

Alternative Normative Concepts*

Matti Eklund

[published in *Analytic Philosophy* 53 (2012)]

Consider the following scenario:

Alternative. There is a linguistic community speaking a language much like English, except for the following differences (and whatever differences are directly entailed). While their words ‘good’, ‘right’ and ‘ought’ have the same evaluative and normative roles as our words ‘good’, ‘right’ and ‘ought’ have, their words aren’t coextensive with our ‘good’, ‘right’ and ‘ought’. So even if they are exactly right about what is ‘good’ and ‘right’ and what ‘ought’ to be done, in their sense, and they seek to promote and to do what is ‘good’ and ‘right’ and what ‘ought’ to be done in their sense, they do not seek to promote what is good and right and what ought to be done.

Sometimes one distinguishes between *moral* terms and *all-things considered* normative terms, by appeal to the seeming non-triviality of worries like: I know that I morally ought to \backslash , but ought I really to \backslash ? Given such a distinction, what I am talking about here are the all-things-considered terms rather than the moral terms. Using the fashionable terminology of ‘thin’ and ‘thick’: the terms should be understood in the *thinnest* possible way.¹

Is a community like this possible? That is, are alternative normative terms like these possible? I will approach the question indirectly. In the next two sections, I will make the following points. First, if *Alternative* is in fact *impossible* then an important class of popular and attractive metaethical theories, or, more specifically, theories of the semantics of evaluative and normative terms – theories I will call *Alternative-friendly* – are false.² Among the *Alternative-friendly* theories are theories central to popular contemporary versions of naturalistic normative realism. Second, if

* I have presented this or closely related material at Brandeis, MIT, SUNY-Albany, and Boston University, and also at the Value Concepts workshop at the University of Leeds, 2010. Thanks to all these audiences for helpful feedback. Thanks also to David Liebesman, Eric Rowe and Giulia Pravato for helpful discussion.

¹ The words considered are of course context-sensitive. I will be focusing on a particular class of uses of these words.

² I will keep talking of “semantics” even where properly I should occasionally rather speak of *metasemantics*, since I will speak of theories of how semantic features are determined rather than theories of what the semantic features are.

Alternative is in fact *possible* then, as we will see, a number of novel thorny problems about normative realism and knowledge of the normative can be raised, and we face an uncomfortable choice between different things to say. The Alternative-friendly theorist must either accept that there are ineffable questions or else accept that the importance of normative questions is deflated.

These points put together serve to raise a dilemma for Alternative-friendly theories. If Alternative is impossible, such theories are false; if Alternative is possible, friends of such theories face the uncomfortable choice alluded to. The rest of the paper will then be devoted to clarifications and responses to various possible objections to this argument.

The kind of problem I will discuss is structurally similar to the Moral Twin Earth argument presented by Terence Horgan and Mark Timmons in various writings (see e.g. their 1992 and 2009), as well as to considerations prominently brought up by R. M. Hare. Later, in section 4, I will explicitly compare Horgan and Timmons' argument as well as Hare's considerations, and note the ways in which the challenge that I present differs.

1. What if Alternative is impossible?

Consider the possibility that Alternative is *impossible*. The reason it would be so is that what our (thinnest) normative words are true of is so intimately bound up with their normative roles that there is no way that words with the same normative role could fail to be coextensive. One can attempt to back this up by appeal to intuition pumps of a kind common in the literature, for example the Moral Twin Earth scenarios that Horgan and Timmons consider: when we imagine a situation of the kind described in Alternative our reaction is that we and the members of the alternative community mean the same thing by 'good', 'right' and 'ought'. However, several different kinds of theories of the semantics of normative terms demand that Alternative is possible. As I will put it, they are *Alternative-friendly*. These theories are all false if Alternative is impossible.

Take first causal normative semantics of the kind associated with so-called Cornell moral realism: normative predicates, like descriptive natural kind predicates, have their reference determined by what their use is appropriately causally linked to.³ Causal links are a different matter from normative role, and it would appear that terms can be causally linked to different properties even while they have the same normative role. Two communities could use predicates with the same normative role, even while tokenings of the two predicates stand in relevant causal relations to different things. Even if the normative role can *help explain* what the use of the predicate is appropriately causally linked to, the causal link is not a function of normative role alone, and that is sufficient for Alternative to be possible. Causal normative semantics is Alternative-friendly. In his

³ This idea is most prominent in Boyd (1988) and Brink (1989).

(1988), one of the main Cornell realists, Richard Boyd, considers an issue related to the one we are bringing up here:

The moral realist—in the guise of the homeostatic consequentialist [the specific view Boyd assumes in the paper], say—holds that what regulate the use of moral terms are facts about human well-being. But this is simply not so. Consider, for example, sixteenth-century discussions of rights. One widely acknowledged ‘right’ was the divine right of kings. Something surely regulated the use of the language of rights in the sixteenth century, but it clearly wasn’t human well-being construed in the way the moral realist intends. Instead, it was the well-being of kings and the aristocratic class of which they were part.⁴

Boyd’s reply to the objection is essentially to say that although the belief in the divine rights of kings was widespread, the use of moral terms was anyway, at bottom, causally regulated by facts about human well-being generally.⁵ This amounts to a piece of – perhaps reasonable – empirical speculation. But even if Boyd’s speculation is correct, there is still a question of whether something like what is alleged in the objection is *possible*; more specifically, whether causal normative semantics is Alternative-friendly. Boyd does not even attempt to provide a reason to think it is not.

The Cornell realists tend to discuss specifically *moral* terms. When I present the theory of how reference is determined as concerning normative terms more broadly, a step is taken. One could in principle adopt the proposed theory when it comes to the reference of the *moral* ‘right’, while – if one acknowledges a separate ‘all things considered’ use – rejecting the theory as a theory of the *all things considered* ‘right’. I will slide over such details. Similar remarks will apply to other theories I bring up in this section.⁶

Second, compare neo-descriptivist normative semantics such as that defended in Frank Jackson’s (1998). On this view, the reference of moral terms is determined by folk morality, our folk theory of matters moral. Take all moral, or generally normative, beliefs among the folk that are sufficiently widely held and firmly entrenched. These together amount to a theory. Then what the

⁴ Boyd (1988), p. 211.

⁵ Boyd (1988), p. 211f.

⁶ Does this mean that the problems I am discussing here can be easily evaded – that a friend of the views under attack can just restrict her views to the specifically *moral* and thus avoid them? No. First, obviously, if she does thus restrict her view, there are significant questions about the normative more generally that she leaves unanswered. Second, both the Cornell realists mentioned and neo-descriptivists like Jackson, who I discuss next, use their semantics to make cognitivism plausible: the accounts of the semantics is used to make plausible that moral sentences are truth-apt. If other central parts of normative language must be left out of the account, one has failed to defend cognitivism about the normative more generally. It is fully compatible with what has been defended that while there is a fact about what one morally ought to do, it is not a factual matter whether one (all things considered) ought to do what one morally ought to do.

moral terms refer to are what satisfy the result of replacing the moral terms by variables. Of course, for this to work, there will have to be sufficiently substantive beliefs linking the moral to the descriptive; and on this view the reference of moral terms is determined by such beliefs. But then Alternative should be possible: two communities can have normative terms which play the same action-guiding role but which have their reference determined by different folk theories in such a way that they are not coextensive.

An apparent complication is that Jackson speaks of the reference of moral terms as determined not by *current* folk morality but by *mature* folk morality, which is an idealized version of current folk morality. But the complication is ultimately irrelevant. If ‘mature folk morality’ was whichever folk morality got the moral facts right, then there couldn’t be two differing folk moralities of different communities determining different referents for moral terms. But such a conception of mature folk morality would trivialize Jackson’s outlook. ‘Mature folk morality’ is more reasonably understood as simply an idealized version of current folk morality: folk morality purged of, for example, incoherencies. But then different communities can have different mature folk moralities.

A third Alternative-friendly theory is that of David Copp (2001). Copp defends what he calls “realist-expressivism”, which is a basically realist view coupled with the idea that use of normative expressions serves, through a mechanism like conventional implicature, to express an attitude.⁷ Copp further attempts to deal with the question of why the truth-conditional contents of normative predicates should be linked to the conventional implicatures associated with these predicates. What he says is that to say of λ -ing that it is ‘wrong’ is to say that there is a “relevantly justified or authoritative” standard that prohibits λ -ing. This supposedly explains the implicature, for if someone says that according to some justified or authoritative standard, λ -ing is prohibited, it is only natural that she thereby should somehow convey the attitude that she is in favor of the standard. Copp thinks that the property the predicate stands for explains what attitude is expressed. It would be natural for someone friendly to all this to further hold that different attitudes are expressed by predicates standing for different properties, and that Alternative is ruled out.

However, Copp’s story seems not to work. Crucial to the story is that what is said is that according to some *justified or authoritative* standard, λ -ing is prohibited. If what was said was only that according to some standard satisfying some *purely descriptive* condition λ -ing is prohibited, the connection between content and implicature would not have been explained. But if Copp means to provide an account of normative language more generally, then what he says goes for ‘right’ and ‘wrong’ goes for ‘justified’ too, for ‘justified’ is also an normative expression. So to say of something that it is ‘justified’ is to say that it has such and such a property and moreover to express an attitude

⁷ Copp is non-committal as to whether the phenomenon he is concerned with deserves exactly the label ‘conventional implicature’. See also his (2009) on this.

about this. This problematizes Copp's account of the connection between content and implicature. For the truth-conditional content of "\-ing is wrong" is then just that there is some F standard that prohibits \-ing, where F is the truth-conditional content of 'justified'.⁸

If Copp cannot explain the conventional implicature by appeal to what property the predicate stands for, it is natural to suspect that there can, given Copp's view, be predicates standing for different properties such that the use of each serves to express the attitude expressing by, for example, the use of 'right'. But then Copp's view is Alternative-friendly. Think of it this way. On Copp's view, there are two different aspects to the meaning of a normative predicate: its contribution to truth-conditional content, and what it conventionally implicates. If one aspect does not explain the other, then these aspects of a normative predicate's meaning can vary independently of each other. For example, the contribution to truth-conditional content can vary while the conventional implicature does not. But this is Alternative-friendliness.

The Alternative-friendly theories discussed are all of the naturalist variety. But Alternative-friendliness and naturalism are not the same thing. There is nothing that immediately rules out that different communities could use their thinnest normative predicates to pick out different non-naturalistic properties. And it could be that an Alternative-unfriendly view allows that the thinnest possible normative predicates stand for naturalistic properties. Indeed, Ralph Wedgwood's Alternative-unfriendly view, which I will bring up later, is of this kind.

So to sum up the main lesson of this section: if Alternative is impossible, that rules out a number of diverse potentially attractive theories of how normative language works. Already that is a good reason for considering Alternative. But things are a whole lot more interesting if Alternative is actually *possible*. I now turn to explore that alternative.

2. What if Alternative is possible?

A first thought is that if Alternative is possible, then there is some sort of live issue as to whether we or the alternative community *get things right*. They do what they do based on considerations about what is 'good' and 'right' in *their* sense; we do what we do based on what is 'good' and 'right' in *our* sense. Since our normative terms aren't coextensive, we act differently. But whose normative terms is it that, as it were, limit the normative structure of reality? What set of normative terms ought to be used when we ask ourselves what to do? It could be ours; it could be theirs.

This thought would in turn encourage a novel kind of skeptical challenge concerning the normative: even if indeed we can come to know what is *good* and what is *right*, there would remain the

⁸ Copp also uses "authoritative". This complicates things but does not affect my criticism. Either "authoritative" too is an evaluative term and then what is said in the text applies to it too, or it is not evaluative but just expresses *that a norm is in play* without in any sense *endorsing* it, in which case it does not explain the conventional implicature.

question of whether it is these notions or the notions of this other possible community that limit the normative structure of reality – and how could we figure out the answer to *that* question? (And is there even an objective answer to that question, even if there are objective answers to questions about what ‘ought’ to be done in our sense and about what ‘ought’ to be done in their sense?)

However, this first thought is *prima facie* problematic, whatever in the end we should say about it. For exactly how should we conceive of what is at issue between us and this other community? To see the problem here, consider the following simple suggestion: the question concerns *which actions it really is right to perform*, those falling under our ‘right’ or those falling under their ‘right’. That would be a silly way to conceive of the supposed question, for – *obviously* – the *right* actions are the ones that fall under our ‘right’. That is just an instance of the general point that the things that are F are the ones that fall under the predicate ‘F’ as I actually use it.⁹

Presumably the proposed way of conceiving of what could be at issue between us and the alternative community is needlessly crude. But there is a general problem that would appear to afflict all statements of what is at issue: in any statement of ours about what is at issue, our normative terms are employed, and in such a way that the question as framed threatens to be trivially settled in favor of *our* terms, while in any statement of theirs about what is at issue, their normative terms are employed and in such a way that the question as framed threatens to be trivially settled in favor of their terms. Below I will discuss whether one might get around this problem. But for now, let us suppose – what I also think is the case – that the problem cannot be avoided. What then?

There are two things one might say. One is that there is something at issue between us and the other community, along the lines gestured at: it is only that the problems pointed to show that it is *ineffable* what the issue is. Even when it has been determined what falls under our ‘right’, ‘ought’, etc., and what falls under their ‘right’, ‘ought’, etc., there is what I will refer to as a *further question*, even if our attempts at stating that supposed question fail, and for principled reasons. While it may be that we are right about what ought to be done and they are right about what ought* to be done – where I use ‘*’ to indicate the alternative ‘ought’ of the other community – either they are wrong in letting their actions be guided by considerations about what ought* to be done or we are wrong in letting our actions be guided by considerations about what ought to be done.¹⁰ I want to discuss the issue in the abstract and not tie the discussion to any particular type of hypothesis about the extensions of the alternative normative concepts, but for illustrative purposes, suppose a deontological theory is true of what ought to be done, but the ‘ought*’ of the other community is consequentialist. Suppose

⁹ Setting aside obviously orthogonal complications related to the liar paradox.

¹⁰ Of course not even the most morally conscientious among us are always concerned with what ought to be done; and even when we are so concerned, we can be wrong, and even radically so, about what ought to be done. And if the members of the alternative community are like us, the same will hold of them, *mutatis mutandis*. In the main text I slide over this, as it doesn’t affect the main points.

further that they are about to perform some action – say, sacrifice an innocent child for the greater good – that is consequentialistically but not deontologically sanctioned. “Wait”, we say, “you ought not to do that!”. “We know”, they reply, “that we ‘ought’ not to do this. But we ought* to do it.” We persist: “But what you ought* to do is not what you really ought to do”. They say: “Right again – but it is also true that what one ought to do is not what one really ought* to do”. Everything they say in reply to us is true; but as the child dies we may still feel that there is something there that they missed. What we feel is that there is a further question there, one about which they are tragically mistaken. (Of course, if we feel that there is a further question we might also feel the nagging suspicion that we are tragically mistaken about it. As we do what we ought to do and don’t maximize happiness, we may worry: maybe it is instead we who are tragically mistaken?... Even if, of course, we *ought* to do what we *ought* to do.)

Another thing one might say, eschewing the idea of a further question, is that once it is settled what falls under ‘right’, etc., and what falls under their ‘right’, etc., all the questions in the vicinity have been settled. There’s what’s “right” in our sense and what’s “right” in their sense, and that is that. The idea that there are any further questions there, whether expressible or not, would be mistaken. I will not attempt here to decide which reaction would be the most reasonable – although needless to say one may be uneasy about the idea of ineffable questions. Instead I will discuss the upshot of each suggestion.

It is fairly obvious what the upshot might be of saying that there is a further question. The skepticism mentioned above is then a live issue. Even if we are perfectly successful in figuring out what is right, good, etc., there remains the fact that there are other possible words with the same normative roles as ours but with different extensions, and maybe we should care instead about what falls under those words. (Of course the ineffability issue rears its head as I try to state the upshot. I use the ‘should’ of my language when discussing this. But don’t begrudge me a pinch of salt when I discuss the idea of embracing the idea of a further question – even if, by the end of the day, this route should be rejected.¹¹) And no matter how successful our methodology for figuring out what is good, right, etc. may be, what reason is there to think that we are also right in thinking that this is what we *should* care about? (Again, don’t begrudge me a pinch of salt.)

I think most would find it more reasonable simply to deny that there is an ineffable further question. But also such a denial has important, and potentially problematic, consequences. Consider again a situation like that of the sacrifice of the innocent child, above. Suppose it is right to \ but right* to / and \-ing and /-ing are incompatible. Suppose further that from our perspective, /-ing in this situation would be abhorrent. Surely what we would *want* to say that they – the others – are

¹¹ The phrase ‘pinch of salt’ was famously used by Frege, when he saw himself running up against the in principle ineffable. See Frege (1892).

just mistaken. *What one really ought to do* in this situation is to \setminus . And of course we *can* say this latter thing; it is not as if, on the option we are now exploring, this is an untrue thing to say. The problem is that *they* can say, with *equal* justification, *what one really ought* to do* in this situation is to \setminus . We may care more about what one ought to do, but they care more about what one ought* to do. We can say that one *ought* to be concerned with what one ought to do rather than with what one ought* to do, and maybe we can say so truly; but *they* can similarly say that one ought* to be concerned with what one ought* to do rather than with one ought to do. And so on. If there is no further question, this is *all there is to it*.

Let the *complacent attitude* be that of despite what has just been stressed resting content, when deliberating about how to act, with being concerned with what one ought to do, and of not being bothered by the fact that there are other things one ought* to do. The friend of the complacent attitude can perhaps attempt to defend her stance by saying that discontent with this attitude stems from the idea that there is a further question, but once we have abandoned that idea, we can and should be content with the complacent attitude. Only if there is a further question can there be an ulterior standard against which what I have dubbed the complacent attitude can be measured and criticized. But whereas such a response might have been in order if it could be supposed that there were only two options – to accept that there is a further question or to adopt a complacent attitude – these are not, in the wider scheme of things, the only two options. One other option is to reject Alternative-friendly views. And yet another option is to retain an Alternative-friendly view, reject the idea of a further question, and instead of adopting what I have called the complacent attitude take the questions about what is good and right and what ought to be done to be *deflated*. If we are inclined to feel good about doing what ought to be done rather than what ought* to be done, and to criticize them for doing what ought* to be done rather than what ought to be done, we might consider that the corresponding reactions on *their* part can be equally efficiently supported.

To stress, the difference between the complacent attitude and the deflationary attitude is this. The complacent attitude is that of according the same importance to questions about what is good and right and what we ought to do as we always have, even given a realization that Alternative is possible and that there is no further question. The deflationary attitude is that the importance of these questions is, precisely, deflated if Alternative is possible and there is no further question. I submit that deflationism seems clearly more reasonable than complacency. Once we have abandoned the idea of a further question, what *basis* can there possibly be for attaching greater objective significance to what we ought to do than to what we ought* to do? We can say that it is right and good to do one what ought to do; but there is the little devil ready to whisper in one's ear what surely is true: that *they* can equally comfort themselves with the thought that it is after all right* and good* to do what one ought* to do. And – if there is no further question – that is all there is to it.

The structure of the foregoing has been somewhat complex. Recapitulation is in order. There is a class of theories of the semantics of normative terms on which Alternative is possible; these theories are Alternative-friendly. If Alternative is in fact impossible, then of course these theories are simply false. If Alternative is possible, then, one may think, there is a *further question* about which normative terms are privileged. I indicated that any such further question would appear to have to be ineffable. If there still is a further question like this, difficult questions arise concerning how it is to be resolved. If there is no further question, then the issue arises of whether the complacent attitude is appropriate, or whether one should instead take the normative issues to be deflated. When it comes to this last issue, I have indicated that the latter option is preferable.

Note how the considerations presented function as an extended argument against the theories I have dubbed Alternative-friendly. Suppose you are an Alternative-friendly theorist. If Alternative is impossible, you are wrong. So suppose Alternative is possible. Is there then a further question of the kind indicated? You might say there is, but then you are committed to the idea of ineffable questions. Such ineffability is something many would feel uncomfortable with. In addition, even if you can say that sentences about what is good and right and ought to be done are capable of objective truth and falsity, that still does not warrant any claim to the effect that the ineffable question has an objective answer. Suppose then that you instead reject the idea of a further question. Then you are stuck with the deflationary attitude; you've ended up deflating the significance of questions of what is good and right and what ought to be done.

My main points have been stated. In the sections to follow, I will fill in some detail and respond to some envisaged objections.

3. Expressing the further question

As earlier indicated, one may think there are ways of expressing what I have called the 'further' question, so that this supposed further question would not have to be ineffable.

One possible model is provided by a suggestion that Ted Sider (2009) discusses favorably: that debates over aesthetic realism should be recast so they are no longer conceived of as being over whether aesthetic sentences have objective truth-values, but over whether some aesthetic predicates are the truly 'natural' or 'fundamental' or 'joint-carving' ones.¹² Sider's idea is that the would-be aesthetic antirealist fights a losing battle if she insists that aesthetic sentences lack truth-values. Instead she should insist that there are no most natural possible aesthetic predicates – different counterparts of, say, 'beautiful' used in different communities. Correspondingly, the aesthetic realist is to be seen as affirming that there are some unique most natural aesthetic predicates.

¹² Sider (2009), p. 401.

Notice the similarity between Sider's proposal concerning aesthetic realism and the themes here brought up. In the dispute over aesthetic realism as Sider conceives of it as being conducted, both the realist and her opponent agrees that the aesthetic sentences of the different languages have objective truth-values: they disagree over whether there is an objective answer to the further question of which aesthetic notions carve the world at its joints. Applied to the general issue at hand here: it can be agreed on all hands that claims about what ought to be done, what ought* to be done, etc., have objective truth-values, but – the idea would be – there is a further question of whether 'ought' or one of its counterparts carves at the world's joints, and, if so, which one does so.

Sider's idea might be attractive to a friend of the idea of a further question who wants nothing to do with absolute ineffability. While attempts at asking the further question which crucially make use of normative terms fail, for reasons indicated, we can non-trivially raise the question of which normative terms are the most *natural* or *fundamental* ones – stand for the most natural or fundamental properties – and thus we can ask the further question after all; or that would be the idea. The further question between us and the alternative community concerns whose terms stand for the most natural or fundamental properties. When we have the sense that they go wrong when using the *-terms for evaluation, and that they had better use our terms, what that amounts to is our holding that our terms carve at the joints.

But Sider's proposal for how to understand the debate over aesthetic realism fails, for the reason that some aesthetic predicates may be the truly natural or joint-carving ones *while this says nothing about their aptness for aesthetic evaluation*. The following example should help make the point. Consider two different communities with different aesthetic predicates – for example, they may have different, non-coextensive predicates 'tasty'. Suppose further that the tastes of one community are such that the extension of 'tasty' in their mouths is more metaphysically natural, for example because there is one particular chemical such that they like – *gustatorily* like – food and drink that contains this chemical in sufficient quantities, so the referent of their 'tasty' is more natural. To take this to be relevant to which community objectively has the aesthetically better taste would clearly be unwarranted. Correspondingly for other normative disputes, including for example moral disputes, even if in other cases it is harder to come up with even *prima facie* compelling examples of greater naturalness. The general lesson is that what is *metaphysically better* – more 'natural' – need not be *evaluatively* or *normatively* better in the relevant dimension.¹³ The point generalizes. Any attempt to get

¹³ There are different conceptions of naturalness; and Sider's own view on naturalness incorporates elements that make the argument just given more difficult to run. But these complications do not at bottom affect things.

First, it is reasonable to insist (as Sider does) that naturalness is not best conceived of as *physical* naturalness. If so, then the greater physical naturalness of one of the 'tasty'-predicates doesn't immediately entail that this predicate is more natural. But this is incidental to the point of the example. Let metaphysical

around the threatening ineffability of the further question by saying that it can be stated in non-normative terms founders on the fact that what is stated then just is not the supposed further question.

4. Moral Twin Earth

Let me now turn to the Moral Twin Earth argument of Horgan and Timmons, as earlier promised. This argument is a prominent argument brought against the causal theory of reference for moral terms, and more recently against neo-descriptivism.¹⁴ As it is typically discussed, Horgan and Timmons' argument is, first, that these theories allow that there should be terms of other possible languages corresponding to our terms in normative role but differing in extension, and, second, by appeal to a thought experiment – the Moral Twin Earth scenario – that this consequence is counterintuitive, for intuitively our terms and these possible terms would have the same meaning and reference. One might think that what I have done in this article is just to bring up this worry in a different form. Let me then compare my argument to that just summarized.

The first claim in the argument from Horgan and Timmons is exactly what I am noting here. But the present challenge is independent of the truth of the second claim. Dialectically, this is important. As stressed in David Merli (2002), the Moral Twin Earth thought experiment is radically underdescribed. It is not explained just *how* the twin-earthian terms come to refer to the different things they refer to, and in the absence of such an explanation one may reasonably suspect that were the thought experiment to be more adequately spelled out, our intuitions would not clearly go against the causal and neo-descriptivist theories. The considerations presented in section 2, concerning what to say if Alternative is possible, illustrate the problems that Alternative-friendly theories face even if the second claim of the Moral Twin Earth argument is false.

Another difference, of course, is that I set my sights on Alternative-friendly theories generally. There is a connection between this difference and the one just noted. Since Horgan and Timmons rely on its being intuitive that we and the twin-earthians mean the same by the relevant

_____ naturalness be anything you like: metaphysical naturalness and aptness for aesthetic evaluation can still come apart.

Second, one can take naturalness to be an all-or-nothing matter rather than a matter of degree. (Again, Sider does this.) But this is relevant only insofar as it can be hard to come up with *any* counterparts of 'tasty' that promise to be natural on an all-or-nothing conception of naturalness, and such that the argument from the main text can be run with those as examples instead. And if the defender of Sider wants to rely crucially on that, she faces a different problem: doesn't the proposed reconceptualization of the debate over aesthetic realism then hand the aesthetic antirealist a too easy victory? The reason would be that no referents of aesthetic predicates are plausibly natural on an all-or-nothing conception of naturalness, for none is plausibly completely natural.

¹⁴ See e.g. Horgan and Timmons (1992) for their argument against causal theories, and their (2009) for the argument directed against neo-descriptivism.

terms, they must consider the theories they target one by one to verify that this really is intuitive. By contrast, the dilemma here focused on can be raised for Alternative-friendly theories wholesale.

As they make clear, Horgan and Timmons take their inspiration from Hare when giving the Moral Twin Earth argument. Hare presents similar thought experiments. What Hare (e.g. 1952, 1997) prominently argued was that naturalism in ethics leads to “relativism”. One reason it is relevant to bring up this claim of Hare’s is that Hare can be seen as attempting, as I am, to investigate the consequences if Horgan and Timmons’ second claim is false.

I think Hare’s charge of relativism is *clearly* false on the most straightforward understandings of what relativism is. First, relativism can be held to involve the claim that *proposition truth* is somehow relative (to an assessor, or a point of view, or...). Naturalism does not entail that. Nor is there anything in Alternative-friendliness that suggests this. (I would myself prefer it if the label ‘relativism’ was reserved for a view of this first kind.) Second, another view sometimes called relativist is a view on which the semantics of the normative is contextualist – so that, to take a simple case, when I say “this is right”, what proposition I express depends on the context, and is something like *this is right relative to standard S* – is sometimes called relativist. But it is hard to see why naturalism would be especially likely to suggest this sort of view, whatever in the end its fate. And again, Alternative-friendliness does not entail a contextualist semantics.

Anyway, what Hare was really after is that naturalism yields that different communities can have words with the same normative role but with different extensions. (In this connection he also makes the further point that this consequence is implausible.) In his (1997), he says,

There are in most languages words which we translate ‘wrong’. These words are....rough equivalents to one another. But the cultures that use these words call quite different things wrong. In one culture, for example, it may be thought wrong not to fight for one’s country, in another more pacific culture it may be thought wrong to fight. The important thing to get hold of is that.... the people in these different cultures....may be using the word ‘wrong’, or its equivalents, in the same sense. Otherwise they would not be contradicting one another, which they clearly are. The people in one culture are saying that fighting is wrong and the people in the other are saying that it is not wrong, *in the same sense of ‘wrong’*, so far as its evaluative meaning goes. But if we follow the naturalists, we shall have to say that the senses of the word in the two cultures are entirely different. This will have the consequence that they are not contradicting one another....If we distinguish the senses by using different subscripts, we can say that one of the

cultures thinks fighting is wrong₁, but that the other thinks it is not wrong₂. But these two opinions may be mutually consistent, if the two senses of ‘wrong’ are different.¹⁵

Hare describes the consequence as ‘relativism’.¹⁶

One may well quarrel with Hare’s reasoning here. The fact that the two cultures are calling different things ‘wrong’ does not immediately entail that they use ‘wrong’ with different senses, on any reasonable descriptivist view. (If this is not immediately clear, note that for any *descriptive* predicate ‘F’ two cultures can call different things ‘F’ and yet use ‘F’ with the same sense: all that would be going on is that they have different *beliefs* about what is F.) Charitably, one can interpret Hare as, rhetorically, making the point that any naturalist theory will entail that some two cultures which intuitively use ‘wrong’ with the same sense are using it with different senses. But there is no proper argument for that claim here.

Even if Hare should be right about what naturalism entails, it is misleading to call the entailed claim ‘relativism’. The consequence that I can truly say “\-ing is wrong” while an interlocutor can truly say “\-ing is not wrong” can sound relativist. But if all that is going on is that we use ‘wrong’ in different senses, this is no more relativist than the point that a Brit and an American can point to the same thing and say “that is chips” and “that is not chips”, respectively, and yet each speak the truth. And I think Hare is clearly wrong to think that all versions of naturalism need have the consequence in question. A naturalist may in principle accept something like Wedgwood’s theory of the metaseantics of normative expressions, which I will present below.¹⁷

Part of what is going on is that Hare is making the point made in the Moral Twin Earth argument. He clearly thinks that the conclusion that the different communities mean different things by ‘wrong’ is false. But when he calls the consequence *relativistic* another theme is arguably brought in. What he says, when elaborating upon the passage quoted, is:

There would be no harm in this if all they were doing were *describing* the act of fighting... The trouble starts when we begin using ‘wrong’ for the purpose for which it actually is used in language, namely for *condemning* acts. [I]t is very natural, since this *is* actually [the use of ‘wrong’], to think that the people in the two cultures are, respectively, condemning and refusing to

¹⁵ Hare (1997), pp. 68ff. In early writings the target of Hare’s criticism was naturalist theories according to which normative terms could be exhaustively analyzed in descriptive terms. But in (1997), he explicitly also includes Cornell realism of the kind earlier described. However, he does not indicate exactly how the argument is supposed to work against Cornell realism.

¹⁶ Hare (1997), p. 65.

¹⁷ Wedgwood himself defends a non-reductive naturalism. See his (2007).

condemn the act of fighting. Then they *are* contradicting each other. But according to the naturalist they may both be right in what they say. There is no contradiction. The naturalist seems to be led to the conclusion that it is both right for one culture to condemn fighting, and right for the other culture not to condemn it. And this is a relativist position.¹⁸

As Hare actually states his worry, he misdescribes his opponent's position. The opponent's view isn't that it is right for one culture to condemn fighting and right for another not to condemn it. It is rather that it is, say, right to condemn fighting and *right** not to condemn it. It can be tempting to just dismiss Hare as being unfair to his opponent. But there is a real concern underlying what Hare says. This concern appears similar to the idea of a deflationary attitude brought up in connection with the suggestion that there just is no further question: that the idea that some action can fall under one 'wrong' but not a normative counterpart *and that this is all there is to it* is incompatible with a non-deflationary realism.¹⁹

Toward the end of his (1988), Boyd brings up something like the relativist worry. He worries that there may be differences in moral beliefs resulting in differences in the extensions of moral expressions. But after having brought up this objection under the heading of 'relativism', Boyd says that the scenario envisaged

...is nevertheless only in a relatively uninteresting sense non-realistic. The dependence of the truth of moral propositions upon moral beliefs envisioned [in the sketch of the relevant type of scenario] would be...an ordinary case of causal dependence and not the sort of logical dependence required by a constructivist conception of morals analogous to a Kuhnian neo-

¹⁸ Hare (1997), p. 69.

¹⁹ Brink (2001, p. 166) calls the kind of relativism that would be at issue "semantic relativism" and brings up the question of why it would be incompatible with realism. His answer is that in the case we are dealing with, the reason our words and those of the Twin Earthlings have different referents has to do with our having different *beliefs* and is not explained by differences in the physical environment alone. This is unconvincing. Why should this matter? For example, on some views, like general neo-descriptivist views on language, beliefs generally have a reference-fixing role. Does that mean that such theories entail a rampant realism-incompatible relativism? If I hold that different communities of scientists refer to different things by 'electron' because the members of these communities have different beliefs associated with the word, am I then committing myself to a realism-incompatible relativism about electrons? Hardly – all I hold is that they refer to different things by 'electron'.

Horgan and Timmons also use the label 'relativism', but one must be careful about what they mean by it. They distinguish two kinds of 'relativism'. First, they describe the upshot of the Moral Twin Earth argument as chauvinistic conceptual relativism. I find the use of the relativist label unfortunate, but be that as it may. The important thing to note is that the relativist charge thus understood is not a separate charge. Second, they note that one way for the target of their argument to respond is to say that even if we and the Twin Earthians refer to different things by our moral words still the words *mean* the same. This would be a genuine relativist view, what they call 'standard relativism'. But what they are considering under this heading is clearly an optional move for their opponent, and not anything that immediately falls out of Alternative-friendly views.

Kantian conception of the dependence of scientific truth on the adoption of theories or paradigms. The subject matter of moral inquiry in each of the relevant communities would be theory-and-belief-independent in the sense relevant to the dispute between realists and social constructionists.²⁰

Boyd is certainly right to distinguish the possibility brought up from traditional antirealism or relativism. But even if traditional antirealism or relativism isn't what is at issue, there remains the issue I have described. If the thin terms employed by different communities can have different referents, and there is no 'neutral' way to settle which vocabulary ought to be used, even in principle, then realism does not achieve all one might have hoped it should achieve.

If there is a real threat to Boyd's type of theory in the vicinity, it is not relativism but the problematic presented here. Specifically, any proponent of an Alternative-friendly theory must either embrace the idea of a further question, one which would appear to have to be ineffable, or face up to the consequences of denying that there is such a further question.

My emphasizing that we are not dealing with anything properly called relativism may seem like a mere terminological point, and not anything worth belaboring. But here is the reason why stressing it is important. Calling the consequence we are concerned with – the possibility of Alternative – “relativism” is to attach a label to it which immediately and misleadingly suggests a certain compromise with what we naturally regard as realism. I think that the possibility of Alternative does present real problems. But it is important to realize that if indeed it does, it does so via the somewhat laborious route here presented: via the discussion of the threat of ineffable questions and the threat of having to accept a deflationary attitude.

5. Sameness of normative role

Throughout I have talked of predicates having the same *normative role*. This may well be a source of considerable uneasiness. What is it for two predicates – two non-coextensive predicates, at that – to have the same normative role? I find the idea of normative role rather intuitive myself. It is characteristic of normative predicates that they are fit to be used in practical deliberation about what to do; it is characteristic of such predicates that their application has, so to speak, practical consequences in addition to merely theoretical ones. They have normative roles. And I think that it is likewise clear that different normative predicates are used differently, as far as such roles are concerned. For example, the normative use of 'thin' predicates is different from that of the various 'thick' ones, and evaluating morally is different from evaluating aesthetically. Different normative

²⁰ Boyd (1988), p. 225f.

predicates have different normative roles. But if one can compare normative roles one can also judge when two predicates have the same normative role. Even if one cannot state in very informative terms what the normative role of a predicate is, one can intelligibly speak of sameness of normative role.²¹

Suppose, however, that I should not be granted this. Suppose either that the idea of sameness of normative role of non-coextensive predicates is not properly intelligible or that the idea of the normative role of a predicate is misguided in the first place. Even so the problems I have been concerned with here do arise, as follows.

Consider a number of different linguistic communities, each with its own set of normative terms, including its own set of 'thin' terms, such that the normative terms of the different communities all fail to be coextensive with each other. Suppose further – along the lines of the doubts just brought up – that it is at best unclear whether any of these terms used by the different communities have the same normative role. It cannot be said that any alternative 'ought' determinately has the same role as ours does. And suppose lastly that *which* of these normative terms you employ matters for action-guiding purposes: the positive normative terms of one language apply to different things from what the positive thin terms of another language apply to. Again the concern arises: how should you react, knowing that while you ought to \setminus , you ought* to \setminus , where *ought** is what is expressed by one of these alternative normative terms? The same puzzles as before arise with respect to *this* question. But then the appeal to sameness of normative role is not essential for the problem I have been concerned with to arise.

6. The extent of the divergence

I have said that some theories of normative semantics are Alternative-friendly: on these theories even the thinnest terms can have the same normative role yet differ in extension. But I have left open the *extent* to which they can differ in extension. It may be thought that if the extensions of predicates with the same normative role can only differ slightly, then the problems I have discussed are not very serious. The normative predicates corresponding to each other in normative role by and large agree

²¹ What I am talking about under the heading of normative role is what Horgan and Timmons (1992) are talking about in different terms when describing their Moral Twin Earth scenario. When describing the case, they say, "Moral Twin Earthlings are normally disposed to act in certain ways corresponding to judgments about what is 'good' and 'right'; they normally take considerations about what is 'good' and 'right' to be especially important, even of overriding importance in most cases, in deciding what to do, and so on" (p. 188). Horgan and Timmons' characterization is so vague that it does not solve any problems anyone might have with the notion of sameness of normative role. But it does illustrate that the idea of such sameness is found in the literature.

Sameness of normative role is also what Hare, from his non-descriptivist perspective, in the passage quoted above discusses under the heading of evaluative meaning, and what he alludes to when calling the different words which we translate 'wrong' "rough equivalents".

with each other in their application, and if some proposed action is in the extension of all the counterparts of ‘right’, then we can say that this action is licensed as right in a thoroughly realist and non-deflationary way. Maybe there are actions such that they are in the extensions of some but not all counterparts of ‘right’, and in *some* sense there may be no ‘objective’ fact of the matter as to whether they really ought to be performed, but so long as these actions are fairly few and far between that need not be of great concern. One can then avoid both the idea of a ‘further’, ineffable question and the threatening deflationary consequences of denying that there is such a question. The whole problematic is, one may think, defused.

Nothing I have said so far blocks this strategy. But if the proponent of a supposedly Alternative-friendly theory attempts this as a way out, she is for all intents and purposes abandoning her Alternative-friendly theory. She is emphasizing that ‘right’ and its counterparts divide actions into three classes – (i) the class of those actions falling under all of them, (ii) the class of those actions falling under some but not all of them, and (iii) the class of those actions falling under none of them – and what is more, she accords a certain importance to this tripartite distinction. There is something important that distinguishes those actions falling under (i); and the question of which actions falling under (i) is privileged over the question of what actions fall under our ‘right’. But if she says this, nothing of importance separates her from a theorist who rather says that the extension of ‘right’ is completely determined by the predicate’s normative role but allows that the normative role does not determine a fully determinate extension. Of course, she says something different about the specific semantics of ‘right’ than does this other theorist. She allows that some actions in category (ii) fall under ‘right’. But by her own lights, the normatively most important question is not about what happens to fall under our ‘right’ but instead about what falls under ‘right’ and each of its normative counterparts; what is guaranteed to fall under all these predicates because of the normative roles of these predicates.

This point actually serves to highlight a subtlety I slid over when introducing the problem. When introducing Alternative, I stipulated that we were to understand the normative terms at issue in the thinnest possible way, and I also stipulated that the normative terms of the alternative community are counterparts of our all-things-considered ‘ought’, ‘good’ and ‘right’. This way of presenting things presupposes that these normative terms of ours are as thin as normative terms can possibly be.²² That is a reasonable presupposition. But it could be false: it could be that *all* our terms are to some extent thick, but perfectly thin terms can in principle be introduced. If indeed there can be a discrepancy like this, the issue to focus on is whether there can be non-coextensive counterparts of the thinnest possible terms.

²² A reminder (compare fn. 1): the relevant terms are context-sensitive. Strictly speaking, I am focusing on a class of *uses* of these terms.

One reason why it is hard to adjudicate the question of whether Alternative is indeed possible is that it is not simply a matter of deciding whether there can be non-coextensive counterparts of our normative terms. Even if that should be so, matters could stand differently with the thinnest possible terms.

7. Alternative-unfriendly semantics

I have talked about potential problems for Alternative-friendly views. But what, one may wonder, might an Alternative-*unfriendly* view be like? One may wonder whether not all views on normative semantics are in fact Alternative-friendly, in which case the problem brought up is rather more general than here brought out. In response, let me briefly describe one prominent view that promises to be Alternative-unfriendly: that of Wedgwood (2001, 2007).²³ Wedgwood focuses on a maximally ‘thin’ predicate meaning roughly ‘ought all things considered to prefer’, whose semantic value is conceived of as completely determined by its *action-guiding role*; its role in *practical reasoning*. More precisely, what Wedgwood proposes is that the predicate is governed by the following basic rule:

Acceptance of ‘B(x, y, me, t)’ commits one to having a preference for doing x over doing y at time t.²⁴

The semantic value of this predicate is then, Wedgwood argues, the four-place relation between x, y, z, and t (if any) such that “it is correct for z to prefer doing x over doing at t and a mistake for z to prefer doing y over doing x at t if, and only if, x, y, z, and t stand in that relation”.²⁵ The preference is correct if and only if it is in accordance with the goal of practical reasoning; and it is a substantive question if there is such a goal and, if so, what the goal is. The predicate Wedgwood focuses on accordingly has its reference determined in an importantly different way from that in which ordinary descriptive predicates have their reference determined: Wedgwood focuses on the practical commitments incurred by acceptance of sentences where the predicate is the main predicate.

On Wedgwood’s theory, the reference of the thin predicate ‘B(x, y, me, t)’ is determined by normative role. Any predicate with this normative role has the same reference. Wedgwood’s theory is not Alternative-friendly.

The distinction drawn in the previous section, between the question of what our thinnest terms are like and what the thinnest possible terms are like, is relevant to the assessment of Wedgwood’s theory. One possible concern with a view like Wedgwood’s is that attention to the use

²³ I will follow Wedgwood’s (2001) presentation, which is somewhat simpler.

²⁴ Wedgwood (2001), p. 15.

²⁵ Wedgwood (2001), p. 18.

of normative terms and concepts, and our judgments concerning possession of normative concepts, shows that none of our concepts is as thin as the concept that Wedgwood describes is.²⁶ Whatever in the end should be said about this, it is a natural concern to raise. But when thinking about its relevance to the present issue, we must distinguish between two claims that can be made on behalf of an Alternative-unfriendly view like Wedgwood's. One is that it correctly describes some of our normative expressions; a second is that it correctly describes some *possible* evaluative expressions that could be used. For many purposes, the first claim will naturally be in the spotlight. But for our purposes, the second, weaker, possibility claim is the relevant one. So long as this possibility claim is true, the problems for Alternative-friendly views are evaded: when asking which 'right' it is that *ought* to be used, ours or that of the alternative community, we can then use a perfectly thin, Wedgwoodian 'ought'.

Wedgwood's theory certainly faces its own problems. I don't want to attempt to argue in favor of it or any similar theory. But it is the sort of theory that promises to get around all the problems here considered.

8. Conclusion

A quick summary is in order, especially since I have been preoccupied with responses to various sorts of objections ever since section 3. Either Alternative is possible or not. If it is not possible, a number of theories of the semantics of normative terms, the Alternative-friendly views, are false. If it is possible, difficult questions arise concerning what to say about the issue of whether any actual or possible terms are in any way objectively normatively privileged. Discussions in the literature of whether various naturalist theories lead to antirealism or relativism are better cast as discussions of the challenge here posed. These naturalist theories are apparently Alternative-friendly, and given that they are, friends of these theories will have to either accept that there is a 'further' question (which appears to have to be ineffable), or embrace a deflationary attitude.

REFERENCES

- Boyd, Richard: 1988, "How to be a Moral Realist", in Geoffrey Sayre-McCord (ed.), *Essays on Moral Realism*, Cornell University Press, Ithaca, New York, pp. 181-228.
- Brink, David: 1989, *Moral Realism and the Foundations of Ethics*, Cambridge University Press, Cambridge.
- Brink, David: 2001, "Realism, Naturalism, and Moral Semantics", in Ellen Frankel Paul, Fred Miller and Jeffrey Paul (eds.), *Moral Knowledge*, Cambridge University Press, Cambridge.

²⁶ See here Merli (2009).

- Copp, David: 2001, "Realist-Expressivism: A Neglected Option for Moral Realism", *Social Philosophy and Policy* 18: 1-43.
- Copp, David: 2009, "Realist-Expressivism and Conventional Implicature", *Oxford Studies in Metaethics* 4: 167-202.
- Frege, Gottlob: 1892, "On Concept and Object", *Vierteljahresschrift für wissenschaftliche Philosophie* 16: 192-205. Reprinted in Michael Beaney (ed.), *The Frege Reader*, Blackwell, Oxford, pp. 181-93.
- Hare, R. M.: 1952, *The Language of Morals*, Clarendon Press, Oxford.
- Hare, R. M.: 1997, *Sorting Out Ethics*, Clarendon Press, Oxford.
- Horgan, Terence and Mark Timmons: 1992, "Troubles for New Wave Moral Semantics: The 'Open Question Argument' Revived", *Philosophical Papers* 21: 153-75.
- Horgan, Terence and Mark Timmons: 2000, "Copping Out on Moral Twin Earth", *Synthese* 124: 139-52.
- Horgan, Terence and Mark Timmons: 2009, "Analytical Moral Functionalism Meets Moral Twin Earth", in Ian Ravenscroft (ed.), *Minds, Ethics and Conditionals*, Oxford University Press, Oxford, pp. 221-37.
- Jackson, Frank: 1998, *From Metaphysics to Ethics*, Clarendon Press, Oxford.
- Merli, David: 2002, "Return to Moral Twin Earth", *Canadian Journal of Philosophy* 32: 207-40.
- Merli, David: 2009, "Possessing Moral Concepts", *Philosophia* 37: 535-56.
- Sider, Theodore: 2009, "Ontological Realism", in David Chalmers, David Manley and Ryan Wasserman (eds.), *Metametaphysics*, Oxford University Press, Oxford, pp. 384-423.
- Wedgwood, Ralph: 2001, "Conceptual Role Semantics for Moral Terms", *Philosophical Review* 110: 1-30.
- Wedgwood, Ralph: 2007, *The Nature of Normativity*, Clarendon Press, Oxford.