

Jennifer Foster and Mark Schroeder
University of Southern California
January 5, 2022

Defining Moral Realism

*And as imagination bodies forth,
the forms of things unknown, the poet's pen
turns them to shapes and gives to airy nothing
a local habitation and a name.*

William Shakespeare, *A Midsummer's Night Dream*

Wherever philosophers disagree, one of the things at issue is likely to be what they disagree about, itself. So also with moral realism, or meta-normative realism more broadly. In addition to asking whether moral realism is true, and which forms of moral realism are more likely to be true than others, we can also ask what it would mean for some form of moral realism to be true, seeking after definitions of 'moral realism' and each of its standard variants, 'naturalism', 'non-naturalism', and so on. The usual aspiration of such inquiry is to find definitions that all can agree on, so that we can use terms in a uniform way. But we doubt that this aspiration is always possible, or even necessarily desirable. It will be our goal in this essay to sketch out some of our reasons for such skepticism, and to lay out a picture of what philosophical inquiry can look like in metaethics and beyond, even when it is impossible to reach uniform agreement on the terms of the debate.

I Fixing Vocabulary

We all know how important it is to fix the meanings of terms in order to have a clear topic for conversation. If we are having a conversation and Jen insists, "moral realism is not true", while Mark insists "moral realism is true", we are having no disagreement at all, if both Jen and Mark accept that the word "wrong" is synonymous with "fails to maximize happiness" but Mark thinks that this makes realism true, while Jen, following Nagel, Parfit, and others, prefers to use "realism" in a stingier way. Perhaps she thinks, along with Nagel, that

If values are objective, they must be so in their own right, and not through reducibility to some other kind of objective fact. They have to be objective *values*, not objective anything else.¹

¹ Nagel [1986, 138]. Contrast Schroeder [2005].

So in order to fix what we are in disagreement about, we – Jen and Mark – need to fix our terms. In particular, we need to decide whether to use “moral realism” in a way that includes Benthamite analytic utilitarianism or not. So far, so trivial – all of this, we believe, should be uncontroversial.

But when philosophers look for the right definitions to use for classificatory terms like ‘moral realism’, ‘reduction’, ‘naturalism’, ‘noncognitivism’, ‘expressivism’, ‘constructivism’, ‘constitutivism’, and many more, they are generalizing on this uncontroversial observation, to look not just for the definitions that will help to fix terms for the purposes of one conversation or another, but *once and for all* – for all conversations that different people have about this topic, and for all conversations that we may come in the future to have about it. The metaphor of academic publishing as a kind of ongoing conversation – a metaphor of which we are both fond, for many purposes – encourages the sense that this must be an important thing to do. For if it is important for Jen and Mark to fix terms for purposes of their conversation, then if their conversation is just a small part of a much bigger conversation, then surely that much bigger conversation can likewise be productive only if we can fix terms for its purposes, as well.

But we doubt that this is such a good idea. Indeed, we doubt that it is even possible. The core reason for our doubts is simple. It is that there is a key disanalogy between ordinary conversations, which have a limited number of participants and happen at a particular time and place, and the metaphorical “grand conversation” about philosophy to which we all aspire to contribute. The disanalogy is that precisely because ordinary conversations happen at a particular time and place, with a limited number of participants, the conversational participants can draw on features that they have in common in order to fix their terms adequately for purposes of the conversation. The larger the conversation, however, and the less fixed in time and place, the less that the conversational participants will have in common – and the harder it will be to adequately fix terms in a single meaningful way. We should expect, then, that in the “grand conversation” of philosophical inquiry as a whole, this task will be hardest of all: for such a (metaphorical) conversation is limited neither in time nor space, nor in who may join in, nor in which background assumptions those participants bring with them. For the purposes of *that* conversation, fixing terms once and for all in a single, meaningful way will be challenging, indeed.

So how is it, in more ordinary cases, that the common features of the conversational participants make it easier to fix terms? The answer is that it gives them a *common ground* – both a set of assumptions that can be held fixed for purposes of their conversation and a range of imaginative possibilities beyond which they do not need to plan in advance.² Assumptions can be held fixed for purposes of a conversation either

² Stalnaker [2002].

because they are taken for granted by all parties, or because they are accepted for purposes of the conversation by the conversational parties in order to keep things simple. Or they can be held fixed even because *one* of the conversational participants takes the assumptions for granted and the others go along with it because disputing the assumptions would be a conversational distraction – as when an atheist engages in ‘God’ talk to use the Euthyphro dilemma (the presuppositions of which she does not herself endorse) to challenge a theist interlocutor’s unreflective endorsement of the divine command theory.

When conversational participants can avail themselves of a fixed set of background assumptions, they can have fruitful and meaningful engagements even if they have not fully fixed the meanings of their terms, and even if they would choose to fix them differently, were they to engage with interlocutors who accept a different set of background assumptions. It is possible to have meaningful engagements – even quite precise engagements – without converging on a single meaning, because sometimes different meanings *have the same conversational upshot*, holding the common ground fixed. Even when two definitions of a terms are not equivalent, they can be *conditionally* equivalent, given some assumptions. And so when those assumptions are part of the common ground, it will not matter for purposes of the conversation that the speakers have not converged on a single “once and for all” meaning for their terms. Likewise, it can make perfect sense for a single speaker to latch onto different definitions for different conversations – and to do so, even, without equivocating in what they care about, because each of their definitions can each be conditionally equivalent to what they care about, conditional on what is held fixed as part of the common ground in each of their different conversations.

In addition to assumptions being held fixed for purposes of conversation, conversational participants also bring to the table a set of imagined possibilities. *Pace* some naïve construals of the Lewis-Stalnaker picture of inquiry as “locating” ourselves in possible world space, we do not start with a fixed understanding of all of the possible ways that things can be and then gradually narrow down. Rather, we start with some clear distinctions that make sense to us, sharpened against some clear foils that we want to deny. The *process* of inquiry then proceeds, dynamically, from there – and as it does, tends to reveal to us theoretical possibilities that we had not yet considered, yielding in turn the need for finer-grained distinctions. Conversational participants could get along very well without deciding, once and for all, whether a spork is a fork but not a spoon, or a spoon but not a fork, or both, or neither, so long as none of them have ever imagined a spork.



Finally, disagreement is pervasive in philosophy – both actual, and potential. Wherever we use a common vocabulary to define our terms, that vocabulary becomes an object of philosophical study, and its properties will become controversial. It is a consequence of this that sometimes what is at issue between two philosophers is in part a matter of what each of them is committed to in virtue of the theory that they accept. If these philosophers can appeal to at least *some* common ground assumptions that they both share (or are willing to countenance for the purposes of the conversation), then they can try to use those shared assumptions to triangulate on what is really at stake between them. Even philosophers who disagree on *a lot* can make progress this way, if they are sufficiently willing – for the time being – to grant, bracket, or otherwise “gloss over” orthogonal points of dispute. Triangulating via common ground thus requires a kind of *pro tem* flexibility or provisionality of framing: the more rigid we are at the outset of inquiry, and the more necessary we take it to be that our dispute be cast, *now*, in terms which will survive any controversy *later*, the less in common we will have to triangulate with in the meantime.

In the next three sections we will survey how each of these ways in which the common features of a limited number of conversational participants facilitates successful engagement with ideas have played out in debates – both substantive and terminological – about moral realism.

2 Conversational Common Ground

Moral realism seems like it *should* be easy to define. However we end up ironing out the details, surely the *gist* is clear enough: that some moral claims are *objectively* true, that there are moral properties *out there*, that those properties *ground* moral facts, and that those moral facts are true *independently* of what we think or feel about them. Indeed, this is essentially everyone’s first pass – here are just a few representative examples:

- *Michael Rea*: Moral realism is the view that there are objective moral facts. There are objective moral facts only if the following two conditions are met: (i) there are moral properties— e.g., properties like being a right action, being a wrong action, being praiseworthy, being depraved, and so on—at least some of which are exemplified by actual objects or events, and (ii) the exemplification of a moral property *p* does not entail that anyone has beliefs about what exemplifies *p*, about whether *p* is exemplified at all, or about the conditions under which *p* is exemplified. Condition (ii) is meant to express part, but only part, of what many philosophers aim to express by phrases like ‘moral properties are not mind-dependent’ or ‘moral facts are not theory-dependent’.³

³ Rea [2006, 215-216].

- *David Enoch*: Robust Realism is an objectivist, response-independence view of normativity. [...] Whether or not a given normative statement applies (for instance) to a given action does not depend on what attitudes regarding it – cognitive or otherwise – are entertained by those judging that it is (or is not) or by anyone in their environment, nor does it depend on the attitudes, desires, and the like of the agent whose action it is or of anyone in her environment.⁴
- *Schroeter and Schroeter*: What is crucial to the realist position – whether Moorean or naturalist – is the claim that normative terms have a determinate reference and signify a specific property. [...] Call this the *Univocity thesis*. It's also distinctive of the realist position that speakers' opinions about what falls into the extension of normative terms are fallible. One's judging an action to be right does not make it so: there is an independent standard of correctness for normative judgments to which speakers are answerable. Call this the *Objectivity thesis*. Moorean and naturalist realists accept both the Univocity and Objectivity theses and take them to be constraints on an adequate realist account of the signification of normative terms.⁵

And why not start here? We have to start the conversation somewhere, after all, and this seems as natural and unproblematic a place as any — right?

As a sociological fact about philosophers, or at least the ones doing metaethics, “property-talk” seems indeed to come naturally to almost everyone. So *common* is the belief that properties exist at all that, for most metaethicists, it looms far in the conversational background, out of focus. But that something is even *widely* “common philosophical ground” does not, it's worth highlighting, actually make it *neutral* philosophical ground. Before most debates about moral realism even get off the ground, nominalists who deny the existence of any properties at all are bracketed out.

David Enoch, whose characterization above was the only one of our three not to use the word ‘property’ explicitly, makes this point explicit before going on to offer his more official characterization of realism in terms of properties:

Let me say, then, that according to Robust Realism, *and general doubts about properties aside*, there are irreducibly normative properties; similarly, *general doubts about facts aside*, there are irreducibly normative facts; and so on.⁶

That he gets away with this is no surprise (and for the same reason, no big deal) given that, *for the purposes of most conversations about metaethics*, it is common ground among the conversational participants that nominalism is

⁴ Enoch [2011, 3]

⁵ Schroeter and Schroeter [2005, 5].

⁶ Enoch [2011, 5 (*italics added for emphasis*)].

false, and that at least *some* properties exist – whatever else might be debated about “moral” ones. And...that’s fine!! Indeed, it’s precisely the point. What we wish to highlight, here, is not that philosophers like Enoch have *failed* in their metaethical theorizing because their definitions of “moral realism” are not sufficiently neutral between *every* participant in Philosophy’s “grand conversation”, but that our metaethical disputes can (and already do!) proceed perfectly fine *without* such a definition. We do not need to fix on a single, once and for all, meaning for metaethical inquiry to get off the ground.

So it is perfectly fair for people who agree that nominalism is false to characterize the terms of their disagreement in a way that presupposes the existence of at least some properties. They can do so without deciding whether one of them – the irrealist – would automatically win if nominalism is true and there are no properties at all, and they can even do so if they agree that their disagreement would survive the discovery that this presupposition is false, and nominalists could have essentially the same dispute as they are having, cast in slightly different terms. Since they each believe that this presupposition is true, however, they believe that there is no need for them to go down the route of having to work out how to recast the terms of their dispute until they actually begin having it with someone who rejects this presupposition.


But more: even if we know how to formulate what is at stake over whether moral realism is true in a way that is orthogonal to whether any properties exist at all, it is likely that whatever we build into this formulation in order to be sufficiently careful about this issue will only make our formulation more complicated. But complicated things are more difficult to think about, and so formulating the issue in these terms will only make it harder for us to decide whether moral realism is true. And if it is common ground among us that some properties do exist, then this is a cost that buys us no benefit that any of us think is worth having! So it is not only possible, but often advisable, to take advantage of what we accept as common ground in order to simplify the terms of our dispute.

3 Failures of Imagination

In the last section we argued that there is nothing wrong with characterizing the terms of your disagreement with someone in terms that presuppose something that is, or for all that you are sure could be, orthogonal to what you are actually in disagreement about. This is a feature of ordinary conversation, and so it is no wonder that it should be a feature of philosophical conversations, which after all are just ordinary conversations that happen to be about philosophy. But in addition to knowingly presupposing things that you either actively agree about or are willing to grant for the purposes of conversation, conversational

participants also frequently – indeed, we would argue, invariably – leave unresolved issues that they have not yet even imagined as possibilities that need to be distinguished.

It is so easy to fail to imagine even relatively obvious possibilities that this fact can be easily exploited for humorous effect:



This works as a joke because (sorry, as philosophers we have to explain this) as you read it the words ‘facing each other’ prime you to expect that indeed, one person’s left is the other person’s right, and so when the author describes people as “thinking” this, you think ‘of course they think it – it’s true’. And then the humor comes from obviousness in retrospect that this is not, in fact, generally true, but only true under the ordinary kind of circumstances that we imagine when we imagine two people facing one another.

The same thing, we know, has happened repeatedly in the recent history of metaethics. Before Mackie [1977], it was easy to define moral realism in contrast to noncognitivism or subjectivism, but after Mackie we needed to incorporate some specification that insists on whatever we take error theorists to deny – for example, that there really are moral properties, that some (positive) moral claims really are true, or the like. When we see Michael Rea [2006] explaining that some moral properties must actually be exemplified, we all understand that the error theory is what he is trying to rule out. And so we get on, pretty well, most of the time and for most purposes, by understanding that he is trying to rule out views that are relevantly similar to Mackie’s from counting, without even needing to worry about whether his definition actually successfully rules out all such views.

And that is a good thing! For even once we have imagined the possibility of an error theory and decided that it is inconsistent with what we meant all along by ‘moral realism’ or at least what we should have meant, much room remains to imagine different ways in which the moral error theory can be understood or developed. The differences between different sorts of views in this space and the grounds on which they get to count as sufficiently relevantly similar to Mackie’s⁷ to count as versions of the error theory are relatively

⁷ Or even whether Mackie’s own view is determinate between various possible interpretations.

diverse, and this diversity requires some finesse in formulating a single condition that does not just say “but nothing relevantly similar to Mackie’s view is true”.

For example, on the presuppositional version of the error theory formulated by Perl and Schroeder [2019], moral claims *do* attribute properties that actually exist and in many cases *are* actually exemplified – they just carry presuppositions *about* those properties that are false. Perl and Schroeder’s presuppositional formulation of the error theory is not ruled out by Rea’s formulation, but we conjecture that it is intended to be. Is that a problem? We say: no – it was not a problem while we didn’t imagine the possibility of the presuppositional form of the error theory, and it isn’t a problem after, so long as it is either common ground that the presuppositional error theory is false, or else that we should interpret Rea’s formulation as successful at what it was obviously *trying* to do – to rule out views relevantly similar to Mackie’s error theory.

The same pattern arises again with respect to metaethical contextualism. Both Rea’s [2006] and Enoch’s [2011] formulations of moral realism are consistent with Stephen Finlay’s [2014] radically contextualist end-relational semantics, and on such grounds Finlay classifies his own view as realist. On Finlay’s view, there *are* properties ascribed by moral terms in context, in context it is true to say that these are “properties like being a right action, being a wrong action, being praiseworthy, being depraved, and so on”, and in the vast majority of contexts it is true to say that whether moral statements are true does not “depend on the attitudes, desires, and the like of the agent whose action it is or of anyone in her environment.” Yet on Finlay’s radical end-relational theory, these properties could be virtually any ordinary property, in the right sort of conversational context, and it *does* typically depend on our attitudes or desires which of these attitude-independent properties we happen to be talking about in each conversational context.

Rea and Enoch are not thinking about the possibility of a view like Finlay’s. Both of them, we predict, would be happy to rule it out, but they are just not quite imagining why the things that they have said do not suffice to rule it out. In contrast, it is clear that Schroeter and Schroeter [2005] *do* intend to rule out views like Finlay’s. That is why they require moral terms to be *univocal*. But the matter of how to rule out all such views without mistakenly committing moral realism to something obviously false is somewhat delicate. As Finlay points out, many words that we use to make moral claims are in fact *not* univocal, but on the contrary can make moral claims in one context while being used to make non-moral claims in another context. This is certainly true for ‘ought’, ‘must’, ‘may’, and ‘reason’, for example. So if moral realism is not to be inconsistent with the obvious observation that some words have both moral and non-moral uses, then even more care is required, in order to say what kinds of context-dependence are and are not inconsistent with moral realism.

It is no doubt a worthwhile task, for some purposes, to try to finesse such refinements, to find the right formulation of moral realism that can rule out Finlay's end-relational semantics without saying anything obviously false. But for many purposes we can get on just fine by understanding that Schroeter and Schroeter intend to rule out similar views, and go from there, without worrying about whether they have succeeded at securing a precise definition that does the correct ruling out.

And more importantly, if it takes such care to carefully enough formulate moral realism to rule out the error theory and sophisticated forms of contextualism even after we do imagine them, we should not be very confident in our ability to successfully formulate conditions that will rule out the possibility of new creative views that exploit yet other possibilities that we have been so far unable to imagine. For example, the idea that a genuine form of relativism about truth might be coherent and applicable to metaethics was not taken seriously by many until the last twenty-five years. As a result, most philosophical defenses of 'moral relativism' in the last three decades of the twentieth century precisified what they were defending as a kind of contextualism. Hence attempts to define moral realism to rule out relativism have often tried only to rule out contextualism, instead. But in the last twenty-five years, undeniably coherent forms of genuine relativism have been developed that have different commitments from contextualism, and *these* must now be ruled out by any formulation of realism that means to rule out relativism.⁸

Or consider the research program of dynamic semantics, according to which the meaning of a sentence is given by its potential to change the context. Dynamic semantics offers a more flexible semantic framework with some similarities and differences from both relativism and expressivism, but despite several decades of history in linguistics, it has only in the last ten or fifteen years begun to enter the consciousness of metaethical theorists⁹, and so it is no wonder that many attempts to rule out similar views from counting as 'moral realism' do not yet rule it out. And of course there are many other possibilities we have not yet thought of. If it is hard enough to carefully rule out the things we know about, our confidence that we can successfully rule out the things that we don't even know about in advance should be very low.

And again, we say: that is okay. Indeed, it is more than okay. The kind of care and sophistication required in order to rule out sophisticated and creatively surprising views about moral talk, thought, and reality can only obscure what is really at stake. To the extent that our definition of what is at stake incorporates complexities that are there only for the purpose of ruling out possibilities that we are assuming are not true, these complexities can only get in the way of our thinking clearly about the issues. So it is better,

⁸ See especially MacFarlane [2014].

⁹ Compare Starr [2016].

really, to stick to simpler formulations. Until, that is, those complexities become relevant. But we can take that as it comes.

Sometimes, of course, once we imagine new possibilities – just like once we change our minds about things we had both taken for granted as part of the common ground – we find ourselves pushed in different directions about how to use classificatory terms like ‘moral realism’. Finlay, for example, is moved to continue to use it in a way that includes his own view, while Schroeter and Schroeter are moved to try to count it out. That’s fine – that just means that once this possibility is taken seriously they care about different things. It doesn’t follow that they couldn’t have been caring about the same thing so long as they were not taking this possibility seriously!

4 Irreconcilable Worldviews

And this brings us to the crux of the matter. We have seen in the last two sections not only that we *can* have substantive and precise disputes even without formulating them in terms that absolutely anyone could accept, but that it is often advisable to do so. First and foremost, relying on our common ground allows us to formulate what is at stake between us in ways that render it more tractable. This reason applies even when we are fully aware that there are other people who would reject the assumptions that we hold in common and even if we think that they could also, in principle, join in what would be recognizably the same dispute that we are having. But second and also important, our capacity to continually imagine new and creative possibilities for what philosophical views in some domain might look like means that it will likely be futile to even try, in advance, to iron out the wrinkles in a formulation that we might reasonably hope could stand the tests of both time and controversy.

This brings us to the problem that nearly everything is up for grabs in philosophy. Even if we wanted to formulate an issue in terms that could be accepted by everyone no matter what their other philosophical views, we would need some agreed-on terminology in which to formulate it. But since everything is ultimately up for grabs among philosophers, there is and can be no such agreed-on terminology.

For some purposes we might hope that little hangs on this. Surely, we might think, arcane questions in the metaphysics of properties and of propositions, the theory of grounding, the semantics of truth, the theory of explanation, and so on must be independent of the central questions of metaethics, and there *must* be ways of re-formulating our questions to get around whatever complications arise within the special concerns of each of these domains. But sadly, it turns out that we have excellent inductive evidence that this is not the case. Even where questions like these are *logically* independent of the concerns that animate us in

metaethics, it frequently turns out that novel and initially surprising (to some of us, at least) views about fundamental topics in other areas of philosophy end up being key to defending the plausibility of metaethical theories. This is one of the central contributing factors to what makes philosophy *hard*, and for many of us, it is an essential part of why thinking about metaethics in particular, can be so rewarding.

In the last two sections we argued first that we *can* conduct perfectly respectable inquiry without fixing terms once and for all, and second that it is *advisable* to do so, even if we could work out the kinks in advance. But the fact that philosophers can disagree about everything has an even more striking consequence: it is *impossible* to define our terms in a way that can make our dispute proof against the development of future kinks. For every dispute must be formulated in some way or other. And any way that we formulate it will use some words. But there are no words that constitute a privileged basic vocabulary in philosophy whose significance or connection to each other cannot be contested. So no matter which words we use to formulate moral realism, there will be possible views formulated in terms of these words that we have either not anticipated or whose defensibility we have not anticipated, which, if taken seriously, would undermine our attempt to use those words to formulate our original dispute in a way that is agreeable to all parties.

Many philosophers, of course, do work with a basic set of ideology within which to formulate their own sense of the issues. And some of these philosophers believe either that the ideology that they so use – or at least, some ideology which they hope is the one that they use – is privileged in some way that makes formulations of the issues couched within this privileged ideology to be privileged ways of understanding the issues.¹⁰ We both have some healthy skepticism for this idea, however formulated or developed. But what is important for our purposes here is just that *even if* there is some privileged vocabulary in which disputes ought to be formulated in order for us to see perspicuously what is at stake, the fact that some parties to the dispute do not *themselves* accept this privileged way of formulating it means that it cannot be used as common ground in order to triangulate what is at stake between them. And yet, for all of that, there may still be something else – something non-privileged, according to the proponent of the privileged vocabulary – that we *can* use to triangulate with both parties to the dispute, because it can be agreed on or at least accepted for purposes of the conversation, by both parties.

So we conclude that it is not only inadvisable, but in general impossible, to fully fix the terms of our disputes in advance. Philosophers who look like they are trying to do so should be charitably interpreted as just trying to fix them for a wide enough audience for the purposes at hand. And that, we say, is good enough – and often more than is strictly speaking needed.

¹⁰ Compare Sider [2011].

5 Pro Tem Inquiry

So what, then, does it look like, to conduct philosophical inquiry in metaethics or about any other topic without being able to fix once and for all the terms that we are using to ask our questions and in which to express what is at stake between us? Our answer is that it looks, more or less, exactly like what we actually observe. Indeed it must, because as we have argued, there is no other way of doing things.

Since emotivism became widely visible in English-speaking philosophy in the 1930's, it has been common ground among many philosophers interested in the questions of metaethics that whatever moral realism is, it isn't that. And so the question of what is at stake between views that fall in some sense into the same family as emotivism – noncognitivist views, for lack of a better name – came to be very closely associated with what is at stake over moral realism. That realism entails cognitivism, in the very loose sense where 'cognitivism' is just the opposite of whatever Ogden and Richards [1926], Ayer [1936], Stevenson [1937], Hare [1952], Blackburn [1984], Gibbard [1990], Ridge [2014], and Charlow [2015] have in common, has become common ground in attempts to formulate the thesis of moral realism – all such attempts make some attempt that is believed or hoped to encompass a sufficient condition for noncognitivism, so understood, to be false.

Yet we know from experience that what this condition has been taken to be has varied over time, as it has become increasingly difficult to secure enough common ground in order to say what is at stake over noncognitivism itself.¹¹ In the early heady days of emotivism, of course, when the readily accepted definition of 'noncognitivism' was that moral sentences cannot be true or false, people were taking for granted that the concerns of metaethics were sufficiently independent of philosophical theories of truth that we can use the vocabulary of truth in order to mark important distinctions. But once disagreements between correspondence and deflationary theorists about truth became prominent enough in metaethics, it became impossible to conduct this debate in terms of truth – too little was common ground.¹²

Those who accepted the correspondence theory of truth could still rationally believe that formulations in terms of truth were adequate to their ends, and they could still succeed at marking distinctions in ways that were communicatively successful in speaking to one another. But even if they are right that the correspondence theory of truth is true, speaking in this way simply came to exclude a larger

¹¹ Compare Schroeder [2010, chapters 2-4].

¹² See, in particular, how this issue played out in Boghossian [1990], Wright [1992], Horwich [1993], and Smith [1994a] (just to take a sample of the literature engaging with this topic at this time).

and larger number of other people from the conversation. And so in order to include those people, new ways of characterizing what was at stake over noncognitivism had to be invented.¹³

It is our contention that this event in the recent history of metaethics was not a breakdown or failure of inquiry, but rather an example of inquiry going well. True, it took some struggle in the 1990's to re-characterize the issues about noncognitivism in ways that were independent of considerations about truth and still captured the disagreements that people had already been having. And true, much of this work is still ongoing. But the right time to do this work was when it became important in order for cognitivists and noncognitivists to still be able to talk to one another – not when this was a bare abstract possibility.

What we mean by this is that it was fine to bracket deflationism about truth – even for those of us who officially thought that deflationism was actually true – so long as it gave us a way of characterizing what was at stake that was not prejudicial to either side. But now that deflationism about truth has come to be commonly accepted by noncognitivists, and come to be used as a *defense* against some arguments against noncognitivism leveled by cognitivists, we need to use something else to triangulate in order to have a conversation to which both cognitivists and noncognitivists – and hence to which both moral realists and their opponents – can be parties.

6 **Morals**

Some theorists seem to be genuinely troubled by the thought that if we're not *all* talking about *the very same thing* when we talk about, e.g., “normativity,” or about “moral properties,” then our disputes in metaethics about such topics are somehow “not real disputes” – that we are ultimately “talking past each other”.¹⁴ And it is this troubling thought, more than (just) an intrinsic desire for clarity and rigor, which we take to be drive recurrent ambitions in metaethics to define moral realism once and for all. And others, starting from the observation that we can't all be talking about just one thing, recommend that we simply distinguish between the many *different* things that we could talk about that are often run together, and recognize that each of these is, in its own right, an interesting thing to talk about.¹⁵

We do not claim that this ambition to fix terms once and for all is universal – indeed, many of the authors whose definitions of “moral realism” we cited in this paper are very careful to steer clear of any such ambition. Nevertheless, it is an alluring pitfall that we see authors fall into again and again. And if we have any major takeaway to offer, here, it is that the troubling thought behind it may be safely put to rest: we do

¹³ Compare Smith [1994b].

¹⁴ See, for example, Parfit [2011], Finlay [2019].

¹⁵ See, for example, McPherson [this volume].

not need to fix on one meaning, or one universal definition, for our disputes about “moral realism” to be meaningful, or for us to make progress. Perhaps it should be the *aim* of our metaethical theorizing *eventually* to arrive at such a “clear and distinct” articulation. But we do not *need* to establish it before we get started; and more to the point, we should not *expect* ourselves to be able to.

The principal obstacle, after all, to getting things precise is that we often don’t know *which things need refining* until their (relative) crudeness becomes salient to us. This has been the lesson of the previous sections: that in spite of even our most valiant efforts to precisify and pre-empt, we can – and usually do – have cognitive “blindness” to the deficiencies and excesses of our present definitions. We can fail to see where they are ambiguous, where they are overly-general, and where they are under-imaginative – and we can do so for utterly prosaic reasons. Perhaps it is because we share so *much* common ground with so *many* interlocutors that we mistake our background assumptions for universal givens. Or perhaps we share so *little* common ground with our perceived “opponents” that we have failed to imagine how they, on a worldview so different from our own, might accept the very definitions we proffered to exclude them. Perhaps *none* of us, yet, have thought to think about the relevant parts of logical space.

These are possibilities not just in conversations about moral realism, but about any topic on which philosophical inquiry has not yet been concluded. They are, however, especially *salient* in the context of moral realism, where the conversation about what “counts” as realist has evolved (and continues to evolve) in surprising and even dramatic ways, in part because the philosophical problems of metaethics are so challenging that they have served as a perpetual engine of innovation in what philosophical assumptions might be used to address them. We should not conclude from this, though, that we cannot successfully talk *about* “moral realism,” or that the conversations we *are* having about it are somehow futile *unless and until* we can settle (for real this time!) *exactly* what we mean.

On the contrary, we should conclude that our conversations about “moral realism” are just like our conversations about anything else: that they are going to be only as productive as we are able to triangulate on the common ground between us and our conversational partners; and that the definitions we employ to *do* that triangulating are almost always going to be provisional at best. We ought to be wary, then, of being too high-minded or literal about the meanings of words like “realism,” both in our own use but especially in our interpretations of others. Likewise, we should be modest as we offer characterizations which, even to us, may seem glaringly inadequate in only a few years’ time.

References

- Ayer, A.J. [1936]. *Language, Truth, and Logic*. London: Gollanz.
- Blackburn, Simon [1984]. *Spreading the Word*. Oxford: Oxford University Press.
- Boghossian, Paul [1990]. 'The Status of Content.' *Philosophical Review* 99: 157-184.
- Charlow, Nate [2015]. 'Prospects for an Expressivist Theory of Meaning.' *Philosophers' Imprint* 15: 1-43.
- Enoch, David [2011]. *Taking Morality Seriously*. Oxford: Oxford University Press.
- Finlay, Stephen [2014]. *A Confusion of Tongues*. Oxford: Oxford University Press.
- [2019]. 'Defining Normativity.' in D. Plunkett, S. Shapiro, K. Toh (eds.), *Dimensions of Normativity: New Essays on Metaethics and Jurisprudence*. Oxford University Press, 187-219.
- Gibbard, Allan [1990]. *Wise Choice, Apt Feelings*. Cambridge, MA: Harvard University Press.
- Hare, R.M. [1952]. *The Language of Morals*. Oxford: Oxford University Press.
- Horwich, Paul [1993]. 'Gibbard's Theory of Norms.' *Philosophy and Public Affairs* 22: 67-78.
- MacFarlane, John [2014]. *Assessment-Sensitivity: Relative Truth and Its Applications*. Oxford: Oxford University Press.
- Mackie, J.L. [1977]. *Ethics: Inventing Right and Wrong*. New York: Penguin.
- McPherson, Tristram [this volume]. 'Metaphysical Structure for Moral Realists.'
- Nagel, Thomas [1986]. *The View from Nowhere*. Oxford: Oxford University Press.
- Ogden, C.K., and I.A. Richards [1926]. *The Meaning of Meaning*. New York: Harcourt Brace.
- Parfit, Derek [2011]. *On What Matters*. Oxford: Oxford University Press.
- Perl, Caleb, and Mark Schroeder [2019]. 'Attributing Error Without Taking a Stand.' *Philosophical Studies* 176(6): 1453-1471.
- Rea, Michael [2006]. 'Naturalism and Moral Realism.' In Thomas Crisp, David VanderLaan & Matthew Davidson (eds.), *Knowledge and Reality: Essays in Honor of Alvin Plantinga (Philosophical Studies Series)*. Dordrecht: Springer, 215-242.
- Ridge, Michael [2014]. *Impassioned Belief*. Oxford: Oxford University Press.
- Schroeder, Mark [2005]. 'Realism and Reduction: The Quest for Robustness.' *Philosophers' Imprint*.
- [2010]. *Noncognitivism in Ethics*. New York: Routledge.
- Schroeter, Laura and Francois Schroeter [2005]. 'Is Gibbard a Realist?' *Journal of Ethics and Social Philosophy* 1(2): 1-18.
- Sider, Ted [2011]. *Writing the Book of the World*. Oxford: Oxford University Press.
- Smith, Michael [1994a]. 'Why Expressivists About Value Should Love Minimalism About Truth.' *Analysis* 54(1): 1-11.

- Smith, Michael [1994b]. 'Minimalism, Truth-Aptitude, and Belief.' *Analysis* 54(1): 21-26.
- Stalnaker Robert [1984]. *Inquiry*. Cambridge: Cambridge University Press.
- _____ [2002]. 'Common Ground.' *Linguistics and Philosophy* 25: 701-721.
- Starr, William [2016]. 'Dynamic Expressivism About Deontic Modality.' In Nate Charlow and Matthew Chrisman, eds., *Deontic Modality*. Oxford: Oxford University Press, 355-394.
- Stevenson, C.L. [1937]. 'The Emotive Meaning of Ethical Terms.' *Mind* 46(1): 14-31.
- Wright, Crispin [1992]. *Truth and Objectivity*. Cambridge, MA: Harvard University Press.