THE CASE FOR

The Case for Normative Realism

Paul Boghossian

1. *Introduction*

In this paper, I use a detailed critical discussion of Matti Eklund's famous challenge (2017, 2025) to normative Realism as a basis for defending a certain way of understanding the view. Eklund's argument is not only meant to challenge our standard way of *formulating* normative Realism; it is meant to raise a serious doubt about whether there could be a coherent doctrine there in the first place. Seeing how to plausibly block Eklund's argument promises to teach us much about how normative realism should be understood.

Along the way, we will have occasion to engage important issues about how normative concepts should be individuated, about how their reference is fixed, and about whether there are, *pace* Quine (1951), some things that are not up for grabs.

2. *Normative realism and Eklund's two claims*

Normative realism is standardly formulated as follows:

(i) Fundamental normative principles are capable of truth and falsity.
(ii) Some such principles are true.
(iii) And those that are true are made true by mind-independent facts.

Let us call this the Standard Formulation of Normative Realism, or SFNR, for short; and (following Eklund) call any normative realist who is satisfied by SFNR a 'mere realist'. Now, 'normative realism' is, of course, a technical term; what's the *intuitive* idea that SFNR is trying to capture? Naively, we might put it like this: we want the world, and not something about us, to determine which fundamental normative claims are correct. We may call this the desire that normative claims be *objectively adjudicated*. SFNR is intended to capture this intuitive idea, and it does so by insisting that normative facts are *mind-independent*.

Analysis Vol. 00 | Number 0 | XX 2026 | 1–32 doi: https://doi.org/10.1093/analys/anag072
© The Author(s) 2026. Published by Oxford University Press on behalf of The Analysis Trust. All rights reserved. For commercial re-use, please contact reprints@oup.com for reprints and translation rights for reprints. All other permissions can be obtained through our RightsLink service via the Permissions link on the article page on our site—for further information please contact journals.permissions@oup.com.

Downloaded from academic.oup.com/analysis/advance-article/doi/10.1093/analys/anag072/8780913 by New York University Libraries user on 14 November 2026

> (Objective Adjudication) We want something *not ourselves* to determine the *correctness* of basic normative (e.g. moral) principles.[1]

Now, SFNR is compatible with a deflationary understanding of 'true', 'fact' and other semantical notions. And there are some philosophers who have felt that mere satisfaction of SFNR, with deflationary construals of those central notions, would be less than what they were after in being realists. I used to think that myself, though I am no longer sure. Anyhow, in my 1989 article on rule-following (Boghossian 1989), I introduced the label 'robust realism' to designate a view that combined SFNR with non-deflationary construals of the key semantical notions. Since then, David Enoch has made this notion famous in its application to the moral domain (I believe he understands it in much the same way I did).

Early on in his book, Eklund introduces yet another character: the ardent realist. This is what he tells us about him:

> If you are a realist of the kind I want to consider, you want to say that Bad Guy – someone who does bad things and is motivated by bad desires – is somehow objectively out of sync with how to conduct one's life: in no way are our differences with Bad Guy merely a disagreement in taste, or a matter of having different desires. Introducing a label, you are an ardent realist if you are a realist and this is part of your motivation. You think any acceptable normative realism should deliver the result that there is something important that one of you … objectively gets right and the other … objectively gets wrong. (2017: 1)

Going by this definition, the ardent realist wants any adequate formulation of 'normative realism' to capture what I just now called 'objective adjudication'. Since the mere realist wants this as well, it looks as though there is no daylight between mere realism and ardent realism.

Eklund's first important claim, though, is that in fact mere realism cannot capture objective adjudication and so cannot be considered ardent in the relevant sense.

> (Eklund's First Claim) Objective adjudication cannot be captured by SFNR.

This is already interesting enough. But Eklund has something much more damaging up his sleeve. And this is:

1 Of course, as we know, Eklund's argument has the intended effect of problematizing our use of any normative expression and so, in particular, our use of 'correctness' here to capture objective adjudication. However, it is perfectly legitimate to use our standard normative vocabulary at this stage since their use has not been problematized yet. This is an important dialectical point in the present context: we are entitled to believe that our current normative framework, which seems to be working well, is in good order unless we are given a powerful reason for doubting it.

Downloaded from academic.oup.com/analysis/advance-article/doi/10.1093/analys/anag072/8780913 by New York University Libraries user on 14 November 2026

Downloaded from academic.oup.com/analysis/advance-article/doi/10.1093/analys/anag072/8780913 by New York University Libraries user on 14 November 2026

(Eklund's Second Claim) There is nothing we can add to SFNR that will enable it to capture objective adjudication.

Obviously, if these two claims were allowed to stand, normative realism would be in a lot of trouble, since it would mean that it was impossible to coherently capture its main animating idea. Hence the need to assess the strength of Eklund's argument for his two claims. To explain his argument, we first need to say something about the notion of a 'normative role'.

3. *Normative roles*

Consider someone judging the familiar chestnut:

(No Abuse) It is wrong to inflict recreational suffering.

Let 'wrong' here mean 'morally wrong'. In making this judgement, an agent is not *just* making a claim but is *proscribing* inflicting purely recreational suffering. He is evaluating such an act negatively, recommending against it, looking down on it. It is those aspects of the judgement that make it a 'normative' judgement. It is natural to ask: in what does that proscribing function of the judgement consist?

As we shall see, an assumption of Eklund's argument is that there is some well-defined aspect of *wrong*'s conceptual role that grounds this proscribing function. This is not uncontroversial, but I am sympathetic to it and so will assume it myself. What might this aspect of the conceptual role look like?

A familiar thought, common in discussions both of judgement internalism and of expressivist views, is that it consists in the *downstream* consequences of making the judgement. We may call these downstream consequences the term's 'normative role'. There are several different ways of elaborating this idea.

Theorists in the naturalist, expressivist tradition tend to think of that conceptual role in dispositional terms. To pick a version to fix ideas: someone judging No Abuse would be *disposed* to avoid performing such an act, to criticize those who do, and to experience negative emotions towards those who do, including oneself.

To the extent to which I am sympathetic to the idea of a normative role, my preference would be to characterize the downstream consequences in *normative* terms, not in naturalistic ones. If there is a single claim that virtually everyone working in metaethics is agreed upon, it is that it is impossible to *reduce* normative *concepts* to naturalistic concepts. (This is a limited remark about the reducibility of normative concepts, not of normative properties or facts.) And although a dispositional account of the conceptual role of normative concepts is not, strictly speaking, a conceptual *analysis* of normative concepts in dispositional terms, it is close enough to incur some of the same difficulties as an outright reduction. For example, and as we know from critical discussions of judgement internalism (see Svavarsdóttir 1999), it is not impossible that someone should be fully competent with 'wrong', judge No Abuse, and yet not be motivated to avoid the cruel act or feel the requisite negative emotions.

These difficulties are avoided by the normative version of such a view. Such an alternative account might go something like this:

> (CRNA) To judge No Abuse is to be *committed* to avoiding inflicting recreational suffering, criticizing those who do, feeling negative emotions towards them.

To put the matter in terms of introduction and elimination rules, the idea is that the concept *wrong* is at least partly specified by being governed by the following kind of *elimination-rule*: if you judge that *x* is wrong, you are committed to giving *x* a certain kind of role in your thinking, motivations and emotions.

Of course, if our account of *wrong*'s conceptual role is to be compatible with a realist account, there had better also be some aspect of the conceptual role that grounds its *referential purport* – the fact that it aims to name a property and does not just perform an expressive function (see Schroeter and Schroeter 2021). For my purposes here, I will take that for granted.

Appropriate normative roles of a similar type could be specified for some of the other canonical normative expressions. For example, the normative role of 'is epistemically justified' might be given by a link between making the judgement that believing *p* is epistemically justified and being committed to believing *p*, relying on it in reasoning and praising it in normative discussion. Wedgwood (2007) has offered a normativist account of the normative role of the 'ought' of practical deliberation, *O*, the all-things-considered ought that settles what is to be done. It goes like this:

> Acceptance of the first-person proposition 'O<me,t>(p)' – where 't' refers to some time in the present or near future – commits one to making p part of one's ideal plan about what to do at t. (97)

For illustrative purposes, I have presented some particular normativist views of particular normative roles. However, the arguments that follow are compatible with any number of alternative views about the nature of these roles, including ones that incorporate such Kantian notions as universalizability, or such contractualist ideas as interpersonal justification.

4. *Eklund's challenge*

With this notion of a normative role in hand, let's turn to Eklund's argument for his two claims.

To begin with the first, why is SFNR alleged to be inadequate for capturing objective adjudication? You make a normative claim, say, a moral one. That claim has mind-independent truth conditions. And, moreover, or so we may suppose, there are mind-independent facts that settle whether those truth conditions are satisfied; and some of them are satisfied. Does not that amount to the objective adjudication of normative claims?

Downloaded from academic.oup.com/analysis/advance-article/doi/10.1093/analys/anag072/8780913 by New York University Libraries user on 14 November 2026

Downloaded from academic.oup.com/analysis/advance-article/doi/10.1093/analys/anag072/8780913 by New York University Libraries user on 14 November 2026

Eklund argues that more is needed because of the following possibility:

> (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' are associated with the same normative roles as our words 'good,' 'right,' and 'ought,' 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. (2017: 18)

Let us designate their concepts with an asterisk. Three points are crucial to Eklund's use of Alternative in an argument against the adequacy of SFNR to capture ardent realism.

First, the normative roles of the starred and unstarred concepts are exactly the same, so that when they say that something is wrong* they are proscribing it in exactly the same way as we do when we say that something is wrong.

Second, *wrong* and *wrong** are *not co-extensive*; while it may be wrong to inflict pain on a person just for fun, it may not be wrong* to do so. Crucially, this is not just a matter of the two communities having divergent *views* about what falls in the extensions of these terms; it is that the *extensions* of the two expressions are genuinely distinct. Why is this crucial? Because if it is a matter of straightforward disagreement about what falls under the concept *ought,* which we can be presumed to share, then the facts postulated by SFNR can be presumed to settle which party is correct. There would be no challenge to the adequacy of SFNR in such a scenario. For the argument to work, the starred and unstarred concepts must differ in extension. (As we will see below, though, that aspect also makes it difficult for the argument to work.)

Finally, from our point of view, their judgements about what is 'right' and 'wrong' are repugnant to us; we are in serious disagreement about moral issues. From our point of view, they are 'bad guys'.

If we grant all this, we would be able to argue against the adequacy of SFNR as follows:

(1) It is possible for there to be two communities, each making correct 'ought' judgements (relative to the 'ought' concepts that they respectively deploy), while disagreeing with each other.
(2) On SFNR, there is nothing in the world that can adjudicate between these conflicting 'ought' judgements because there is nothing that can adjudicate between these competing 'ought' concepts – i.e., there are no facts by virtue of which you ought to use *ought* versus *ought** (*No Privilege*).
(3) Hence, SFNR cannot capture objective adjudication.

Downloaded from academic.oup.com/analysis/advance-article/doi/10.1093/analys/anag072/8780913 by New York University Libraries user on 14 November 2026

To elaborate a bit, SFNR has no resources with which to prefer the unstarred judgements over the starred ones (or vice versa), since it has no resources with which to adjudicate between these competing sets of normative concepts, say, ruling one of them correct and the other incorrect. By assumption, both sets of judgements are capable of being true, and mind-independently so.

Depending on which set of concepts you are using, then, you will be committed to promoting utterly different act types, ones that are not only incompatible with each other, but rebarbative from the other's perspective. Hence, SFNR cannot capture objective adjudication and so cannot capture ardent realism. Thus, we have an argument for Eklund's First Claim.

But there is of course more in the offing. For what could we *add* to SFNR to get us the desired worldly adjudication of all moral matters? What further condition could we add to SFNR to capture the idea that one set of normative concepts is privileged over its counterparts?

Eklund argues that there is a dilemma here. If the added condition is normative, then it too will have a normative counterpart, and so a further question will arise about it: is 'privilege' the thing to have or 'privilege*'?

On the other hand, if it is not normative, it is not clear how it helps. For example, suppose we introduce the descriptive idea that some normative properties are more *joint carving* than others. This opens the possibility that our concepts may be more joint carving than theirs and so may provide a basis for choice.

By itself, though, this would do nothing. We would need to add the *normative premise* that our concepts should be as joint carving as possible. However, once we add that we would be back confronting the question of whether we should be 'should' users or 'should*' users, where should* users would aim at joint-carving* properties; and so we would not have got very far.

So, the additional condition required for SFNR to capture ardent realism can be neither normative nor descriptive. Thus, there is nothing for this additional condition to be and we thereby have an argument for Eklund's Second Claim.

This, then, is the challenge Eklund has presented us with. Obviously, the threat is serious. Can objective adjudication be saved?

5. *Does the argument work?*

The main claims I will argue for are these:

(A) Alternative is not possible: it is not possible for two terms with the same normative role to have distinct extensions.

(B) Even if it were possible to have two communities with normative terms possessing the same normative roles but with distinct extensions, there would be no clear sense in which the two communities could be said to be in genuine *disagreement* with each other, a claim that is also crucial to Eklund's challenge.

(C) Even if Alternative were possible, one cannot assert No Privilege, and so the argument fails in any case.

I am going to discuss these in reverse order. Although I will occasionally discuss normative concepts associated with particular normative domains – such as moral concepts or epistemic concepts – the main focus of the discussion here will be, and should be, the thinnest and most basic of our normative concepts, concepts like the 'ought' of practical deliberation or the concept of 'is a reason for'. If the argument does not work for those most basic concepts, it will have failed. It is one thing to show that our moral concepts have competing alternatives, and another to show that our all-things-considered ought has a competing alternative. For reasons that will emerge, I do not believe either claim. But it is important to bear the point in mind.

6. *The 'alarming symmetry'*

It is very tempting to describe Eklund's argument as I have done above (let us suppose we are talking about the 'ought' of practical deliberation):

(1) It is possible for there to be two communities, each making correct 'ought' judgements (relative to the 'ought' concepts that they respectively deploy), while disagreeing with each other.
(2) On SFNR, there is nothing in the world that can adjudicate between these conflicting 'ought' judgements because there is nothing that can adjudicate between these competing 'ought' concepts – i.e. there are no facts by virtue of which you ought to use *ought* versus *ought**. (Call this *No Privilege*.)
(3) Hence, SFNR cannot capture ardent realism (objective adjudication).

But there is an obvious difficulty with this way of stating the argument, one that Eklund is fully aware of (I may have short-changed him on this point in a previous discussion; Boghossian 2021). For how could anyone rationally assert No Privilege? Surely, we *ought* to use *ought* rather than *ought**; after all, *ought** contains horrible mistakes from our point of view as *ought* users. Any *ought* user like us would find No Privilege obviously false. How else to assert, though, what No Privilege is trying to get at if not with the use of our standard basic unstarred normative concepts? Whatever else No Privilege is to be, it has to be some kind of normative proposition – it has to say something like: there are no facts about which 'ought' concept we *ought* to be using, or that it would be *correct* for us to use, or that it would be *right* for us to use. But the minute it is stated this way, using our standard normative vocabulary, it seems to refute itself.

So, it looks as though there is a serious dilemma here: either we state Eklund's argument by using something like No Privilege, in which case the argument does not get off the ground; or we cannot state his argument.

Downloaded from academic.oup.com/analysis/advance-article/doi/10.1093/analys/anag072/8780913 by New York University Libraries user on 14 November 2026

Now, it is usually taken to be a big problem for an argument that it cannot be stated without self-refutation. For example, a very effective consideration against the viability of global relativism – the view that all facts are relative to something – is that it is self-refuting (see Boghossian 2006). Why is this any different? If you cannot state the problem without refuting yourself, why should anyone worry about it? Recall: the burden of proof here is squarely on the challenger.

Eklund does not see it that way. He believes that it is possible to recognize that there is a problem here without our being able to state it. About an earlier statement of the point I have been making, he says:

> Boghossian's main critical remarks concern … No Privilege. The main point he makes is the one to which I have just called attention: must not someone giving my argument, and hence affirming No Privilege, try to state something ineffable? And my response is the one given above. No, one must not rely on anything like that. The point is rather to bring out that, assuming the possibility of Alternative, the ardent realist wants to affirm a given claim of privilege, and to raise problems for this. To make this point one need not commit to No Privilege. (2025: 121)

We do not need to rely on the problematic claim of No Privilege for the argument to work, says Eklund. What should we rely on instead? He says:

> The main thing I can do in response to concerns on this score is just to lay out the challenge from alternative concepts and see what response I get. I can describe the possibility of conflicting alternatives to our actual concepts, note the ways in which there is a seeming symmetry between our actual normative concepts and possible alternative systems of concepts – and insofar as some realist upon reflection finds the symmetry alarming, she is an ardent realist. One possibility, of course, is that when doing so I fail to get the desired reaction of alarm. Since statements of realism have tended to be compatible with mere realism, self-described realists have not explicitly committed to what I am calling ardent realism. But luckily, I do often enough get the desired reaction. Theorists happily describe themselves and their opponents as ardent realists. The responses to the challenge from alternative concepts are not that it is misconceived and targets a strawman. (119)

Surely, though, there is something deeply unsatisfactory about prosecuting Eklund's argument in this way. What is it that we are supposed to think? That there is a deep challenge here if we feel it? And not, if we do not? Surely, what Eklund really believes is that we *should* feel that there is a deep challenge here. However, the minute he puts it that way, the problem disappears.

What we are stumbling across here is the fact that, concerning our most fundamental normative concepts, the ones that we need to use to raise the very fundamental questions being raised here, there is no rational way to raise

Downloaded from academic.oup.com/analysis/advance-article/doi/10.1093/analys/anag072/8780913 by New York University Libraries user on 14 November 2026

Downloaded from academic.oup.com/analysis/advance-article/doi/10.1093/analys/anag072/8780913 by New York University Libraries user on 14 November 2026

a question about their correctness. It is not possible for us to be rationally talked out of our basic normative concepts, even if Alternative is possible; and so there is no good way to ask whether they are the ones we ought to use. The question answers itself as soon as it is posed.

The pressing question is: what is the significance of this fact? Does it show us something deep about the role of the normative in structuring and framing all inquiry, including inquiry into normative matters? Or is it just a *pragmatic* problem akin to our not being able to assert 'No one knows anything' without self-refutation, while it might still be true that no one knows anything?

The reason Eklund, and various sympathetic commentators on his problem, have felt that this cannot be a fact of any great philosophical significance is that they are impressed by the alleged *symmetry* of the situation. We say about them: 'They are thoroughly wrong to do what they do'; and they say about us: 'They are thoroughly wrong* to do what they do'.

On its face, though, this does not look like such an alarming symmetry. First, strictly speaking there is no symmetry: we are saying that they are wrong, and they are saying that we are wrong*. What Eklund needs is for there to be a *normative* symmetry between us. But 'normative' is itself a normative word, and so cannot be used to say in what the symmetry consists.

But even if we set this point aside, there is a difficulty in seeing that the symmetry is 'alarming'. If we came across a community that held that counter-induction was the correct inductive practice (assuming that had not led to their early demise), we would conclude that they were nuts. We would not take our situations to be normatively symmetric. If they insisted that counter-induction is *constitutive* of what they mean by 'correct inductive practice', we might conclude that they mean something completely different by 'correct', and, therefore, that we have no real idea what they are saying, or whether it conflicts with anything we believe (more on this type of option below). Either way, that will not amount to rational pressure on us to change our ways.

So, the first and perhaps most serious problem with Eklund's position is that it cannot be stated without refuting itself. This puts it in the same company as other dubious views like global relativism and extreme scepticism ('No one knows anything').

However, let us set this important difficulty aside and assume that Eklund finds some way of convincing us that there is a difficulty here. Do the other aspects of his argument work as advertised?

7. *Is Alternative possible?*

Eklund distinguishes between two main strategies for responding to his challenge – the *direct* strategy and the *metalinguistic* strategy. The direct strategy would concede the possibility of Alternative but would try specifying some invidious distinction between the starred and unstarred concepts that would show one set to be superior to the other. He sees little hope for this

strategy because of his dilemma. He thinks the realist's best hope is the metalinguistic strategy, which involves denying that Alternative is possible. But he worries that it cannot be pulled off. I am more optimistic on that score.

As we saw above, there are two aspects to the use that Eklund's argument puts Alternative to. The first is that our most basic normative concepts can have normative counterparts – concepts, that is, which possess the same normative roles as those of our concepts, but which have different extensions. The second is that the judgements made with the starred concepts *conflict* with those made with the unstarred ones, so that We and They are in genuine disagreement about normative matters.

How could two normative terms that have the same normative roles end up with distinct extensions? Clearly, the thought must be that a term's normative role is not sufficient to fix its extension. Let us follow Eklund in dubbing the thesis that the normative role of a normative term *is* sufficient to determine its reference the thesis of 'referential normativity'.

To suppose that referential normativity is false is to assume that the correct theory of reference for such terms must appeal to some *further* factor, in addition to their normative roles, to fix their extensions. What could such a further factor be?

Two possibilities are salient from the extensive literature on the theory of reference, and both are invoked by Eklund: causal connections to the environment; or associated sets of 'platitudes', as on the so-called Canberra Plan, with the reference determined by fit with these associated descriptions.

The former idea construes normative terms on the model of natural kind terms. Just as 'water' gets its reference not solely in virtue of fit with its associated conceptual role, but also in virtue of what it is causally connected to in the context of the acquisition of the concept, so 'wrong' gets its extension in part in virtue of what it is causally 'tracking' in the context in which it is acquired or used. (Notice that I switch here to talking about moral concepts, rather than the basic ought of practical deliberation. This follows Eklund's practice and also gives his argument its best shot. It is much harder to make sense of these options with the most basic normative concepts.)

On the alternative Canberra view, normative terms have 'platitudes' associated with them that contribute to determining their reference. The thought here is that the starred and unstarred concepts would have different, and mutually incompatible, platitudes associated with them. While our 'wrong' would have 'It is wrong to inflict suffering just for one's own amusement' associated with it, their 'wrong' would not have that associated platitude but might, for example, claim that 'It is wrong to help a person in need even if it costs you nothing to do so'.

Before I turn to examining whether we need to reject referential normativity, I want to argue that even if we did allow that our concepts have normative counterparts, it is very hard to see how we might secure the other crucial aspect of Alternative, namely, that there is a genuine *disagreement* between Us and Them.

Downloaded from academic.oup.com/analysis/advance-article/doi/10.1093/analys/anag072/8780913 by New York University Libraries user on 14 November 2026

Downloaded from academic.oup.com/analysis/advance-article/doi/10.1093/analys/anag072/8780913 by New York University Libraries user on 14 November 2026

8. Genuine disagreement?

On the Canberra option, it is pretty immediate that We and They are simply talking past each other. Yes, the normative roles of their normative terms are the same as those of ours. But our concept *wrong* incorporates such platitudes as: 'It's wrong to inflict severe suffering on someone just for fun' and 'It is not wrong to help someone in need, other things being equal'.

While their *wrong** incorporates, *qua platitudes* about *wrong** – that is to say, *qua* things that individuate it as the concept that it is – either the *negation* of these platitudes, or sentences that immediately entail the negations of these platitudes.

But that immediately implies that *wrong* and *wrong** are very different concepts from one another, despite sharing a normative role, and hence that there is no guarantee that what they express with 'This is wrong' will conflict with what we express using the very same sentence.

It is easy to miss this point because it is easy to conflate the view that is under discussion (the Canberra view of the relevant normative terms) with the following very different view:

> Our *wrong* and their *wrong** are both individuated *solely* in terms of their normative roles; however, *wrong* users tend to believe that inflicting recreational suffering is wrong, while *wrong** users tend to deny this.

In other words, it is easy to conflate the Canberra view with the very different view according to which it is possible to have two communities that hold *straightforwardly conflicting* views about what is wrong. On the latter option, there would be genuine *disagreement* between the two communities.

However, and as I emphasized above, the straightforward disagreement scenario would be useless for Eklund's purposes, because it does not exclude the possibility envisaged by SFNR, that there are facts that settle which of the two communities is correct in their normative judgements. And so, on this scenario, there would be no argument against the adequacy of SFNR. Eklund's argument crucially depends on the *wrong* and *wrong** concepts being distinct; but that is also what gets in the way of their grounding genuine conflicts between the starred and unstarred judgements.

How do things look on the causal relations option? This time we are not varying the associated platitudes, but rather the properties in the environment to which the users of the two communities are differentially linked.

It is hard for me to imagine what this could mean, but I will let that pass for now, coming back to it below. Once again, though, if the concepts expressed by 'wrong' are individuated in terms of *both* their normative roles *and* their extensions, then *wrong* and *wrong** are distinct concepts and there is, therefore, no guarantee that our judgements will conflict with theirs.

There remains a third option. On this option, we individuate the concepts solely in terms of their normative roles. This would be analogous to individuating 'water' solely in terms of its narrow conceptual role, rather than in

terms of its extension, in a Twin Earth thought experiment. On this option, *wrong* and *wrong** will be the *same* concept. Will not that secure real disagreement?

Yes, but only if referential normativity is true. If referential normativity is not true, then our 'This is wrong' and their 'This is not wrong' need be no more in conflict with one another than 'It's hot here' said by me here and 'It's not hot here', said by you over there. Tokens of 'wrong' would behave like indexicals, picking up different referents depending on their respective contexts of acquisition; they would not have intrinsically associated satisfaction conditions.

Philosophers disagree about whether such narrowly individuated indexical thoughts are genuine thinkables. But however that pans out, the *wrong* and *wrong** judgements would once again not necessarily be in conflict with one another, despite being expressible by sentences that are negations of one another.

If we individuate the concepts in question solely through their normative roles, the only way in which the relevant judgements could be in conflict with one another is if referential normativity were true, so that 'wrong' and 'wrong*' picked out the *same* referent. But if that were so, Alternative would be impossible, and so Eklund's argument would fail anyway. So, no matter which way the concepts are individuated, there seems no good way to secure the key claim that They *disagree* with Us.

What Eklund may be counting on is that we will think that what he has described is a straightforward scenario in which we have a community promoting and proscribing things in the same way as we are, but disagreeing with us about what is to be promoted or proscribed. Such scenarios are of course possible.

But that is not what he has described. That other scenario requires referential normativity, a thesis which is not compatible with his argument. Without referential normativity and given minimal assumptions about what it takes for two judgers to be disagreeing with each other, we lose the disagreement.

9. *Referential normativity – causal relations*

Let us turn next to the most important aspect of Alternative, namely, the claim that each of our basic normative concepts has a possible normative counterpart – a concept that shares its normative role, but which has a different extension.

As we saw, for this to be possible, a normative term's normative role cannot suffice for determining its reference. It can contribute to determining it, but it needs help from some other factor if it is to secure a determinate referent. And as we also saw, there are two salient candidates for what that further factor might be: causal relations to the environment (the causal relations view), and associated platitudes (the Canberra view). In my opinion, neither

Downloaded from academic.oup.com/analysis/advance-article/doi/10.1093/analys/anag072/8780913 by New York University Libraries user on 14 November 2026

Downloaded from academic.oup.com/analysis/advance-article/doi/10.1093/analys/anag072/8780913 by New York University Libraries user on 14 November 2026

of these views is at all plausible in application to basic normative terms. As a result, I think that it is implausible that Alternative is possible.

Let us begin first with the causal relations view. The idea here is that there is a close analogy between the way in which our natural kind concepts have their reference fixed and the way in which normative terms have their reference fixed. In the case of the former, the reference is fixed partly by whatever natural property happens to be regulating the use of 'water' in the context of acquisition (thus, H2O in our case, and XYZ in the case of the Twin Earthers); and the thought is that much the same could be true of our normative terms. How would that work? Well, could not we say, as various writers have suggested, that we could have one community's use of 'right' regulated by a deontological property, and a Moral Twin Earther's use of 'right' regulated by a different, say, consequentialist property, even as both words share the same normative role?

The idea that normative terms might be Twin-Earthable in this way has been vigorously and, in my view, very effectively disputed by the Moral Twin Earth literature (see Horgan and Timmons 1991, and others). Their target was Cornell-style naturalist moral realism, but the points generalize to our scenario.

But there is a different problem with the bruited analogy, one that I have not seen discussed. The notion of 'regulation by a property' would have to work in very different ways in the natural kinds case and the normative concepts case.

In the 'water' case, the fact that our use of the word is 'regulated' by H2O and the Twin Earther's use is regulated by XYZ is, so to speak, entirely *blind*: the Earthly and Twin Earthly users of 'water' need do nothing active to have their use regulated by these different properties. All they need to do is use it according to its narrow conceptual role in the settings in which they acquired them; the world takes care of the rest.

But this surely cannot be the way in which our use of 'right' might be regulated by a deontological property and their use by a consequentialist one. It is not as though we are living in a deontological environment and they are living in a consequentialist one. (Moral facts, if they exist, are presumably necessary.) Rather, the thought must be that when you look at our applications of 'wrong', the best explanation for the pattern of applications that you see is that we are trying to follow some deontological rule, even if we cannot state it if asked; whereas when you look at their use of 'wrong', the best explanation of the pattern of applications you see is that they are trying to follow some consequentialist rule, even if they cannot state it if asked. We are trying to track a deontological property, and they are trying to track a consequentialist one.

This is more closely analogous to a scenario in which one community is instinctively drawn to, say, classical logic, while another is instinctively drawn to intuitionistic logic, than it is to the water case. In the water case, we are following the *same* rule but picking up different referents in the two environments. And this is so precisely because the rule we are following tells us: there

is a something, I know not what, that is exemplified by these local exemplars and that is what 'water' is a name for. And then, lo and behold, it turns out that that something is different in the two environments. In the water case, no amount of looking merely at the usage of the word 'water' will tell you that two different referents are involved. You would have to know, empirically, what the natures of the two substances are.

But that is not how it would be in either the logic or the normative cases. In those cases, you would be able to tell just by looking at the pattern of applications that one user is trying to follow a deontological rule, and the other one is trying to follow a consequentialist one. That is not a case of two users following the same rule and picking up different referents given the environment. It is a matter of them using different rules. Thus, the water analogy does not help us make sense of how the causal relations view might apply to the normative case.

10. The Canberra view and minimalism about normative concepts

Let us turn to the Canberra view, the other salient meta-semantic account of normative terms that might underwrite the rejection of referential normativity.

This view, too, seems false as applied to the normative case. It is incompatible with a thesis that I have argued for elsewhere on independent grounds, namely *minimalism* about normative concepts (see Boghossian 2025). Since minimalism will be important later, I will spend a bit of time rehearsing the reasons for endorsing it.

According to minimalism about normative concepts, basic normative concepts are individuated *solely* in terms of their normative roles. A normative role, as we have seen (and as is needed for Eklund's argument) does not carry any information about what falls under it: it consists simply in an elimination-rule rather than an introduction-rule. As such, it can be intelligibly applied to any act-type, however plausibly or implausibly – it can be used to coherently prescribe or proscribe anything whatsoever. Nothing in the nature or character of this normative role prevents it from being coherently used to make outrageously false moral claims – such as, for example, that it is not wrong to inflict suffering just for one's own amusement. Such a claim may be false, but it would not be conceptually incoherent.

Why should we believe minimalism about normative concepts? Let us call any moral concept that incorporates a view about which acts or traits fall in its extension a *Loaded* moral concept. By way of contrast, I will call any moral concept that consists merely in its outward facing normative role a *Minimal* moral concept. Thus, the question about minimalism is a question about whether our basic, thin moral concepts are Minimal or Loaded. How can we decide this question?

I believe that the fact that they are Minimal rather than Loaded is shown by the fact that we understand, and feel the need to respond substantively to, the

Downloaded from academic.oup.com/analysis/advance-article/doi/10.1093/analys/anag072/8780913 by New York University Libraries user on 14 November 2026

Downloaded from academic.oup.com/analysis/advance-article/doi/10.1093/analys/anag072/8780913 by New York University Libraries user on 14 November 2026

most extreme sceptical hypotheses about normative matters. We are not satisfied with purely 'semantic' solutions to such sceptical challenges.

Take the case of epistemology. If anything is a platitude about epistemic rationality, I suppose, it is that enumerative induction is rational and enumerative counter-induction irrational. On the Canberra view, it would have to be *analytic* of our concept of 'epistemic justification' that induction is justified and counter-induction is not. But we do not find satisfying a response to Hume's sceptical challenge about induction to say, as Strawson (1952) tried to do: inductive reasoning is just what we *mean* by 'epistemically justified reasoning'. However, if the Canberra view were true, this should be a perfectly satisfying answer to Hume's challenge, just as satisfying as it is to answer a 'sceptic' about the three-sidedness of triangles by saying that that is just what we mean by 'triangle'.

What is the difference between the two cases? Why is a semantic solution appropriate for the 'triangle sceptic' but inappropriate for the induction sceptic? The difference between them is that what we are really asking in the induction case is:

> Does induction deserve to have a certain kind of place (the kind specified by the normative role) in our reasoning and thinking?

This question can be asked by employing a completely Minimal concept of epistemic justification, one that was fully individuated just by its normative role and had no admixture of substantive doctrine. So, even if we could ask this question using a Loaded concept of 'epistemic justification', one on which the answer to our question would be analytic and trivial, that is not satisfying because we also have to hand a Minimal concept with which to ask the substantive question, and the answer to that would not be advanced in any way by the trivial answer to the Loaded question.

Someone might worry that all that this shows is that we can suspend belief in one platitude at a time without compromising our grip on the normative concept in question – in this case, that of epistemic justification – but not that we can suspend them all at once. Perhaps the situation here is like that of Neurath's boat – we can repair one plank while standing on the others, but we cannot repair them all at once.

But this seems not to be true in the normative case. It is not as though we must suspend scepticism about the external world or other minds in order to intelligibly contemplate a scepticism about induction. They can intelligibly be in play at the same time.

The same goes for morality. If the Canberra view were true, we should be able to respond to a challenge to justify No Abuse simply by saying that that is part of what we mean by 'wrong'. Instead, though, we understand the question and believe that it deserves a substantive answer. Of course, justifications do come to an end somewhere, so at some point we will be forced to say that some basic normative claim is just primitively true. But there is a difference

between appealing to a synthetic primitive normative truth and appealing to one that is alleged to be analytic.

If these considerations are right, they show that our basic normative concepts are Minimal, or that, if they are not, they may as well be, without the loss of anything that really matters to us in our use of them. What we care about, in asking a normative question about x, is whether x deserves to have a certain place in our thinking, deliberations and so forth. The freer we are in asking whether things deserve that role, the better served will our purpose be in having such a concept.

To put the point another way: even if our *actual* normative concepts are Loaded, there are clearly perfectly good *possible* normative concepts that are Minimal, and that would be enough to block Eklund's argument.

11. *Alternative is not possible*

The upshot of this discussion, then, is that neither the causal relations model nor the Canberra model is plausible for basic normative concepts. And if that is so, there is a very real question, it seems to me, whether Alternative is genuinely possible, since it is hard to see what other meta-semantic view could underwrite the possibility of Alternative.

We are not yet done, however, because Eklund maintains that even if normative role is taken to fix reference, and so that there are no normative counterparts to our basic normative concepts, this will not make his challenge go away, since there are challenges similar in spirit that can be raised regardless. He mentions two such challenges.

12. *A capricious God?*

The first involves a capricious God:

> A different problem that remains even if normative role fixes reference is a kind of wrong type of reason problem. To present the problem starkly, consider the possibility that a capricious God solely for his amusement has decided to let all concepts with the same normative role as OUGHT refer to some arbitrarily chosen property. Then normative role does indeed fix reference. But it still remains that knowledge of what ought to be done fails to provide the kind of normative guidance one might have expected. (Eklund 2025: 116)

This God is clearly in violation of No Abuse, but we will let that pass. What Eklund is thinking here is that referential normativity could be satisfied but, as it were, for the wrong type of reason. *Ought* and *ought** do co-refer, but they co-refer to some arbitrary property, one that, I suppose, need not even be normative in nature. If we could make sense of such an eventuality, there would indeed be a question why we should take correct ought judgements to normatively guide our actions.

Downloaded from academic.oup.com/analysis/advance-article/doi/10.1093/analys/anag072/8780913 by New York University Libraries user on 14 November 2026

Downloaded from academic.oup.com/analysis/advance-article/doi/10.1093/analys/anag072/8780913 by New York University Libraries user on 14 November 2026

But I cannot make sense of the stipulated scenario. First, there is presumably some fact or other about how the reference of normative concepts is fixed. And whatever that fact is, it is presumably a necessary fact. To suppose that God can make the terms refer to some arbitrarily chosen property that they do not in fact refer to is to suppose that God can do something that is metaphysically impossible. But he cannot do that (he cannot make triangles with more than three sides or bachelors who are married). Second, I do not really understand what it would mean for a concept to have some arbitrarily associated satisfaction conditions. We have the concept, say, *watermelon,* but God makes it the case that this concept is satisfied by all and only tables. I cannot wrap my mind around that.

13. The embarrassment of riches objection

Eklund's second idea under this rubric is somewhat more promising. He calls it the 'embarrassment of riches' problem.

What Eklund has in mind here is the possibility of pairs of normative terms which differ somewhat in their constitutive normative roles – enough to determine different extensions, but not so much that they cannot be seen to be competing alternatives. If we can have such pairs of terms, we will have terms that can compete with one another, even while not being normative counterparts.

What would this look like? Let us go back to discussing this issue primarily in terms of the thinnest and most basic of normative terms, such as the 'ought' of practical deliberation, the last 'ought' before action, as Gibbard puts it.

Eklund says that Gibbard (2003) overlooks the possibility that there might not be a unique last such ought, that there might be several different notions that could play the same role while differing in their normative roles in small but significant ways. He also comments on Richard Joyce's discussion of the practical 'ought'. Joyce wrote (as quoted by Eklund):

> Practical reasons are, by definition, those which guide our actions when everything has been taken into account: if the rules of chess are telling me to move my rook straight, the rules of etiquette telling me that I ought to let my opponent win, the rules of prudence telling me that I ought to move the rook diagonally (thereby annoying my opponent thereby winning a bet), then practical rationality is not a *further* consideration. Rather, when I have weighed all these claims and come up with an answer to what is to be done, then that is the judgment of practical rationality. (Joyce 2001: 50–51; quoted in Eklund 2017: 57)

About this Eklund writes:

> Following Joyce, let us use 'should_p' for the 'should' of practical rationality. Then there is a question of whether there are alternative terms that are in normative competition with 'should_p' – whether some

> community could use some notion 'should_p*' such that while this too is an all-things-considered notion, it is not coextensive with 'should_p.' If there is such an alternative possible notion … maybe considerations of what one 'should_p' ought to be abandoned in favor of what one 'should_p*' do. (2017: 57–58)

Of course, here we have once again the problem that, even if we concede the existence of these alternative notions, the question is no sooner stated than answered: what is that 'ought' in that last sentence if not 'should_p'?

But set that aside. Notice that Eklund raises the existence of these alternative last 'ought's as an abstract possibility: so far as I know, he has never described a concrete example. I think there is a reason for that. And that is that this notion, along with a few others, is so thin that it is hard to imagine what a competing alternative to it could be. Recall: the alternative had better be different enough in its conceptual role so that it is plausible that it has a different referent, but not so different that it no longer looks like the 'ought' of practical deliberation. It is very hard, it seems to me, to see how to satisfy both conditions in this case.

Look at Wedgwood's account of the 4-place predicate occurring in sentences of the form

> *x* is (all things considered) a better thing for *z* to do at *t* than *y*.

He proposes, as Eklund presents the rule, that the predicate is governed by the following basic rule:

> (Rule for Better Option) Acceptance of B(x, y, me, t) commits one to having a preference for doing x over doing y at time t. (2017: 55; cf. Wedgwood 2007)

What could we vary in this rule so as to get a competing alternative to this notion of 'better option all things considered'? What could we vary enough so that it is not just a tweak that would not result in a different extension, but not so much that it no longer looks like a notion of 'better option all things considered'? Perhaps I am not being imaginative enough, but I cannot see what it could be.

Similar remarks apply to the bare notion of 'being a reason for'. A reason to believe *p* is a positive consideration in favour of believing *p*; a reason for doing *A* is a positive consideration in favour of doing *A*. Once again, it is very hard to see how this bare, basic notion could have a competing alternative.

Of course, folks who grasp the concept of 'being a reason for belief' may disagree amongst themselves about what *counts* as a reason for belief. Some might be dogmatists about perception, holding that a perceptual seeming that *p* by itself, even in the absence of a justified belief about its reliability, could count as prima facie justification for believing *p*; others might be conservatives, holding that a background belief about reliability is required.

Downloaded from academic.oup.com/analysis/advance-article/doi/10.1093/analys/anag072/8780913 by New York University Libraries user on 14 November 2026

Downloaded from academic.oup.com/analysis/advance-article/doi/10.1093/analys/anag072/8780913 by New York University Libraries user on 14 November 2026

Some might hold that there can only be epistemic reasons for belief (reasons that bear on the truth of *p*), whilst others may hold that there can be non-epistemic, pragmatic reasons for belief. But these would be substantive disagreements about what falls under 'is a reason for belief', disagreements that we may assume would be settled by the normative facts postulated by SFNR.

What an Eklund-style argument needs, if it is to work, is two distinct concepts of 'is a reason for believing', which nonetheless count as competing alternatives. And the point I have been belabouring is that it is hard to see how to get two notions of 'is a reason for' that are sufficiently distinct from one another to determine distinct extensions, while being recognizably distinct versions of 'is a reason for'. This is hard to get, given the thinness of the normative roles at issue.

14. Enoch's reasons/schmeasons objection to Scanlon

What I have been saying here may seem to run counter to a famous objection that David Enoch has made to Scanlon's (2014) 'ontologically thin' conception of normative reality.

> In previous work (McPherson 2011; Enoch 2011, Section 5.3), we offered what we might call the 'schmeasons' (or 'counter-reasons') challenge to Scanlon. The basic idea is this. Consider two linguistic communities – the 'reasoners' and the 'schmeasoners.' Each community has a certain term – 'reason' and 'schmeason,' respectively – that they take to be central to their normative practices. And in each community there are well-developed practices of criticism and evaluation that use the relevant term. But the reasoners and schmeasoners tend toward quite different substantive views in their practices. If we suppose that these practices are coherent, and constitute their own domains, then both communities might be functioning quite well relative to their respective domains. But intuitively, we want to say that all is not well: it is bad that the schmeasoners are sensitive to schmeasons rather than reasons. This, however, looks like an objection that can be raised perfectly symmetrically from within each of the two domains. For the schmeasoners can point out that it is 'schbad' that we are sensitive to reasons rather than schmeasons. Our suggestion was that a metaphysically committal realism does a better job capturing what we intuitively want to say. This is because the metaphysically committal realist can say that only the reasoners – and not the schmeasoners – track the normative structure of reality. (Enoch and McPherson 2017: 835)

All of this is highly reminiscent of Eklund's argument. We have two communities, each with its own set of normative concepts, each doing well relative to its own practice, but where, intuitively, one of them is 'worse' off than the other. And obviously here too we face some of the same issues: can we so

much as state what is 'off' about the schmeasoners? Can there be such schmeasons? And if there can, are these two communities really disagreeing?

The problem in assessing whether Enoch faces the same issues as Eklund is that it is, as he and others have pointed out, entirely unclear what Scanlon means by a 'domain'. If we suppose, as Enoch and McPherson seem to be doing in this passage, that if a practice is internally coherent it therefore constitutes its own domain, then the objection to Scanlon can be pressed without replicating the problematic features of Eklund's setup. All that would be needed is for the reasoners and the schmeasoners to agree in the normative roles that they assign these words but to have, as Enoch and McPherson themselves put it, 'quite different *substantive* views in their practices'. In that case, normative reality will decide between them.

If, however, for whatever reason, that will not do, and we do have to deploy the apparatus of normative counterparts, then many of the same issues that got raised for Eklund's argument will arise here as well.

15. *Referential normativity defended*

So far, I have tried to make the case that Eklund's challenge can be resisted in various ways. Even if everything I have said is correct, it is not as though normative realism is in the clear. Perhaps we have shown that referential normativity has not been refuted; but we have not yet explained how it could be true. How does a normative role determine a reference?

In 2021, Schroeter and Schroeter published a paper entitled 'Bad News for Ardent Normative Realists?' In it they argued that there is in fact no good way to make sense of how a normative role could single out a particular property – such as maximizing happiness – as its referent. They argued that given certain plausible constraints on the assignment of semantic content, constraints having to do with the semantic content's role in psychological explanation and normative assessment, no determinate referential content could be determined on the basis of the normative role alone.

Without going into their arguments in detail, I will just say that I think that one can see their worry if we had to construe, as they in fact do, the normative roles of normative concepts in *descriptive*, dispositional terms. If all we can say is that, if you judge No Abuse, then you are *disposed* to avoid certain behaviours, to criticize them and feel negative emotions towards those who exhibit them, it is hard to see what specific property it would be appropriate to assign to 'wrong' as its semantic value. The property of being disgusting might qualify, and so would several others.

Things look more promising if the normative role is specified using normative materials. Recall my toy account of the conceptual role of 'morally wrong'.

> (CRNA) To judge No Abuse is to be *committed* to avoiding inflicting recreational suffering, criticizing those who do, and feeling negative emotions towards them.

Downloaded from academic.oup.com/analysis/advance-article/doi/10.1093/analys/anag072/8780913 by New York University Libraries user on 14 November 2026

Downloaded from academic.oup.com/analysis/advance-article/doi/10.1093/analys/anag072/8780913 by New York University Libraries user on 14 November 2026

If we were to ask what truth conditions we could attribute to a belief that would be expressed by No Abuse if those are the downstream commitments of making it, it would not be implausible to say:

> The judgment that it is wrong to inflict recreational suffering is true iff inflicting recreational suffering *deserves* to have a certain kind of role in our deliberation, motivation, normative discussions and emotions.

And at least for the non-naturalist realist, I do not see any great difficulty in supposing that normative reality contains such facts. Wedgwood (2001, 2007) has given a similar argument in the case of his practical ought.

I am not committing myself to the details of any of this; and it is anyway too sketchy to be satisfying. I am simply indicating that the prospects for securing referential normativity are not so dim, if we allow ourselves a normative construal of normative role at the conceptual level, and a non-naturalist normative realism at the metaphysical level, theses that I am independently attracted to.

16. Conclusion

Ever since Nelson Goodman (1955) invented *grue,* philosophers have thought that we can evaluate for correctness not only claims but concepts. After all, does not Goodman's case show that *green* is somehow or other a 'better' concept than *grue*?

In evaluating concepts, we should distinguish between evaluating the fitness of a concept *for a particular task* versus evaluating its correctness regardless of task.

It is easy enough to make sense of the first kind of evaluation. *Green* is better than *grue* for the purposes of formulating lawlike generalizations. I do not, however, believe there is some clear-cut, absolute sense in which *green* is better than *grue*.

Of course, some concepts, like *witch,* are empty; and so you may think that there is at least one clear-cut sense in which a concept can be criticizable absolutely: a good concept should not be empty.

However, even here, there is an implicit relativity to a task. If the task is to describe the world, then an empty concept will not serve you well; but if the task is to describe the propositional attitudes of the Salem witch-hunters, then the empty concept may be exactly the right one.

In any case, there is, as I have tried to argue, a special problem evaluating the correctness of normative concepts in particular, since correctness is itself a normative notion. If our most basic, thinnest normative notions had intelligible, competing alternatives, this fact might have led us into paradox, with there being compelling reasons both to dismiss the question about their correctness as well as reasons to take it seriously. But if what I have been arguing above is correct, the most basic of our normative concepts do not have intelligible, competing alternatives. If you are going to be in the normativity business at all, you will need the bare notion of 'a positive consideration in favour of', and that of the 'best thing to do overall'.

Of course, people will have different views about what counts as a positive consideration in favour of, and what counts as best thing to do overall. We can make sense of that as competing *substantive* views about what counts as a reason and what counts as having most reason to do. But we should not try to make sense of that by saying that we can have distinct but competing concepts of 'is a reason for' and 'best thing to do overall'. On that path lie all the difficulties we have uncovered above.

So far as I can see, the combination of a normativist construal of normative role, a minimalist construal of the individuation of normative concepts, referential normativity and non-naturalist normative realism is a stable and viable solution to the challenges raised by Eklund's important argument against the coherence of ardent realism. It is possible that other combinations of views could work as well.

Pace Quine, it is not the case that everything is up for grabs. In particular, there is a limit to the extent to which we can get underneath our most basic normative concepts and call them into question (for further anti-Quinean examples, see Boghossian 1996 and Boghossian and Wright 2024). As Thomas Nagel says:

> In order to have the authority it claims, reason must be a form or category of thought from which there is no appeal beyond itself – whose validity is unconditional because it is necessarily employed in every purported challenge to itself. This does not mean that there is no appeal against the results of any particular exercise of reason, since it is easy to make mistakes in reasoning. … But the corrections or doubts must come from further applications of reason itself. (1997: 7–8)

I believe that this is a profound insight. If reason is the final arbiter of what we are justified to believe on any topic, the only way to criticize reason is through some use of reason itself. That means that there has to be some fixed immoveable core to reason that is presupposed in any exercise of reason.

It is a fascinating and much debated question exactly what that core contains. Boghossian and Wright (2024), for example, argue that it will include at least some version of *Modus Ponens* and Universal Instantiation. But whatever one thinks about that, it is very hard to see how to do what Eklund is trying to do here, namely, raise a question about the 'correctness' of correctness itself.

Which is not to say that we should not be grateful to him for having pushed the point as far as it is possible to push it.[2]

New York University
USA
paul.boghossian@nyu.edu

2 I am grateful to Matti Eklund, Yu Guo, Peter Railton, Andrew Rubner, Declan Smithies, Josh Schechter, and Crispin Wright, and to audiences at Princeton and NYU for comments that helped improve this paper.

Downloaded from academic.oup.com/analysis/advance-article/doi/10.1093/analys/anag072/8780913 by New York University Libraries user on 14 November 2026

Downloaded from academic.oup.com/analysis/advance-article/doi/10.1093/analys/anag072/8780913 by New York University Libraries user on 14 November 2026

References

Boghossian, P. 1989. The rule-following considerations. *Mind* 98: 507–49.

Boghossian, P. 1996. Analyticity reconsidered. *Noûs* 30: 360–91.

Boghossian, P. 2006. *Fear of Knowledge*. Oxford: Oxford University Press.

Boghossian, P. 2021. Normative principles are synthetic a priori. *Episteme* 18: 367–83.

Boghossian, P. 2025. Minimalism, alternative normative concepts and normative realism. In *New Essays on Normative Realism*, eds. P. Boghossian and C. Peacocke, 163–79. Oxford: Oxford University Press.

Boghossian, P. and C. Wright. 2024. Kripke, Quine, the 'adoption problem' and the empirical conception of logic. *Mind* 133: 86–116.

Eklund, M. 2017. *Choosing Normative Concepts*. Oxford: Oxford University Press.

Eklund, M. 2025. Alternative concepts, ardor, and elusive questions. In *New Essays on Normative Realism*, eds. P. Boghossian and C. Peacocke, 112–40. Oxford: Oxford University Press.

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

Enoch, D. and T. McPherson. 2017. What do you mean 'This isn't the question'? *Canadian Journal of Philosophy* 47: 820–40.

Gibbard, A. 2003. *Thinking How to Live*. Cambridge, MA: Harvard University Press.

Goodman, N. 1955. *Fact, Fiction, and Forecast*. Cambridge, MA: Harvard University Press.

Horgan, T. and M. Timmons. 1991. New wave moral realism meets Moral Twin Earth. *Journal of Philosophical Research* 16: 447–65.

Joyce, R. 2001. *The Myth of Morality*. Cambridge: Cambridge University Press.

McPherson, T. 2011. Against quietist normative realism. *Philosophical Studies* 154: 223–40.

Nagel, T. 1997. *The Last Word*. New York: Oxford University Press.

Quine, W.V.O. 1951. Two dogmas of empiricism. *Philosophical Review* 60: 20–43.

Scanlon, T.M. 2014. *Being Realistic About Reasons*. Oxford: Oxford University Press.

Schroeter, L. and F. Schroeter. 2021. Bad news for ardent normative realists? *Ergo* 7: 1015–52.

Strawson, P.F. 1952. *Introduction to Logical Theory*. London: Methuen.

Svavarsdóttir, S. 1999. Moral cognitivism and motivation. *Philosophical Review* 108: 161–219.

Wedgwood, R. 2001. Conceptual role semantics for moral terms. *Philosophical Review* 110: 1–30.

Wedgwood, R. 2007. *The Nature of Normativity*. Oxford: Oxford University Press.