be in place yet defeated). CRS explains this principle: since, according to CRS, grasping the concept **grass** requires mastering the conceptual roles that define it, an agent who possesses no relevant dispositions cannot be considered to have mastered these roles, hence does not possess the concept. More generally, since concept-possession requires mastering conceptual roles, agents who haven't mastered the roles cannot be counted as possessing the concept, so a fortiori cannot be interpreted as consistently deploying the concept in an irrational way. Once again this argument applies, mutatis mutandis, to the moral case.

2.2. Particular Arguments for CRS for Moral Concepts

The second group of reasons for favouring a CRS approach to moral concepts concerns particular features of those concepts. As Peacocke (1992: 4, 36, 147) notes, a successful theory of a particular concept should be able to explain all the phenomena that are distinctive of it. That CRS seems apt to explain some well-discussed features of moral concepts therefore mitigates in its favour.

2.2.1. Irreducibility of Moral Concepts

The first particular feature of moral concepts that CRS can explain is their irreducibility to non-moral concepts. That **good** and other moral concepts cannot be defined in any non-moral terms is commonly taken to be one conclusion of G.E. Moore's open question argument (the other is that moral concepts denote 'simple and indefinable' properties). On Moore's view, for any proposed definition of **good**, 'it may always be asked, with significance, of the complex so defined, whether it itself is good' and the fact that this question is open—that is, not settled by an understanding of its terms—shows that the proposed definition cannot, after all, define **good**.[20] The point generalises to all moral concepts. What many

philosophers take this to show is that moral concepts cannot be defined, that is, that one cannot offer a necessarily true universally quantified biconditional of the form 'x is M iff. …', where M is a moral concept that the right-hand side gives, in non-moral terms, necessary and sufficient conditions for x to be M.[21]

The irreducibility of moral concepts is not universally accepted. Some hold that moral concepts can be given non-obvious conceptual analyses, others that they can be reduced to normative concepts, such as the concept of a reason (although many of these hold that normative concepts themselves are irreducible).[22] For present purposes it is not necessary to adjudicate this disagreement. This is because *if* moral concepts are irreducible, then a CRS approach to those concepts can explain why. Likewise, *if* moral concepts are reducible, CRS can explain that too. This is an advantage of the account.

To defend these claims, it is necessary to consider the general connection between CRS and conceptual analysis. The key point is that a CRS account of a given concept is compatible with, but does not necessitate, the availability of a traditional-style conceptual analysis of that concept. (Where a traditional-style conceptual analysis is, roughly, a decompositional analysis that yields a finite set of independently graspable and relatively precise necessary and sufficient conditions for the analysandum to apply, and which thus maintains the truth-conditions and cognitive significance of the analysandum; roughly, the type of analysis that has proved so elusive in the case of **knowledge**). The difference between the two types of case concerns the nature of the conceptual roles that define a concept: in some cases these will allow, and explain, the availability of a traditional conceptual analysis; in others they will not. In the first instance, CRS will be compatible with a traditional conceptual analysis when the conceptual roles that define a concept are such that mastering them amounts to (implicit or explicit) knowledge of a theory of the (independently graspable, relatively precise) necessary and sufficient conditions for the application of a concept.[23] Imagine, for instance, that **knowledge** can be analysed into **justified true belief**, the latter being individually necessary and collectively sufficient for the former. CRS can explain why such an analysis is available, for it will propose conceptual roles definitive of **knowledge** that have the following form:

(C1)   Apply **knowledge** to x only if one applies **justified, true** and **knowledge** to x.
(C2)   Apply **justified** to x if one applies **knowledge** to x.
(C3)   Apply **true** to x if one applies **knowledge** to x.
(C4)   Apply **belief** to x if one applies **knowledge** to x.

Mastering these roles amounts to (implicit) knowledge of the necessary and sufficient conditions for x to count as knowledge. Hence, when (and only when) the conceptual roles that constitute a concept have this form, a traditional conceptual analysis will be available—its content derivable from the concept-constitutive conceptual roles.

In the second type of case, however, the conceptual roles definitive of a concept will not generate a traditional-style conceptual analysis. This will be the case whenever the concept-constitutive conceptual roles are not such that mastery of

8

them amounts to knowledge of the (independent, precise) necessary and sufficient conditions for the application of the concept. Suppose, for instance, that mastery of a concept requires grasping that the concept applies to certain prototypical case and a disposition to apply the concept to any object judged relevantly similar with respect to that concept, but where the relation of 'relevant similarity' cannot be systematised in a concise, precise and independent set of necessary and sufficient conditions (and, perhaps, admits of borderline cases). In that case, no traditional-style conceptual analysis will be forthcoming, since there will be no way of specifying traditional-style necessary and sufficient conditions for the concept to apply.[24] More generally, there are many types of conceptual roles other than those neatly linking the application of a concept to the application of other concepts such as exhibited in the form (C1)–(C4). These include conceptual roles that involve judgements about prototypes, relations of family resemblance, connections to perceptual states and connections to motivational states. In all these cases, no traditional-style conceptual analysis will be possible.[25]

This general point includes the moral case as an instance. If moral concepts are irreducible then CRS can explain why: because the conceptual roles mastery of which constitutes possession of them are not of the (C1)–(C4) form. This may be because they are of a 'prototypical' form, explicated in terms of family resemblance, involve connections to perception, motivational states or some combination of these. On the other hand, if (contrary to Moore) moral concepts are definable, then CRS will be able to explain this too: by giving an account of them in terms of conceptual roles that have the (C1)–(C4) form. Either way, therefore, CRS can explain the (un)availability of conceptual analyses of moral concepts. This is a virtue of the account.

2.2.2.   Practicality of Moral Concepts

A second more particular argument for applying CRS to moral concepts is closely related to the first. Many have argued that the reason why moral concepts are irreducible to non-moral concepts is that the former embody a certain *practicality*, or connection to directing action, that the latter cannot supply. As Johnstone puts it, an otherwise rational individual, when faced with the claim that an action is unpleasant, say, can always ask 'So what?', indicating indifference to the alleged fact.[26] But the same agent cannot rationally display the same indifference when faced with the claim that an action is wrong. Moral epithets seem to go some way to answering practical questions which all non-moral descriptions leave open.[27] This practicality would certainly explain why the Moorean questions remain open: since non-moral concepts are compatible with the practical indifference that moral concepts rule out.

There is, of course, much dispute about how precisely the practicality of moral concepts is best explicated.[28] However, many philosophers hold that the best way of elucidating this claim is in the form of some sort of motivational internalism, that is, as the claim that there is some interesting conceptually necessary connection between applying a moral concept to an object and an appropriate (defeasible)

motivation to act accordingly.[29] By a 'conceptually necessary' connection, defenders of internalism usually mean that it is a connection guaranteed by a proper mastery of moral concepts, and such that failure to instantiate a connection disqualifies an agent as a competent user (hence possessor) of moral concepts.[30] Again, not everyone accepts some form of motivational internalism. But again CRS need not take a stand in this dispute. For *if* motivational internalism is true, then CRS can explain why, and *if* motivational internalism is false, CRS can explain that too.

First, accepting motivational internalism, CRS will claim that the conceptual roles that individuate moral concepts are those that tie the application of those concepts to certain motivational states, such as desires, preferences, plans or intentions. So whereas, for example, the presence of **grass** in a propositional content might tie the relevant intentional state to a *belief* containing the concept **plant**, the presence of **good** in a propositional content might tie the relevant intentional state to a *desire* to promote or perform.[31] Second, if motivational internalism is denied, CRS can claim that moral concepts generate connections only to belief-like (or, more precisely, non-motivating) states. Either way, therefore, CRS can explain the practical import of moral concepts. This is another virtue of the account.

Finally, suppose one accepts both the irreducibility of moral concepts and their practicality, cashed out in terms of (some version of) motivational internalism. The CRS-based explanations of these features provide a certain coherence. A CRS account of moral concepts that involves conceptual ties to motivational states explains why moral concepts are motivationally engaged and this in turn explains why moral concepts cannot be reduced to non-moral concepts. Conceptual ties to motivational states explain motivational internalism, which in turn explains irreducibility. If so, then a CRS approach to moral concepts has the considerable attraction of providing a unified explanation of both their practicality and irreducibility.

3. Varieties of CRS for Moral Concepts

CRS is not a single theory but a family of similar theories, all of which accept conceptual role as an essential component of the identity of concepts but disagree over the precise nature of these roles. Below I set out some of the dimensions of disagreement, which will subsequently allow us to distinguish several versions of CRS for moral concepts.[32]

3.1. Conceptual Roles: Wide or Narrow

So far I have focused on conceptual roles that involve connections between intentional mental states of a thinker, such as her beliefs, desires and intentions. These may be called 'narrow' or 'short-arm' conceptual roles: they concern conceptual role 'inside the head' or 'from the skin in'. But versions of conceptual role semantics may also include 'wide' or 'long-arm' conceptual roles in their accounts of some

concepts. These wide roles specify connections between intentional mental states and objects, properties or conditions 'outside the head', in the wider world.[33] Someone might claim, for instance, that part of the conceptual role definitive of **water** is that predicative uses of this concept in thought are dominantly (or normally) caused by instances of the chemical compound H_2O. An agent whose use of the concept does not obey this rule would not count as possessing our concept of water. (This might explain, for example, why apparently similar thoughts had on Earth and on Twin Earth have different truth conditions; see Putnam 1975 and Burge 1979.)

Several points are worth noting here. First, some may baulk at describing causal relations between objects and thoughts as 'conceptual roles'; perhaps the term 'functional role' is more appropriate. But this is unproblematic, since the term 'conceptual role' is here being deployed as a term of art. Second, wide conceptual roles are not inferential roles (since inference is a relation between mental states) and are typically non-normative. In fact, as Block (1987: 178) points out, wide conceptual roles typically 'look like a traditional theory of reference', hence involve relations such as causation, asymmetric dependence and so on.[34] Third, wide conceptual roles can be incorporated into CRS in at least two ways. On the first view—the 'two-factor view'—some concepts and thoughts have two distinct yet complementary types of content: 'narrow content' determined by narrow conceptual roles (and shared, for example, by my twin on Twin Earth) and 'wide content' determined in part by wide conceptual roles (and not shared by my twin).[35] On the second view—the 'non-solipsistic view'—some concepts and thoughts have one type of content, but that content is determined by a combination of narrow and wide conceptual roles.[36] I will not distinguish further between these two views here. Both accept that wide conceptual roles need to be included in a complete account of the contents of some thoughts and concepts. I leave open the vexed question of whether or not such views entail a separable notion of narrow content.[37] Fourth, I will share the common assumption that where an account of a concept includes wide conceptual roles, these are in addition to narrow conceptual roles.[38] Consequently, the choice for a defender of CRS is whether to incorporate wide conceptual roles alongside narrow roles or whether to include narrow roles only. Fifth and finally, there is the issue of how the two types of roles are related. As Fodor (1987: 80–2) puts it: what glues the two factors together? Again, I shall share a common assumption on this issue: that where an account of the concept involves wide conceptual roles, precisely which roles are involved is determined by the narrow conceptual roles, perhaps together with contingent facts about the thinker's environment.[39] So, for example, the narrow conceptual roles associated with **water** determine a reference-fixing description along the lines of as 'the reference of **water** is whatever property it is that is the (dominant) cause of predicative use of that concept in the actual world'. As things contingently stand, this is H_2O (things are different on Twin Earth). Hence, the narrow conceptual roles, plus this contingent fact about the world, determine that the connection between instances of H_2O and applications of the concept **water** is part of the latter's wide conceptual role.

3.2.   Narrow Conceptual Roles: Normative or Non-normative

A second dimension that can be used to distinguish CRS-based accounts of concepts concerns the normativity of narrow conceptual roles. Consider the following conceptual role linking intentional states S1 and S2.

(S)   S1 → S2

On normative versions of CRS we can read '→' as 'makes rational' or 'rationally commits one to' or 'makes rational if S1 is rational' or 'makes rational the inference to'.[40] On this view, (some) conceptual roles specify normative relations between intentional states. Take the belief that *grass is green*. On one normative version of CRS, this thought is partly individuated by the fact that being in this state *rationally commits one* to believing that grass is coloured. On all versions of CRS, possessing a concept requires mastering its conceptual roles; on normative versions of CRS, such mastery is itself (partly) a normative matter.[41]

On non-normative versions of CRS, by contrast, we can read '→' as 'causes' or 'produces via an inferential process' or 'would, in appropriate conditions, produce via an inferential process' or 'produces via a primitively compelling inferential process' or 'is taken by the agent to obviously make legitimate the inference to …' or some combination of these. On this view, (narrow) conceptual roles specify non-normative relations between intentional states, such as inferential relations.[42] Take the belief that *grass is green*. On one non-normative version of CRS, this thought is partly individuated by the fact that agents in this state will find the inference to the belief that *grass is coloured* primitively compelling (in appropriate circumstances).[43] On non-normative versions of CRS, mastering a concept requires a (general, defeasible) disposition to follow its content-determining conceptual roles, where this is a non-normative matter.

One potential problem with normative versions of CRS is that they cannot provide the basis for a reductive account of normativity. A potential problem for non-normative versions of CRS is that they cannot capture the plausible claim that mastering a concept is in part a normative matter.[44] A potential problem with any version of CRS that employs the notion of inference in describing conceptual roles is that it cannot provide a basis for a reductive account of intentionality (or at least, cannot insofar as 'inference' is spelt out in intentional terms).[45] I will not attempt to adjudicate these disputes here. Instead I will assume that both normative and non-normative versions of CRS, including versions that employ the notion of inference, are viable for application to moral concepts. Consequently, the lessons I attempt to draw for applying CRS to moral concepts will not involve taking sides in any of these disputes.

3.3.   Narrow Conceptual Roles: Beliefs or Desires

A third way of distinguishing CRS-based accounts of individual concepts is in terms of whether the narrow conceptual roles in predicative uses involve connections with belief-like states or with desire-like motivating states (or both). For example, a defender of CRS might claim that the belief that *grass is green* is given

its content in part by a conceptual connection to another belief, viz. the belief that *grass is coloured*. If all of the (narrow) conceptual roles that define a thought are connections to beliefs in this way, then we can call the CRS-based account of the thought 'purely doxastic'. Likewise, a concept is purely doxastic if its contribution to the conceptual roles of thoughts (in predicative uses) is only ever to connect those thoughts with beliefs. Alternatively, someone might hold that the thought that *torture is wicked* is given its content in part by a conceptual connection to a non-belief-like state, such as approval or disapproval of torture (the former if 'wicked' means 'great' the latter if it means 'evil'). If some of the narrow conceptual roles that define a thought are connections to motivating states in this way, then we can say that the account of the concept is 'conative'. Likewise, a concept is 'conative' if it contributes connections to such states. Note that under this terminology, 'conative' accounts will also include mixed cases, that is, cases which involve connections both to belief-like and to motivating states.

4. Narrowing Down the Alternatives

CRS-based accounts of concepts can therefore be distinguished in at least three dimensions: according to the width, normativity and conative nature of their conceptual roles. Defenders of CRS in general need not apply exactly the same type of theory to all concepts: they can, for example, allow that some concepts are differentiated by wide and narrow conceptual roles, while others are differentiated by narrow roles only. But for a given moral concept, there are at least eight variations of conceptual role semantics to consider:

(1) Wide and narrow conceptual roles; narrow conceptual roles are normative; narrow conceptual roles are purely doxastic.
(2) Wide and narrow conceptual roles; narrow conceptual roles are normative; narrow conceptual roles are conative.
(3) Wide and narrow conceptual roles; narrow conceptual roles are non-normative; narrow conceptual roles are purely doxastic. (Cornell realism)
(4) Wide and narrow conceptual roles; narrow conceptual roles are non-normative; narrow conceptual roles are conative. (Australian realism)
(5) Narrow conceptual roles only; narrow conceptual roles are normative; narrow conceptual roles are purely doxastic.
(6) Narrow conceptual roles only; narrow conceptual roles are normative; narrow conceptual roles are conative. (Wedgwood)
(7) Narrow conceptual roles only; narrow conceptual roles are non-normative; narrow conceptual roles are purely doxastic.
(8) Narrow conceptual roles only; narrow conceptual roles are non-normative; narrow conceptual roles are conative. (Gibbard, Blackburn, Enoch)

Some of these theories of moral concepts are familiar from the existing literature in metaethics. One version of (3) is elsewhere known as 'Cornell Realism'. This

theory models moral concepts on natural kind concepts such as **water**, which, we saw above, are elucidated in terms of narrow plus wide conceptual roles.[46] Assuming conceptual roles reflect elements in a traditional-style conceptual analysis, (4) includes versions of 'Australian' realism offered by Jackson and Smith.[47] Several versions of (8) have been defended by non-cognitivists, who argue that moral concepts are conceptually tied to motivational states such as states of planning or particular kinds of approval and disapproval.[48] Wedgwood has defended a version of (6), which I discuss below (Section 5).

Several of these versions of CRS for moral concepts face serious problems. Consider first theories that specify narrow conceptual roles that are purely doxastic, that is, (1), (3), (5) and (7). These theories will have trouble accommodating the practicality of moral concepts, discussed above (Section 2). At least, if we accept that moral concepts are interestingly conceptually tied to relevant motivation, then any theory that elucidates those concepts solely in terms of conceptual connections to non-motivational belief-like states will by that token be defective: such theories cannot accommodate motivational internalism.

Consider next theories that include wide conceptual roles alongside narrow roles, that is, (1), (2), (3) and (4). These theories face the problem that precisely the sort of case that motivates the inclusion of wide conceptual roles in the case of (some) non-moral concepts seems to motivate the *exclusion* of such roles in the moral case. The case is Twin Earth.

In the case of non-moral concepts, such as **water**, the argument runs as follows. Me and my twin on Twin Earth are exact physical duplicates of each other, at least from the skin in. The difference between us is that I live in a world where the (dominant) cause of thoughts about water is H_2O and she lives in a world where the dominant cause of thoughts about (twin-)water is XYZ. (The bracket is necessary to leave open the issue whether my twin employs the concept we express by the word 'water'). The watery thoughts of me and my twin have different truth-conditions. For example my thought that *water is wet* is true iff. H_2O is wet, whereas my twin's thoughts that (*twin-)water is wet* is true iff. XYZ is wet. Since truth-conditions are determined by content, our thoughts must have difference contents. But since, by hypothesis, me and my twin are identical in terms of intracranial functional organisation ('inside the head'), the difference in the contents of our thoughts must be determined by relations outside the head, for example, by our causal relations to properties in the world. Hence,—as the slogan almost goes—content ain't (purely) in the head. Hence, in the case of **water**, Twin Earth motivates the inclusion of wide conceptual roles.[49]

Compare Moral Twin Earth.[50] As before, my Moral Twin is an exact duplicate of me from the skin in. The difference this time is that I live in a world where the (dominant) cause of thoughts about **rightness** is the consequentialist property C, whereas she lives in a world where the (dominant) cause of thoughts about **rightness** is the distinct deontological property D. Question: Do our rightness-thoughts have different contents? Suppose my thought that Φ is right and my twin's thought that Φ is right have different contents, hence (assuming that content determines truth-conditions) different truth-conditions. Call this 'the supposition'. Suppose,

for instance that my thought that Φ is right is true iff. Φ is C, and my twin's thought that Φ is right is true iff. Φ is D. Now imagine that I think that x is right and my twin thinks that x is right. According to the supposition, any agreement between us is merely superficial: although it appears as if we are agreeing, we are in fact at cross purposes, having distinct thoughts. Indeed it may be that one of our thoughts is true and the other false (if x is C but not D, for example). Likewise, imagine I think that y is right and my twin thinks that y is not right. According to the supposition, any disagreement between us is merely superficial. In particular, since our thoughts about rightness have different truth-conditions, both our thoughts could be true (if y is C but not D, for example). Both these consequences are unpalatable. Intuitively, me and my twin are in *agreement* about the morality of x—specifically, about its rightness—and in *disagreement* about the morality of y. We should reject any supposition that entails otherwise. Hence, we should reject the supposition that my thoughts and my twin's thought about rightness have difference contents and (hence) different truth-conditions.[51]

Thus, in the moral case, not only does the argument in favour of including wide conceptual roles not go through but the Twin Earth scenarios also suggest that moral content *is* in the head. For if agreement and disagreement about the rightness of Φ require that me and my twin have the same concept of rightness, then these facts about our agreement and disagreement show that our concepts are the same.[52] But then it seems that the best explanation of this similarity in content is similarity of functional roles inside the head. Hence, for moral concepts, meaning just is in the head.[53] If so, any account of moral concepts that includes content-determining relations outside the head—that is, that includes wide conceptual roles—is defective. Hence, we should reject theories (1)–(4).

5. Wedgwood and Robust Reference

Consider, finally, theories (6) and (8). These have distinct advantages. The connection to motivational states allows these views to accommodate the practicality and irreducibility of moral concepts. That same connection allows an explanation of the intuitions at work in the Moral Twin Earth argument. Those intuitions were that me and my twin were in substantive agreement about the rightness of x and in substantive disagreement about the rightness of y. A connection to motivational states locates this agreement and disagreement. In the first case, we are in agreement about x because we have the same attitude towards it (and, perhaps, no motive to change the other's attitude) and in the second case are in disagreement about y because we have opposed attitudes towards it (and, perhaps, motive to change the other's attitude). Thus, a connection to motivational states, such as that posited by theories (6) and (8), explains the thought that me and my twin agree about the rightness of x and disagree about the rightness of y.

I will not consider the differences between (6) and (8), that is, the issue of the normativity of narrow conceptual roles. Instead I will address the issue of the consequences of adopting such a CRS-based account of moral concepts. In particular I

want to answer the following question: Can a CRS account of moral concepts, along the lines of (6) or (8), be plausibly combined with cognitivism about moral concepts? My answer will be 'No'. I conclude that the most plausible form of conceptual role semantics for moral concepts supports non-cognitivism in metaethics.

Wedgwood defends a version of (6), and Enoch (2011: 177-84) adheres to Wedgwood's programme but for his rejection of the claim that 'the intentional is normative'; that is, Enoch adopts a Wedgwoodian version of (8). Wedgwood's focus is on 'thin' moral concepts, specifically the concept expressed by 'x is (all things considered) a better thing for z to do at time t than y' (in Wedgwood 2001) and the concept of **practical ought** (in Wedgwood 2007). I will focus on the former, although the points I make apply equally to Wedgwood's account of the latter.

According to Wedgwood, concepts are individuated by the basic rules of rationality governing their use.[54] These rules are basic insofar as they are 'fundamental truths of rationality, which cannot be explained on the basis of any more basic principles about how it is rational to use the concept … in thought'.[55] For the concept **x is (all things considered) a better thing for me to do at time t than y**—which can be abbreviated to **B(x,y,me,t)**—the sole basic rule of rationality is the following rule of practical reasoning:

(R1)   Acceptance of $B(x,y,me,t)$ commits one to having a conditional intention to do x rather than y.[56]

Where a 'conditional intention' is an intention to do x rather than y if one does either.[57] According to Wedgwood, one intentional state 'commits' one to another intentional state just in case, if being in the first state is rational then so is being in the second (at least if the question of whether or not to adopt the second state arises).[58] So we can read (R1) as:

(R1*)   If acceptance of $B(x,y,me,t)$ is rational then a conditional intention to do x rather than y is rational.

In common with all versions of CRS, Wedgwood holds that possessing a concept individuated by certain roles involves mastering those roles.[59] In particular, he holds that possessing a concept requires a non-derivative mastery, or disposition to follow, the basic rules of rationality that define the concept. A mastery or disposition to follow a set of rules involves generally following those rules non-accidentally (where this involves, at some level, some grasp of the fact that these are the basic rules) and is also (for Wedgwood) partly a normative matter.[60] Hence, for Wedgwood, possessing a concept is not just a matter of actually following certain conceptual roles, it is a matter of it being the case that one rationally ought, in some sense, to instantiate those roles. Finally, a mastery or disposition to follow a rule is non-derivative just when it is not due to having some independent justification for following those rules.[61]

Wedgwood does not rest here. He claims that the conceptual roles definitive of a given concept determine its semantic value, and more particularly, in the case of

moral concepts, the relevant conceptual roles determine that moral concepts refer to robust worldly properties. In this way, Wedgwood combines CRS for moral terms with metaethical cognitivism, and more recently, Enoch has followed this trend.[62] As Wedgwood puts it, a conceptual connection to motivating states manages to conjure up a 'distinctively practical mode of presentation' of a real worldly property.[63] Let us examine this combination of views.

Begin with semantic value. The semantic value of a concept is the contribution it makes to the truth-conditions of the thoughts in which it appears.[64] Hence, the truth-conditions of thoughts are determined by the semantic values of their constituents. Since the constituents of thoughts are concepts, then concepts must determine their semantic values. If we are to give an account of concepts in terms of conceptual roles, then it follows that the conceptual roles associated with a given concept must determine that concept's semantic value.[65]

All this is platitudinous. But, how, precisely, do conceptual roles determine semantic value? To answer this question is to provide a 'determination theory' for a concept (Peacocke 1992: 17–19). More precisely, since in the present context we are concerned only with predicative concepts, how do conceptual roles for predicative concepts determine the semantic values of those concepts?

Consider first a deflationary determination theory. Suppose the truth-conditions of a given thought are given by the deflationary T-schema, thus:

(T) {**x is F**} is true iff. x is F.[66]

The semantic value of the predicative concept **F** is its contribution it makes to the truth-conditions of the thoughts in which it appears. Then, we can say, for a thought of the form {**x is …**}, the contribution made by a predicative concept **F**, when substituted into the '**…**' position, is to generate a sentence which is true iff. x is F. So **F** makes the following contribution to truth-values: it generates a thought which is true when applied to things that are F and a thought that is false when applied to things that are not F (note the disquoting here). Since 'semantic value' is just contribution to truth-conditions, this is **F's** (deflationary) semantic value: the semantic value of **F** is the set of F-things.[67] In short, if truth-conditions are understood in deflationary terms and semantic value is just contribution to truth conditions, then semantic value is also deflationary.[68] Importantly, this account of a semantic value does not require that such concepts be taken to refer to robust worldly properties. We can therefore satisfy the above platitudes about concepts and their semantic values without a robust referential semantics. In particular, we can accept that, if CRS is true, conceptual roles must determine semantic value, without holding that these conceptual roles necessitate reference to robust worldly properties. Understood in deflationary terms, the semantic value of **F** is simply the set of things which are F, that is, the extension of **F**.

Wedgwood's view involves more than the deflationary determination theory. It holds that predicative concepts often refer to or 'stand for' what I am calling robust worldly properties (Wedgwood 2001: 5 cf. Peacocke 1992: 16 ff.; Enoch 2011: 3–5). Hence, the semantic values of these concepts are robust properties. These concepts

are true of their extensions because members of those extensions instantiate those properties. Hence, Wedgwood cannot accept the deflationary theory as the whole story about how concepts determine their semantic values. Consider instead the alternative, 'best sense' determination theory:

(BS)   The semantic value of a concept is the semantic assignment that makes 'correct' or 'makes best sense of' the conceptual roles definitive of that concept.[69]

This needs elucidation. Peacocke explains this as the view that we should assign semantic values in such a way as that the

> … belief-forming practices mentioned in a concept's possession conditions are correct. That is, in the case of belief formation, the practices result in true beliefs, and in the case of principles of inference, they result in truth-preserving inferences. (Peacocke 1992: 19)

Consider the belief or thought that *x is F*, which I shall abbreviate to {**x is F**}. Peacocke is envisaging two types of conceptual roles for this belief. The first is a 'belief-forming role', which specifies a condition, W, not itself a belief, under which the belief is to be formed. That is,

(B1)   W → {**x is F**}

Here, W might be a perceptual experience or a worldly condition. Suppose in this case it is a worldly condition, such as x's being green. Hence, (B1) specifies a wide conceptual role.

The second type of conceptual role that Peacocke considers involves inferential relations between beliefs, such as

(B2)   {**x is F**} → {**x is coloured**}

Peacocke's thought is that we should assign a semantic value to **F** such that in the case of W, the belief that *x is F* is true (this is the requirement provided by (B1)). And such that the inference from *x is F* to *x is coloured* is truth-preserving, i.e. such that in the inference given in (B2), the consequent is true if the antecedent is. A plausible assignment is therefore the property of being green. Hence, the semantic value of **F** is a robust worldly property, viz., the colour green.

Wedgwood modifies Peacocke's theory in two crucial ways. First, by extending it. In (B1) and (B2), the only mental states involved are beliefs. Peacocke's stated theory doesn't cover cases where conceptual roles involve connections to non-belief-like propositional attitudes, such as desires and intentions. Wedgwood extends the theory to cover this sort of case. In particular, Wedgwood claims that we should assign semantic values in such a way as to make the conceptual roles that individuate a concept both *valid* and *complete*.[70] A role linking two intentional states is *valid* just in case it is necessary that if the 'input' state is correct then the

'output' state is uniquely correct. A role linking two intentional states is *complete* just in case the output state is correct *only if* the input state is correct.[71] Crucially, Wedgwood claims, correctness can apply both to motivational states, such as desires and intentions, and to beliefs.[72]

We are now in a position to see how Wedgwood determines the semantic values of moral concepts. Consider **B(x,y,me,t)**, defined by the rule

(R1)   {**B(x,y,me,t)**} → Conditional intention to do x rather than y.

According to Wedgwood, the semantic value of **B(x,y,me,t)** is that property which makes it the case that if instantiated by me, x, y and t, makes it correct for me at t to have the conditional intention to do x rather than y. As Wedgwood puts it:

> The semantic value … must be that four-place relation between, x, y [me] and t, such that, necessarily, it is correct for [me] to prefer doing x over doing y at t and a mistake for [me] to prefer doing y over x at t if, and only if, x, y, [me], and t stand in that relation. (Wedgwood 2001: 19)

And latter, more simply:

> … the reference or semantic value of the [concept **B(x,y,me,t)**] must be the relation of being (all things considered) a better thing to do. (Wedgwood 2001: 21)

Thus, unlike the deflationary determination theory, Wedgwood's view determines that moral predicative concepts have as their semantic values robust properties.

It is helpful to put aside the complexities of Wedgwood's account for the moment and focus the form of its claims. The theory employs a purely narrow concept-constitutive conceptual role—(R1)—to determine a real, robust worldly referent for the related concept. Roles purely inside the head thereby establish a connection between the world and what's inside the head. This is somewhat surprising. As Schroeter and Schroeter (2003: 201-2) point out, if we are to interpret **B(x,y,me,t)** as a genuinely referential concept (that is, as a concept whose semantic value is a robust property), we would imagine that someone who grasps the concept would be disposed (ceteris paribus) to show some sort of sensitivity to the world, specifically some sort of sensitivity to the property that is that concept's semantic value. But Wedgwood's account doesn't demand this, only that the inner conceptual roles are in order. As Schroeter and Schroeter explain:

> A subject who has non-derivative mastery of rule [(R1)] shows no signs of taking the target predicate [or concept] to be answerable to a [robust] normative property. Commonsense interpretative charity militates against importing a [robust] normative element into the semantic assignment when the subject's use of the predicate [or concept] is not answerable to any normative property. [Wedgwood's theory] thus provide[s] the wrong determination theory for an expression governed by [(R1)]: the semantic value of such an expression [or concept] is not a [robust] normative property (2003: 202).[73]

Indeed, this point can be extended to an objection to any view, such as Wedgwood's, that combines a purely *narrow* conceptual role semantics with the additional claim that the concepts so defined have as their semantic values robust properties (for example, Enoch's view). Most generally, it is implausible to suppose that a concept refers to a worldly property when competence with that concept is compatible with no disposition whatsoever to be sensitive to any worldly property.

Something has gone wrong in Wedgwood's account. How can purely narrow conceptual roles inside the head determine a robust worldly referent? Note this is *not* something that Peacocke holds is possible, even though Wedgwood takes Peacocke as his inspiration for his determination theory. For Peacocke's general determination theory allows for the possibility of *wide* conceptual roles (roles that start outside the head)—this is why I noted earlier that Wedgwood makes *two* modifications to Peacocke's account of determination theories.[74] The second modification is to give a general form of a determination theory that *also* covers cases where concepts are individuated in terms of purely narrow conceptual roles. This is what creates the problem, in such cases, of bridging the gap between narrow roles and robust worldly referents.

How might Wedgwood bridge this gap? To the above criticism from Schroeter and Schroeter, he might argue:

> I accept that the account of moral concepts in terms of conceptual roles such as (R1) doesn't *by itself* suffice to interpret those concepts as referring to robust properties. But there are good independent arguments for thinking that moral concepts refer to such properties. Once those arguments are in place, the question is not "*Do* moral concepts refer to robust properties?" but "*Which* robust properties do they refer to?". To this latter question, my general determination theory gives an answer: they refer to those properties which make the conceptual roles that individuate them both valid and complete. Thus my argument has two stages: first establishing in general the claim moral concepts are robustly referential; second triangulating their reference via their content-constitutive conceptual roles.

This is in fact the strategy that Wedgwood adopts (and Enoch follows): in both his defences of CRS for moral concepts Wedgwood notes that there are *independent* arguments for taking the semantic value of moral concepts to be robust properties.[75] But this reply faces an uncomfortable dilemma. Either the semantic value of moral concepts is solely determined by narrow conceptual roles, such as (R1), or it is not. If it is then there is no justification for taking this semantic value to be any worldly property; indeed, there is good reason to think that the semantic value is *not* any worldly property, since competence with robustly referential concepts typically requires some sort of sensitivity to the world (this is Schroeter and Schroeter's point). Alternatively, if the semantic value of moral concepts is *not* solely by narrow conceptual roles then either it is determined by such roles *plus wide conceptual roles* or it is determined by considerations outside of the theory of what it is to possess concepts, that is, outside of the theory of

conceptual roles. The former would involve moving to a theory of CRS for moral concepts that includes wide conceptual roles, which is both implausible (given the Moral Twin Earth argument) and not compatible with Wedgwood's repeated claim that a complete account of moral concepts can be given in terms of purely narrow roles.[76] The latter is also implausible. On this view it is not the case that the content of concepts determines their semantic value; for example, it is not the case that the content of moral concepts determines their robust referents. As we saw above (Section 2), a theory of concepts is a theory of what it is to possess a concept, and according to Wedgwood, what it is to master moral concepts is to master conceptual roles. But mastering such roles does not by itself, on the current view, determine the referent of the concept. It would be possible, for instance, for there to be two concepts, both individuated solely by rule (R1), but one of which refers to a robust property and the other of which does not. This is implausible. It is implausible to think that two concepts could contribute in different ways to the truth-conditions of the thoughts in which they appear, and yet the mastery of those concepts consists in exactly the same propositional-attitude-forming abilities. Most generally it is implausible to deny, as this part of the dilemma must, that the contribution of a concept to truth-conditions is determined solely by its possession conditions.[77]

Some further points about this dilemma may serve to elucidate it. First, the first horn of the dilemma (i.e. the claim that content *is* solely determined by narrow conceptual roles, which *do not* determine a robust referent) seems to be the correct place to sit in the case of the concepts of the logical constants, such as **negation**, on which Wedgwood (perhaps unfortunately for his cognitivism) models his account of moral concepts.[78] The conceptual roles that define the concepts of logical concepts are purely narrow, but it is also the case that these concepts are not understood as referring to robust properties. Although it is sometimes said, in a Fregean spirit, that logical constants 'refer' to truth-functions, we must not confuse this terminology for a substantive account. As Field notes:

> … it is hard to see how to give any account of "reference to" truth functions *unless we construe the claim that "¬" refers to a certain truth-function as really a claim about the role that "¬" plays in our conceptual scheme*. And if we do so construe the claim, then we can see that "reference to truth-functions" is an unnecessary diversion ….(Field 1977: 402)

Second, the first horn of the dilemma is in fact a more general (and more ancient) problem for CRS. This is the problem of the 'word-world' linkage, that is, the problem of how conceptual roles inside the head can link up with properties in the world. It was versions of this problem that in part motivated the move to include wide conceptual roles in a complete account of content. But this is the move, in the moral case, which Wedgwood explicitly rejects. Peacocke, as previously noted, can accept this face-saving move, since he is happy to include wide conceptual roles in his theory of concepts.

In sum, the problem Wedgwood faces is this. The narrow conceptual roles that he takes to define moral concepts are insufficient, by themselves, to determine a

robust referent for those concepts. Wedgwood may argue that robust reference is secured by considerations outside the theory of conceptual possession-conditions, but this would be to sever the plausible connection between the theory of what it is to possess concepts and the theory of what determines their semantic value.

6. Conclusion

I have argued that

- For moral concepts, a plausible CRS (a) excludes wide conceptual roles (b) involves narrow conceptual roles linking moral judgements to motivational states (Sections 2, 4).
- That such an account cannot plausibly be combined with the claim that moral concepts refer to robust properties (Section 5).

This leaves metaethical non-cognitivism as the true beneficiary of conceptual role semantics for moral concepts.[79]

Neil Sinclair
University of Nottingham
UK
neil.sinclair@nottingham.ac.uk

NOTES

[1]   For the application of the distinct theory of teleosemantics to the moral case, see Sinclair (2012).

[2]   I use bold to refer to concepts. By a 'robust property' I mean a property that is genuine constituent of the world, something that will necessarily be quantified over in our best account of what types of thing the world contains. I intend this term to be neutral with respect to debates about the correct methodology for ontology and also be consistent with the view that there are various types of robust property (e.g. natural vs. artificial, reducible vs. irreducible, metaphysically fundamental vs. metaphysically derivative, causally inert vs. causally efficacious and so on). By contrast there is a minimal property of Fness whenever there is a predicate 'F' that can be warrantedly applied to some objects. In the minimal sense, properties may be no more than 'the semantic shadows of predicates' (Blackburn 1993: 8).

[3]   Of course, both beliefs and desires can be implicated in motivation. By referring to 'motivating' states I mean to refer to those that have a connection to action characteristic of states with a desire-like 'direction of fit'. See Humberstone (1992).

[4]   See Wedgwood (2007: 59 ff). The notion of semantically evaluable content comes from Fodor (1987: 11).

[5]   Henceforth I use italics to refer to propositional contents.

[6]   Peacocke (1992: 3); and Wedgwood (2007: 59, 81).

7   Peacocke (1992: 2); and Botterill and Carruthers 1999: 132.

8   See Peacocke (1992: 3, 32–3; 1996: 441). This is controversial. Harman (1982) offers a conceptual role semantics that takes all linguistic meaning to derive from the expression of concepts. This debate is tangential to those considered in this paper, since all sides in the current dispute agree that moral expressions express moral concepts.

9   Peacocke (1992: 2).

10   Field (1977), Horwich (1994),Wedgwood (2001) and Enoch (2011: 177–84) focus on linguistic meaning; Harman (1982), Block (1986), Peacocke (1992), Botterill & Carruthers (1999) and Wedgwood (2007) focus on mental content. Block (1987) notes the possibility of a theory applying to all representations. The following are all logical possibilities: (i) conceptual role semantics applies independently to both mental concepts and items in public languages (cf. Horwich 2005: 7–9, ch. 7); (ii) conceptual role semantics applies directly only to concepts, with linguistic expressions deriving their content (meaning) from expressive relations to concepts or thoughts (cf. Harman 1982); and (iii) conceptual role semantics applies directly only to items in a public language, with concepts deriving their content from relations to these.

11   One obvious connection is provided by the broadly Lockean view that '[t]he meanings of linguistic expressions are determined by the contents of the concepts and thoughts they can be used to express' (Harman 1982: 242; see also Wedgwood 2007: 82–3). For doubts, see Loewer (1982: 308–9) and Davidson (1975).

12   Of course, particular pieces of theoretical and practical reasoning are often parts of some larger reasoning process: see Broome (2009).

13   Botterill & Carruthers (1999: 177). For ease of reading, I will henceforth remove the qualification concerning 'actual and hypothetical' conceptual roles.

14   Field (1977: fn. 19); Block (1987: 164; 1986: 667).

15   Wedgwood (2007: 161–5).

16   For this argument, see Field (1977: 389–90); Loewer (1982: 310–11); Fodor (1987): 73–6; Botterill and Carruthers (1999: 178).

17   Explicit statements of this argument are rare, but see Loar (1982: 276); Loewer (1982: 310); Peacocke (1992: 1–40); and Millar (1994: 74) for intimations.

18   As Peacocke notes, this claim is an analogue in the theory of mental content of Dummett's claim that a theory of linguistic meaning is a theory of understanding: see Dummett (1975). Compare Millar (1994), who cites Wittgenstein as inspiration.

19   Block (1987: 165). See also Wedgwood (2007: 164).

20   Moore (1903: §13); Ayer (1946: 138–9); Snare (1975); Wedgwood (2007: 68–9). Note that Moore's argument is more commonly presented as an argument against reductive definitions of moral terms, but it is here transposed as an argument against reductive analyses of moral concepts.

21   Wedgwood (2007: 66).

22   For the former, see Jackson (1998: 150–3); for the latter, see Scanlon (1998: ch. 1).

23   See Peacocke (1992: 149) and Heal (1996: 414).

24   See Peacocke (1992: 12). Note that some philosophers hold that most, if not all, concepts are 'prototypical' in this way, hence do not submit to traditional conceptual analysis. The history of failed conceptual analyses would seem to support this view. See Botterill and Carruthers (1999: 185), Stich (1992) and Tye (1992). Another example concerns CRS as applied to logical constants, where the very notion of traditional conceptual analysis (that is, as a particular completion of the schema: 'x is F iff …') seems misconceived. See Peacocke (1987) and Wedgwood (2007: 82–8).

25   For this point as it applies to connections to motivational states, see Wedgwood (2001: 27–30); Schroeter and Schroeter (2003: 195). Compare Gibbard (2003: chs. 1–2).

²⁶   Johnstone 1989: 157 ff.

²⁷   This 'practical' open question argument can be traced back at least to Hare (1961: 41–3, 81). See also Jackson (1998: 153) and Wedgwood (2007: 69–72) and Darwall/Gibbard/Railton (1992). Blackburn (1998: 70) employs a similar premise in an argument for expressivism.

²⁸   See Sinclair (2007).

²⁹   See, for example, Svavarsdottir (1999) and Wedgwood (2007: 23–34).

³⁰   See, for example, Smith (1994: 71–6). For doubts, see Merli (2009).

³¹   See Wedgwood (2007: 97) for one particular version of this general form of argument.

³²   One variation not discussed here concerns holistic and non-holistic versions of CRS. The former take *all* of a thought's conceptual roles to determine its content, the latter only a proper subset of those roles. The versions of CRS that I discuss below are non-holistic. For discussion, see Fodor (1987: 55–95), Fodor and Lepore (1992: 163–186) and Botterill and Carruthers (1999: 181–4).

³³   See Block (1987: 170–9); Harman (1987); Loar (1982); Field (1977); McGinn (1982); Peacocke (1992: 113, 185); Horwich (1994).

³⁴   For asymmetric dependence, see Fodor (1987: 107–9).

³⁵   Block (1987: 167–72).

³⁶   Harman (1987).

³⁷   Some have argued that narrow content is problematic, which may mitigate in favour of the non-solipsistic view. For arguments of this kind, see Burge (1979,1986). For defences of narrow content, see Fodor (1987: 27–53) and Botterill & Carruthers (1999: 131–60).

³⁸   See, for example, Field (1977), Block (1987) and Harman (1987). Note that Wedgwood (2007: 61–65) considers and rejects an account of moral concepts that involves only wide conceptual roles. For the sake of argument, this paper accepts Wedgwood's argument.

³⁹   See Block (1986: 643–4); Loar (1982: 279); Schroeter and Schroeter (2003: fn. 2).

⁴⁰   For normative versions of CRS, see Wedgwood (2001, 2007), Millar (1994) and Brandom (1994).

⁴¹   See Wedgwood (2001: fn. 17, and 2007: 84–6, fn. 10, 161–73).

⁴²   See Block (1987: 160–1), Field (1977), Peacocke (1992: 139), Horwich (1994) and Millar (1994). Enoch (2011: 177–184) also seems to favour a non-normative version of CRS.

⁴³   For Peacocke, an inference is primitively compelling for an agent just in case (i) the agent finds it compelling, (ii) the agent does not find it compelling because it can be derived from other premises or principles and (iii) the agent does not treat the veracity of the inference as answerable to anything else (1992: 6, 134–7). For critical discussion and a proposed modification, see Millar (1994: 76–8).

⁴⁴   See Millar (1994: 79–80), Rey (1996: 420–1) and Papineau (1996: 431). See Horwich (1994: 151–2; 2005: 104–133) and Peacocke (1996) for replies.

⁴⁵   Block (1987: 160–1). See Botterill and Carruthers (1999: 185–90) for a response.

⁴⁶   See, for instance, Boyd (1988). Compare Wedgwood (2007: 58–66), which takes Cornell realism to include *only* wide conceptual roles.

⁴⁷   Smith (1994) and Jackson (1998).

⁴⁸   See Blackburn (1998) and Gibbard (2003).

⁴⁹   For the Twin Earth argument, see Putnam (1975). For its extension to non-natural-kind terms (or concepts), see Burge (1979). As a motive for including wide conceptual roles, see Block (1987) and Field (1977).

⁵⁰   The following argument is adapted from Horgan and Timmons (1992).

⁵¹   One could also, of course, reject the idea that content determines truth-conditions and claim that the thoughts of me and my twin have the same content (hence are the same

thoughts) but have different truth-conditions (see, for example, Botterill and Carruthers 1999: 137). I do not want to question this move here, for it accepts the current point, namely, that moral content is in the head (since it is only 'inside the head' that my twin and I are similar).

[52]   The general validity of this move—from disagreement to sameness of concept deployed—has been questioned by Plunkett & Sundell (2013), given the possibility of some cases of disagreement (e.g. metalinguistic negotiations) with no shared concept. However, it is not obvious (and Plunkett and Sundell do not present an argument for thinking that) the particular disagreement in the Moral Twin Earth case is of this type, rather than of the type where disagreement is based in shared concepts. For more general doubts about the value of intuitions regarding disagreement for metasemantic theory, see Dowell (2016).

[53]   One might resist this argument by claiming that agreement and disagreement do not require me and my Twin's concepts to be *exactly the same*, so long as they are sufficiently similar to frame dispute (see, for example, Copp 2000 and van Roojen 2006). But this seems not to capture the full force of the intuitions in the Twin Earth case: for me and my Twin agree (in the first case) and disagree (in the second case) about precisely this: the rightness of Φ. We are not somewhat in agreement about the rightness and somewhat not in agreement.

[54]   Wedgwood (2001: 6, 2007: 82).

[55]   2007: 84. See also 2001: 7.

[56]   It is a little unclear why this amounts to a rule of practical *reasoning*, since the connection between the two mental states is one of 'rationally committing' rather than 'inference'. I put this point to one side here.

[57]   Wedgwood (2001: 15–16).

[58]   Wedgwood (2007: 84).

[59]   Wedgwood (2007: 161–73).

[60]   Wedgwood (2001: 8) uses the terminology of 'mastery'; 2007: 84–6 uses the terminology of 'following'. Both note that the notion is ultimately normative: 2001: fn. 17 and 2007: 86 fn. 10.

[61]   Wedgwood (2001:8 and 2007: 85–6).

[62]   Wedgwood (2001: 18–19 and 2007: 5–6, 99–100); Enoch (2011: 177–84).

[63]   Wedgwood (2007: 80). I use the term 'conjure' because it is unclear how a connection to motivational states manages to provide a distinctive mode of presentation, as opposed to a distinctive connection to motivation.

[64]   Wedgwood (2001: 5). Some non-cognitivists, of course, will prefer a broader conception of semantic value as contribution to context-independent content (or something of that sort), but for the time being I am happy to work within Wedgwood's framework.

[65]   For this argument, see Peacocke (1992: 16); Horwich (1994: 141–51); Schroeter and Schroeter (2003: 192).

[66]   I here express the T-schema in terms of the truth conditions of thoughts rather than sentences. For more standard expressions and defences, see Horwich (1990) and Field (1986).

[67]   This argument is a particular application of a general form of argument given in Horwich (1994: 141–4).

[68]   Alternatively, one might say that, on this view, the semantic values of predicates are deflationary properties, where deflationary properties are just 'the semantic shadows of predicates'; see note 2.

[69]   See Peacocke (1992: 19, 26, 133–43); Wedgwood (2001: 9, 18–19, 2007: 86–7, 99); Enoch (2011: 179); and Schroeter and Schroeter (2003: 192, 195–6).

[70]   Wedgwood (2001: 9–11, 2007: 86–7).

[71]   Wedgwood (2001: 10–11, 2007: 87). See also Schroeter and Schroeter (2003: 196–7).

[72]   Wedgwood (2001: 19–20).

[73]   Wedgwood (2007, 99–100) does consider Schroeter and Schroeter's (2003) critique but not this part of it.

[74]   See Peacocke (1992: 113, 185).

[75]   Wedgwood (2001: 5–6, 2007: 97); Enoch (2011: 177).

[76]   Wedgwood (2001: 15; 2007: 97).

[77]   See Peacocke (1992: 22–3). This view also involves denying that Peacocke calls 'The Identification', namely, that possessing a concept is knowing what it is for something to be its semantic value.

[78]   Wedgwood (2001: 6-12, 2007: 82–9); cf. Enoch (2011: 178–9).

[79]   Fellowship from the Arts and Humanities Research Council, UK. Grant number: AH/I002707/1

REFERENCES

Ayer, A. J. (1946), *Language, Truth and Logic*. London: Gollancz.

Blackburn, S. (1993), *Essays in Quasi-realism*. Oxford: Oxford University Press.

—— (1998), Ruling Passions. Oxford: Oxford University Press.

Block, N. (1986), 'Advertisement for a Semantics of Psychology', in P. A. French, T. E. Uehling and H. K. Wettstein (eds) *Midwest Studies in Philosophy X*. Minneapolis: University of Minnesota Press.

—— (1987), 'Functional Role and Truth Conditions', *Proceedings of the Aristotelian Society, Supplementary Volume*, 61: 157–81.

Botterill, G. and Carruthers, P. (1999), *The Philosophy of Psychology*. Cambridge: Cambridge University Press.

Boyd, R. (1988), 'How to be Moral Realist', in G. Sayre-McCord (ed.) *Essays on Moral Realism*. Ithaca: Cornell University Press: 181–228.

Brandom, R. (1994), *Making it Explicit*. Cambridge, Ma.: Harvard University Press.

Broome, J. (2009), 'Unity of Reasoning', in S. Robertson (ed.) *Spheres of Reason*. Oxford: Oxford University Press.

Burge, T. (1979), 'Individualism and the Mental', in P. A. French, T. E. Uehling and H. K. Wettstein (eds) *Midwest Studies in Philosophy IV*. Minneapolis: University of Minnesota Press: 73–121.

—— (1986), 'Individualism and Psychology', *Philosophical Review*, 95: 3–45.

Darwall, S., Gibbard, A. and Railton, P. (1992), 'Toward *Fin de siècle* Ethics: Some Trends', *Philosophical Review*, 101: 115–189.

Davidson, D. (1975), in S. Guttenplan (ed.) *Thought and Talk Mind and Language Mind and Language*. Oxford: Oxford University Press.

Dowell, J. (2016), 'The Metaethical Insignificance of Moral Twin Earth', in R. Shafer-Landau (ed.) *Oxford Studies in Metaethics, volume 11*. Oxford: Oxford University Press.

Dummett, M. (1975), 'What is a Theory of Meaning?', in S. Guttenplan (ed.) *Mind and Language*. Oxford: Oxford University Press.

Enoch, D. (2011), *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford University Press.

Field, H. (1977), 'Logic, Meaning and Conceptual Role', *Journal of Philosophy*, 69: 379–408.

—— (1986), 'The Deflationary Conception of Truth', in G. MacDonald and C. Wright (eds) *Fact, Science and Morality*. Oxford: Blackwell.

Fodor, J. (1987), *Psychosemantics*. Cambridge, Ma.: MIT Press.

Gibbard, A. (2003), *Thinking How to Live*. Cambridge, Ma.: Harvard University Press.

Harman, G. (1982), 'Conceptual Role Semantics', *Notre Dame Journal of Formal Logic*, 23: 242–56.

—— (1987), '(Non-Solipsistic) Conceptual Role Semantics', in E. Lepore (ed.) *New Directions in Semantics*. London: Academic Press.

Heal, J. (1996), 'Belief, Simulation and the First Person: Comments on *A Study of Concepts* by Christopher Peacocke', *Philosophy and Phenomenological Research*, 56: 413–17.

Horgan, T. and Timmons, M. (1992), 'Troubles for New Wave Moral Semantics: The Open Question Argument Revived', *Philosophical Papers*, 21: 153–175.

Horwich, P. (1994), 'What Is It Like to Be a Deflationary Theory of Meaning?', *Philosophical Issues 5: Truth and Rationality*: 133–54.

—— (2005), Reflections on Meaning. Oxford: Oxford University Press.

Humberstone, L. (1992), 'Direction of Fit', *Mind*, 101: 59–83.

Jackson, F. (1998), *From Metaphysics to Ethics*. Oxford: Clarendon Press.

Johnstone, M. (1989), 'Dispositional Theories of Value', *Proceedings of the Aristotelian Society, Supplementary Volume*, 63: 139–74.

Loar, B. (1982), 'Conceptual Role and Truth-conditions', *Notre Dame Journal of Formal Logic*, 23: 272–82.

Loewer, B. (1982), 'The Role of "Conceptual Role Semantics"', *Notre Dame Journal of Formal Logic*, 23: 305–15.

Merli, D. (2009), 'Possessing Moral Concepts', *Philosophia*, 37: 535–556.

McGinn, C. (1982), 'The Structure of Content', in A. Woodfield (ed.) *Thought and Object*. Oxford: Oxford University Press.

Millar, A. (1994), 'Possessing Concepts', *Mind*, 103: 73–82.

Moore, G. E. (1903), *Principia Ethica*. Cambridge: Cambridge University Press.

Papineau, D. (1996), 'Discussion of Christopher Peacocke's *A Study of Concepts*', *Philosophy and Phenomenological Research*, 56: 425–32.

Peacocke, C. (1987), 'Understanding Logical Constants: A Realists' Account', *Proceedings of the British Academy*, 73: 153–99.

—— (1992), A Study of Concepts. Cambridge, Ma.: MIT Press.

—— (1996), 'Can Possession Conditions Individuate Concepts?', *Philosophy and Phenomenological Research*, 56: 433–60.

Plunkett, D. and Sundell, T. (2013), 'Disagreement and the Semantics of Normative and Evaluative Terms', *Philosophers' Imprint*, 13: 1–37.

Putnam, H. (1975), 'The Meaning of "meaning"' in H. Putnam *Mind, Language and Reality*. Cambridge: Cambridge University Press.

Rey, G. (1996), 'Resisting Primitive Compulsions', *Philosophy and Phenomenological Research*, 56: 419–24.

Scanlon, T. (1998), *What We Owe to Each Other*. Cambridge, Ma.: Harvard University Press.

Schroeter, L. and Schroeter, F. (2003), 'A Slim Semantics for Moral Terms?', *Australasian Journal of Philosophy*, 81: 191–207.

Sinclair, N. (2007), 'Expressivism and the Practicality of Moral Convictions', *The Journal of Value Inquiry*, 41: 201–20.

—— (2012), 'Metaethics, Teleosemantics and the Function of Moral Judgements', *Biology and Philosophy*, 27: 639–662.

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

Snare, F. (1975), 'The Open Question Argument as Linguistic Test', *Ratio*, 17: 122–9.

Stich, S. (1992), 'What is a Theory of Mental Representation?', *Mind*, 101: 243–61.

Svavarsdottir, S. (1999), 'Moral Cognitivism and Motivation', *Philosophical Review*, 108: 161–219.

Tye, M. (1992), 'Naturalism and the Mental', *Mind*, 101: 421–41.

Wedgwood, R. (2001), 'Conceptual Role Semantics for Moral Terms', *Philosophical Review*, 110: 1–30.

—— (2007), The Nature of Normativity. Oxford: Oxford University Press.