

The End-Relational Theory Does Not Vindicate Naturalism¹

Abstract: Analytic reductivism in metaethics has long been out of philosophical vogue. In *Confusion of Tongues: A Theory of Normativity* (2014), Stephen Finlay tries to resuscitate it by developing an analytic metaethical reductive naturalistic semantics for ‘good’. He argues that an end-relational semantics is the simplest account that can explain all of the data concerning the term, and hence the most plausible theory of it. I argue that there are several further assumptions that a reductive naturalist would need to make about contextual parameter completion, in order to derive reductive naturalism from an end-relational semantics – assumptions that nonnaturalists might forcefully resist. I also argue for the claim that an end-relational semantics could provide surprising resources for nonnaturalists to address semantic worries about their views – the upshot of which paints the way for a new and sophisticated nonnaturalism about the semantics of moral discourse. Nonnaturalists have long suspected that they need not worry about semantics. This paper ultimately supports that suspicion.

Introduction

The term ‘good’ appears in a bewildering variety of sentences. Some of these sentences seem to imply no more than that a pencil is sturdy (e.g. ‘This is a good pencil’) and that strawberries are sweet (e.g. ‘Strawberries are good’), while others seem to make distinctive claims (e.g. ‘Minimizing animal suffering is good’) regarding matters of value or importance (e.g. ‘It’s good that Osama Bin Laden was killed’). This makes it challenging to determine whether ‘good’ has a core semantic element that it contributes to the meanings of sentences in contexts of use. Some philosophers, including G.E. Moore² (1903, p. § 9-10), Russ Shafer-Landau³ (2003, p.66) and Derek Parfit (2011, p. 38-42), even suggest that the meaning of the term is fundamentally ambiguous.

Unfortunately for Moore, Shafer-Landau, and Parfit, the view that ‘good’ is ambiguous faces serious problems. The seriousness of these problems has led some to

¹ Special thanks to Steve Finlay, Mark Schroeder, and Caleb Perl for providing feedback on drafts of this paper.

² “Good...is incapable of any definition, in the most important *sense* of that word.” (emphasis mine)

³ “...non-naturalism comprises two essential claims: a metaphysical claim, to the effect that moral properties are sui generis, and not identical to any natural properties, and a semantic claim, to the effect that moral terms cannot be given a naturalistic analysis.”

deride this sort of view, including J.L. Mackie⁴ (1977, p.51), who found the idea “most implausible”, citing cross-linguistic evidence as counting against it, and Paul Ziff⁵ (1960, p.203), who thought even less of it, noting that it would be “absurd” for “...one would be forced to say the same about a great many other adjectives.” These problems alluded to by Mackie and Ziff suggest that a theory of the meaning of ‘good’ has to entail that the term makes a single kind of semantic contribution to the meanings of sentences in which it appears, across all contexts of use, on pain of falling short of accounting for the intuitive and empirical data concerning the meaning of the term.

In *Confusion of Tongues: A Theory of Normativity*, Stephen Finlay (2014) uses this Mackian/Ziffian observation in order to indirectly argue for a surprising form of analytic reductive naturalism in metaethics.⁶ According to Finlay, the view of the meaning of ‘good’ which makes the best sense of how it could make a uniform semantic contribution across such a diverse range of uses appeals to what he calls an “end-relational” semantics. According to this semantics, the word ‘good’ contributes a single naturalistic relation to the proposition that a sentence expresses in a context of use, moral or otherwise. He suggests that this vindicates a kind of analytic metaethical reductive naturalism, because it follows from a claim about ordinary linguistic meaning (making it analytic) that even in moral cases (making it metaethical), ‘good’ picks out a

⁴ “...it would be most implausible to give to the word ‘good’ in moral uses a sense quite unconnected with its sense or senses in other contexts. There cannot be two or more words ‘good’, mere homonyms of one another, like ‘bank’ (of a river) and ‘bank’ (a financial institution); for ‘good’ in English has counterparts in many other languages that have much the same range of moral and non-moral uses. We must hope to find either a single general meaning that the word has in both moral and non-moral contexts, or at least a core meaning of which its other senses are outgrowths.”

⁵ Ziff observes that terms like ‘heavy’ and ‘tall’ exhibit many of the same characteristics as the term ‘good’. Ziff asks his reader to consider the sentences ‘That is a good strawberry’ and ‘That is a good lemon’. To be a good strawberry is first and foremost to be sweet and to be a good lemon is to be sour. Now suppose these simple observations lead one to conclude that ‘good’ is ambiguous between sweet and sour uses. Such a person would then also have to claim that ‘heavy’ is similarly ambiguous, because it exhibits the very same kind of behavior in the sentences ‘That is a heavy car’ and ‘That is a heavy pencil’ as ‘good’ does in the first two sentences. The conditions under which a car is heavy are much different than the conditions under which a pencil is heavy, just like the conditions under which a strawberry is good are very different than the conditions under which a lemon is good. Since it is implausible to think so many of our terms are ambiguous in such ways, according to Ziff, it is bad idea to embrace any sort of ambiguity for the term ‘good’.

⁶ “The default hypothesis should be that ‘good’ has a single, unified semantics, however, especially since the moral/nonmoral distinction is both systematic across general normative vocabulary and robustly cross-linguistic.” Finlay (2014, p.19)

naturalistic relation defined by his End-Relational Theory (making it reductively naturalistic).

I am going to be arguing that the support that an end-relational semantics for ‘good’ can provide reductive naturalism has to be much less direct than Finlay acknowledges, and reductive naturalists in general might hope it to be. I’ll do this by highlighting several assumptions that a reductive naturalist would have to make about the underlying logical form of ‘good’ and its interaction with context, in order to secure a fully naturalistic reduction.

In order to fix ideas, I will start in § 1 by reviewing a canonical Moorean variety of metaethical nonnaturalism, in order to make clear which package of semantic and metaphysical views Finlay takes considerations of semantic parsimony to militate strongly against. In § 2 I’ll outline a few semantic assumptions that Finlay, and any like minded reductive naturalist, more generally, would have to adopt, in order to tighten the connection between an end-relational semantics for ‘good’ and reductive naturalism – assumptions that haven’t received the attention they merit in the context of this debate. In § 3 I will zero in on a surprising result – though one that has not gone completely unnoticed by Finlay, namely, that an end-relational semantics for ‘good’ is available to everyone, both naturalists and nonnaturalists alike. I’ll wrap things up in § 4 by addressing a few objections and replies.

1 Nonnaturalism, end-relationalism, and parsimony

At the highest level of abstraction, nonnaturalism in metaethics is the metaphysical view that there are nonnatural moral entities that are not identical with nor reducible to those entities that are the subjects of the natural sciences. According to nonnaturalists, in order to make complete sense of a world that appears to exhibit moral (or more broadly, normative) features, we have to embrace a crowded ontology of natural and nonnatural kinds, even in the full and familiar light of Occam’s Razor. Of course, however, no nonnaturalist countenances a metaphysically extravagant ontology like this, as friends and enemies of nonnaturalism alike openly characterize it, without thinking there is good reason for doing so. One recent but highly influential motivation for this picture of the

world has been the thought that moral entities are *just too different* from natural ones to be identical with or reducible to them.⁷

While nonnaturalists are quite explicit about their metaphysical commitments, openly nonnaturalistic accounts of the semantics for paradigmatic moral words such as ‘good’, ‘ought’, and ‘reason’ are relatively harder to find. Much of the little that nonnaturalists have had to say about the semantics for such words, however, suggests that many of them are hospitable to the idea that their meanings reflect the metaphysically diverse nature of the world. This is to say that many nonnaturalists seem to imply that words like ‘good’ have one sort of meaning in some contexts (e.g. an utterance of the sentence ‘Promise-keeping is good’ in the context of a debate about the value of promise-keeping) and a different sort of meaning in other contexts (e.g. an utterance of the sentence ‘This is a good hammer’ during a conversation about hanging a painting).

In other words, many nonnaturalists suggest that ‘good’ is fundamentally ambiguous, in that it contributes a nonnatural constituent to the propositions that sentences containing it express in *moral* contexts of use, and a natural constituent to the propositions that sentences containing it express in *nonmoral* contexts. Shafer-Landau (2003, p.66), for example, is forthcoming about this when he writes that “From the non-identity of moral and natural properties, it follows that *moral* terms cannot be naturalistically analyzed.” And Derek Parfit (2011, p.38-42, emphasis mine) strongly suggests this much in distinguishing between ‘good’ in the “reason-involving *sense*” (emphasis mine) and non-reason-involving sense.⁸

Stephen Finlay turns this line of thinking upside down, and in doing so is led to embrace the historical rival of nonnaturalism: metaphysical naturalism. Unlike the nonnaturalist who seems to think the semantics for ‘good’ reflects the diversity of the

⁷ See Terence Cuneo (2007, p. 219-223), David Enoch (2011, p.107-108), FitzPatrick (2008, p.176), Michael Huemer (2005, p.94-95), and Derek Parfit (2011, p.343-344) for arguments against various forms of reductive naturalism. Although each offers a distinct argument against reductive naturalism, what is common between each of them is that each takes the intuition that the moral or normative domain is just too different from the nonmoral/nonnormative domain seriously.

⁸ Michael Huemer (2005, p.209-210) is the only nonnaturalist of which I’m familiar who advocates that ‘good’ makes a nonnatural contribution to the meanings of both nonmoral and moral sentences.

metaphysics, Finlay argues that the metaphysics reflects the uniformity of the semantics. He models his view about the semantics for ‘good’ on the semantics for ordinary, uncontroversially nonmoral predicates such as ‘tall’.⁹ Like ‘good’, the predicate ‘tall’ appears in sentences that find use in a diverse range of contexts. So, for example, the term can appear in sentences in the context of professional basketball players, like the sentence ‘LeBron James is tall’, and in sentences in the context of philosophers, such as ‘Derek Parfit is tall’. But no one would ever be led by this observation to conclude that ‘tall’ predicates one kind of thing of LeBron James, such as the property of *being basketball-tall*, and another kind of thing of Derek Parfit, such as the property of *being philosophy-tall*.

Instead, many would conclude that ‘LeBron James is tall’ in a context of use expresses the relativized proposition that LeBron James is tall *for professional basketball players* and that ‘Derek Parfit is tall’ expresses the relativized proposition that Derek Parfit is tall *for philosophers*. On this familiar, and widely-accepted view about ‘tall’, the term is semantically uniform, but exhibits flexibility, via its sensitivity to sentential constituents, such as prepositional phrases (e.g. ‘for professional basketball players’), and sensitivity to salient contextual parameters (e.g. the average height of professional basketball players).

Finlay argues that what is uniform about the semantic contribution of ‘good’ is that it is an incomplete relational predicate with an underlying logical form consisting in a relation of ‘promotion’¹⁰, with parameters or argument-places for two kinds of proposition: propositions concerning objects (object-propositions) to which goodness is attributed, and propositions concerning ends (end-propositions) to which goodness is relativized.¹¹ It is the parameter for end-propositions which allows ‘good’ to exhibit the

⁹ The natural hypothesis is therefore that ‘good’ is a member of this family of incomplete predicates [e.g. ‘tall’, ‘old’, ‘cold’, ‘eager’], implying that the various kinds of goodness are relational properties.” Finlay (2014, p.21). His line of thought here is indebted to Ziff (1960) and Thomson (1997).

¹⁰ More clumsily but more accurately, it’s the relation of ‘increasing the probability of’. Roughly, the idea is that a proposition *p* promotes *e* just in case the instantiation of *p* raises the probability of *e* obtaining. Even though Finlay explicitly denies the ‘promotion’ gloss on ‘raising the probability of’, I will continue to use it, for ease of exposition.

¹¹ “A wide grammatical variety of ‘good’ sentences can be attributed the same logical form...by which the word expresses a relational property schema...predicating of a proposition *p* a property of standing in ...relation *R* to an end-proposition *e*.” Finlay (2014, p.38)

same kind of semantic flexibility as ‘tall’, by taking end-propositions as values for semantic completion, either from sentential constituents or that which is contextually salient.

To see how this view allows Finlay to claim that ‘good’ makes a single kind of contribution to the meanings of sentences in which the term appears, across a diversity of uses, first consider an end-relational analysis of the sentence ‘It’s good that the computer with Avid installed was used to edit the film’, in the nonmoral context of a discussion in a post-production house. Roughly, this sentence expresses the proposition that use of the computer with Avid installed promotes an end-proposition that is salient in this context, which in this particular case, is likely the end-proposition that the film is successfully edited. Thus, the goodness of using the computer with Avid installed is relativized to the promotion of the end of successfully editing a film.

The big payoff of this analysis is that it applies consistently to sentences in moral contexts of use, such as the sentence ‘It’s good that you’re no longer supporting factory farming’ in a discussion about the permissibility of eating meat between an ethically motivated vegan and nonvegan. On Finlay’s account, the meaning of this sentence involves relativization of the object-proposition that you’re no longer supporting factory farming to the end-proposition that is salient in the context. Given the description of this particular context, the end-proposition that animal suffering is reduced would likely be taken as argument. Thus, this sentence would express the proposition that you no longer supporting factory farming promotes the reduction of animal suffering.

Finlay’s End-Relational Theory of the semantics for ‘good’ entails that the term ultimately receives the very same semantic treatment in moral and nonmoral contexts of use, just as the term ‘tall’ makes the same contribution in contexts concerning both professional basketball and philosophy. This explains how ‘good’ is *not* ambiguous, despite appearances.

2 From semantics to metaphysics?

Finlay’s End-Relational Theory is a simpler (i.e. more unifying) view about the semantics for ‘good’ than the view that many nonnaturalists implicitly seem to hold about its

meaning. On Finlay's view, the term always makes a single kind of contribution to the propositions that sentences express across all contexts, whereas nonnaturalists like Moore, Shafer-Landau and Parfit seem to think that the term sometimes makes one kind of contribution (in nonmoral contexts) and another kind of contribution other times (in moral contexts). Since 'good' always makes a *single* kind of contribution to the propositions that sentences containing it express, and since this contribution is the relation of 'promotion', understood as the *naturalistic* relation of 'probability-raising', Finlay's end-relational semantics for 'good' would appear to support reductive naturalism. That is, it would appear as though the simplest theory that can accommodate all of the linguistic data concerning 'good', and hence best theory of it, supports a reductive naturalistic metaphysics.¹²

This would be a compelling case for reductive naturalism, if one could guarantee that *every* application of the end-relational semantic analysis of sentences containing the term 'good' in contexts of use yields thoroughly naturalistic analyses. The problem is that no one could guarantee this with an end-relational semantics for 'good', without making several auxiliary assumptions about the underlying logical form of the term and its interaction with context – assumptions that many nonnaturalists would hesitate to grant. So, the support that an end-relational semantics for 'good' may lend reductive naturalism has to be considerably less direct than reductive naturalists might hope it would be.

Over the course of the rest of this section we will see that there are at least three assumptions that need to be built into an end-relational semantics for 'good', in order for reductive naturalism to follow from it. The goal in outlining these is not to suggest that Finlay, nor any other relational semanticist with sympathies to naturalism, is unaware of them. In fact, as we'll see, Finlay explicitly addresses at least one such assumption. Rather, the goal is to raise awareness concerning just how wide the gap really is between an end-relational semantics for 'good' and a full-blown reductive naturalistic metaphysics

¹² "I shall employ an analytic method, seeking a metaphysical analysis of normative facts, properties, and relations by means of a conceptual analysis of the meanings of the normative words by which we refer to them. The *result* will be a form of analytic reductionism (or "analytic *naturalism*"), an explanatory reduction of normative properties (etc.) into complexes of non-normative properties by reductively defining normative words and concepts in entirely non-normative terms." (Finlay 2014, . p.4, emphasis mine)

– a gap that, to my mind, demands more attention than it has received thus far in the literature.

Assumption 1, Availability: No nonnatural entity plays any role in determining which end-propositions are *available* in a context to be taken as values by the contextual parameter of ‘good’.

To illustrate this assumption, imagine a conversation about the permissibility of flipping the trolley switch in a version of Philippa Foot’s (1967) classic case. Sheena doesn’t believe it would be good to flip the switch in order to save five people over one, whereas Ramona does. Now, suppose that at the time of Ramona’s utterance there were only two end-propositions available in this context to supplement the contextual parameter of ‘good’ for semantic completion: A proposition concerning the Utilitarian end of maximizing happiness, and another proposition concerning the Kantian end of doing only that which can be universally willed.¹³ Although the proposition that ‘Pulling the switch is good’ ultimately expresses in this context of use will be entirely naturalistic, because the only two end-propositions available for supplementing the contextual parameter of ‘good’ are naturalistic, the thing to notice is that an end-relational analysis of this sentence won’t be entirely naturalistic, if some nonnatural entity is responsible for making these two end-propositions available in the first place.

Of course, however, nonnaturalists might not be willing to accept Availability, and so it is incumbent upon end-relational semantic theorists with naturalistic ambitions to either show, impossibly, that no nonnatural entity plays a role in fixing which end-propositions are available in contexts, or show that a naturalistic explanation of the phenomenon is the best available. Finlay opts for the latter route. He argues for the truth of general pragmatic principles governing conversational exchange, such that uses of

¹³ I’m here assuming that this Kantian end is naturalistic.

unrelativized sentences containing ‘good’ are implicitly relativized to ends saliently desired by either the speaker or audience, by default.¹⁴

Interestingly, however, some philosophers¹⁵ have recently argued for precisely the exact opposite conclusion, namely, that unrelativized sentences containing terms like ‘good’ receive a default, “moral” interpretation – one that is distinct from “bouletic” (desire- or preference-relative) readings. The point in mentioning this controversy is not to settle it. Rather, it is to center our attention on an important claim – one that has mostly been addressed indirectly, by theorists tackling other issues in the literature. The point is to suggest that the nature of contextual parameter completion is very much up for grabs between naturalists and nonnaturalists and must not be taken for granted by those who would try and derive naturalism from an end-relational semantics.

Assumption 2, Saliency: No nonnatural entity plays any role in determining which available end-proposition is *salient* in a context, such that it is taken as a value by the contextual parameter of ‘good’.

Saliency is a close relative to Availability. It states that there aren’t any nonnatural entities that are responsible not for determining which end-propositions are *available* for completing the contextual parameter of ‘good’, but which end-propositions, among those available, are *salient* in a context. Return now to the exchange between Ramona and Sheena, in order to illustrate this thought. Suppose again that Ramona has uttered the sentence ‘Pulling the switch is good’ and that the only end-propositions available for semantically completing ‘good’ in this context is the naturalistic Utilitarian end of maximizing happiness and the naturalistic Kantian end of doing that which can be universally willed. We can even further suppose, pace Finlay, that the explanation for why these are the only two end-propositions available is entirely pragmatic and hence naturalistic. This is to say, then, that it is because Ramona desires these ends that they are available in this context.

¹⁴ For all of the details regarding Finlay’s pragmatics, see his (2014, Chapter 5).

¹⁵ Janice Dowell (Forthcoming, p.26)

Given the setup of this scenario, it certainly looks like an end-relational analysis of the meaning of ‘Pulling the switch is good’ in this context of use supports reductive naturalism. After all, each constituent of the proposition expressed by the sentence in this context of use is a natural one (e.g. the object-proposition concerning pulling the switch, the relation of promotion, and the end-proposition concerning either the Utilitarian or Kantian ends). But such appearances are misleading. In order for this to be the case it has to be assumed that no nonnatural entity plays a role in making either the Utilitarian or Kantian end *salient* in this context. The problem though is that an end-relational analysis in no way guarantees this, and no nonnaturalist has to grant it. So, the naturalist faces a pair of familiar options: either show, impossibly, that no nonnatural entity can influence which end-propositions are salient, or show that a naturalistic explanation of the phenomenon is the best one available.

Finlay opts for the latter route, by again arguing that the truth of pragmatic norms on conversational exchange explain how it is that ends become salient in contexts. As Finlay notes, “Ends can of course be salient in a conversation in many different ways,” including his preferred way, namely, by being the shared object of desire between conversational participants.¹⁶ Now, it may very well be the case that the content of our desires can determine that which is salient in a context in the way Finlay describes. The thing to notice about this pragmatic explanation, however, is that it in no way rules out nonnatural entities from playing a role in this explanation, at least indirectly, if we can have desires, beliefs, and the like with nonnatural contents. Since conversational participants could plausibly share a desire to promote a nonnatural end, citing the content of our attitudes as that which makes ends salient in contexts isn’t enough to establish Saliency, for the simple reason that the nonnatural contents of our beliefs may help determine which ends are salient in context.¹⁷

Assumption 3, Occupation: No end-proposition
containing nonnatural
constituents can occupy the

¹⁶ Finlay 2014, p.136

¹⁷ Nor would it be enough to establish Availability.

contextual parameter of
'good'.

The final assumption that a proponent of naturalism would have to make in order to secure a fully naturalistic end-relational semantic treatment for 'good' is that the term has a parameter in its logical form taking not contextually salient end-propositions as values, but contextually salient end-propositions containing only *naturalistic* constituents, or, more simply, naturalistic end-propositions. While it might be true that the *term* 'good' always contributes a single and natural relation to the propositions that sentences express, it might still be the case that *sentences* containing the term express propositions with at least one nonnatural constituent, because a prepositional phrase or a context could very well supplement 'good' with an end-proposition containing nonnatural constituents, or, again more simply, *nonnatural* end-propositions, for completion.

To see how 'good' could be completed in both of these ways, such that a sentence containing it would express a nonnatural proposition in a context of use, first consider a sentence with an explicitly relativizing prepositional phrase (e.g. 'It's good for being just that we abolish the institution of private property'). Suppose that 'good' contributes the relation of promotion to the proposition that this sentence expresses, just as Finlay would claim it does. In this case, the relation of promotion would hold between the object-proposition that we abolish the institution of private property and an end-proposition concerning whatever it is that 'for being just' picks out. But, importantly, it could be that the entity referred to or described by this phrase is nonnatural, which would entail that the sentence as a whole expresses a proposition with at least one nonnatural constituent.

The same goes for sentences containing 'good' that lack relativizing prepositional phrases in contexts of use. Recall the conversation between Ramona and Sheena. Since the sentence 'Pulling the switch is good' is not explicitly relativized it requires that the context supply an end-proposition to complete the predicate 'good', in order for the sentence to express a complete proposition. If some nonnatural proposition were salient in this context, then 'good' would take it for semantic completion, such that the sentence would express a proposition with at least one nonnatural constituent.

These examples suggest that sentences containing the term ‘good’ can express propositions with nonnatural constituents. Of course, the success of these examples is contingent upon both the availability of nonnatural propositions and the saliency of them. So, the examples are not offered as actual cases of sentences containing the term ‘good’ expressing propositions with nonnatural constituents in their contexts of use. Instead, they are being offered to show what kind of cases would need to be ruled out in order for an end-relational semantics for ‘good’ to support reductive naturalism.

But if the example sentences above (e.g. ‘It’s good for being just that we Abolish the institution of private property’ and ‘Pulling the switch is good’) can express propositions that contain nonnatural constituents, it should be acknowledged that the *source* of such constituents would not be the term ‘good’ itself. Rather, the sources of such nonnatural constituents would be the phrase ‘for being just’ and context, respectively. This is important to note because it amounts to something of a partial victory for philosophers like Finlay who are sympathetic to reductive naturalism, since it might be thought that one of the goals of an analytic reductive naturalistic project in metaethics is to show that moral terms like ‘good’ are not, after all, in any way special, and that, accordingly, they only make naturalistic contributions to the propositions that sentences containing them express.

This victory is small progress, however. This is because the primary objective of any analytic reductive naturalistic project in metaethics is to secure reductive metaphysical analyses, by providing reductive analyses of ordinary linguistic meanings of terms like ‘good’ in their contexts of use. But *sentences* are the primary vehicles for expressing the meanings of moral terms in their contexts of use. So, in order to establish reductive naturalism via an end-relational semantics, one of the very first things that has to be assumed is that sentences containing terms such as ‘good’ cannot take propositions with irreducibly nonnatural entities as values for semantic completion. Of course, the problem with this is that no nonnaturalist would ever come on board with this assumption.¹⁸

¹⁸ Finlay seems willing to grant this result, since he writes that “The end-relational theory is conservative with regard to the content of first-order normative or ethical theory...It is therefore compatible with major moral theories including utilitarianism, Kantianism, contractualism, and Divine Command.” Finlay (2014, p.254).

3 End-relational semantics for nonnaturalists

Though it might be thought that reductive naturalism scores highest on the plausibility charts, because the end-relational semantics that it can enlist yields the simplest semantics for ‘good’, I have tried to suggest that this isn’t quite right.¹⁹ In fact, an end-relational semantics for ‘good’ leaves the metaphysical nature of goodness relatively open, such that a non-Moorean, *sophisticated* nonnaturalist about the semantics of moral terms can agree with Finlay that ‘good’ always contributes a natural relation to the propositions that sentences containing it express, but also maintain that in some cases, the term may receive a nonnatural proposition as an argument for semantic completion from either a prepositional phrase or a context.

Finlay’s end-relational semantics is an account of the underlying logical form of the term ‘good’ – one that merely states that ‘good’ is an incomplete predicate with an argument place that exhibits sentential and contextual sensitivity. The account does not, however, completely dictate the metaphysical nature of the propositions that this argument-place may take as argument, contra Occupation. This is unfortunate for naturalists like Finlay, since one positive upshot of this observation is that he has unwittingly breathed life into the nonnaturalist research program, by providing those sympathetic with it new semantic resources. No longer does a nonnaturalist have to concede that the term ‘good’ is fundamentally ambiguous between moral and nonmoral contexts – a concession that many philosophers have long regarded as an embarrassment for the view.

Moreover, a sophisticated nonnaturalism can provide resources for bolstering recent arguments against their opponents. For example, in Part VI of *On What Matters*, Derek Parfit adopts a unique argumentative strategy for defending metaethical nonnaturalism. It might be thought that the deep and intractable disagreements between philosophers working in ethics is evidence against the possibility of moral progress. And if this is right, then one way to help establish the possibility of moral progress is to

¹⁹ See Enoch (2011) for this conception of theory score-keeping.

dissolve such disagreements. On this understanding, then, Parfit is trying to show that in at least his case, the disagreement between him and many other philosophers is merely apparent, because those philosophers who are thought to disagree with him don't share the right concept of morality (or normativity, more generally) to do so. Thus, for example, the expressivist Alan Gibbard, the Humean Bernard Williams, and the error theorist J.L. Mackie are not in any disagreement with Parfit, because Parfit's concept of normativity is essentially reason-involving, whereas Gibbard, Williams, and Mackie understand normativity in a broadly motivational, non-reasoning-involving sense.²⁰

Unfortunately for Parfit, even if it were true that Gibbard, Williams, and Mackie didn't share his concept of normativity²¹, it's not at all obvious *how* his goal of dissolving apparent disagreements between him and his rivals is well served by an appeal to differences in understandings about the nature of normativity. It might be better served, however, in at least some instances of apparent disagreement, with the help of a sophisticated nonnaturalism about the semantics moral discourse.

Take the apparent disagreement about the nature of the rationality of desires between Parfit – who is a nonreductive nonnaturalist – and Richard Brandt – who is a reductive naturalist, as an example.²² Brandt took himself to have held the view that a desire is rational just in case an agent would still have it after reflection on all the relevant facts, or just in case the desire would survive what he calls *cognitive psychotherapy*. And correspondingly, for our purposes, he might be taken to have also held the view that the meaning of the word 'rational' is something like 'fully informed'.²³ Parfit, on the other

²⁰ Here is a characteristic quote from Parfit on Williams on the nature of reasons: "When I have earlier claimed that Williams did not understand this external concept of a reason, some people have urged me to be more charitable. These people suggest that, like Scanlon, I should assume that Williams had this concept, and was merely making different claims about which facts give us reasons. But this assumption would, I believe, be less charitable. If Williams did understand the external normative sense, why does he so often call this sense mysterious and obscure?...Williams rejects this idea, I believe, because this kind of goodness is reason-involving, and Williams thinks of reasons, not as facts that count in favour of our having some desire or acting in some way, but as facts that might motivate us." (Parfit 2011, p.434-435)

²¹ See Phillips (forthcoming) for an argument for why Mackie clearly did share Parfit's concept of normativity.

²² Parfit (2011, p.368-376)

²³ Officially, Brandt takes himself to be offering a "reforming definition" of the term. But the letter of Brandt's view is not important for our illustrative purposes in this section.

hand, takes himself as holding that a desire is rational just in case there are sufficient irreducibly nonnatural reasons for having it and that the “ordinary sense” of the word ‘rational’ is reason-involving or irreducibly nonnaturalistic.

It could be that Parfit and Brandt don’t genuinely disagree about the nature of the rationality of desires, on the truth of an end-relational semantics for the evaluative term ‘rational’²⁴, if it turns out that Parfit and Brandt have been using the term in *different contexts*. This could be the case, if the contextual parameter of ‘rational’ were to take a naturalistic end-proposition as its value out of the mouth of Brandt, but a nonnaturalistic end-proposition as its value out of the mouth of Parfit,²⁵ and, importantly, *expressing propositions with nonnatural constituents is just what it takes to be in a moral context*.

On this way of framing their dispute, Brandt and Parfit don’t disagree *about morality*, because Brandt never made a *moral* claim that conflicted with Parfit’s claims about this domain.²⁶ And this is thus how it could be that the disagreement between Parfit and Brandt about the nature of the rationality of desires in *moral* contexts is merely apparent. The truth of an end-relational semantics for moral discourse allows us to better understand the surprising thought that Brandt may never have even expressed anything in opposition with Parfit’s views about the nature of normativity.

4 Responses and replies

There are a variety of responses available to reductive naturalists of an end-relational semantic persuasion for tightening the connection between an end-relational semantics and reductive naturalism. Consider again the sentence ‘It’s good for being just that we

²⁴ One might deny that ‘rational’ is an evaluative term, and one might also worry about exactly how this term could be end-relational. If so, then insert a different evaluative term into the example. The point being made here is a general one, and does not depend on the details of the analysis of ‘rational’.

²⁵ Why think ‘rational’ is takes a naturalistic end-proposition as value for Brandt, but not Parfit? Well, if the content of our attitudes could determine which end-propositions are salient in contexts, as Finlay thinks, then the term ‘rational’ out of Brandt’s but not Parfit’s mouth might have always taken naturalistic propositions for completion, because of his naturalistic views about these matters. For another conception of the content of context on which the content of attitudes can count as determining it, see Stalnaker (1978).

²⁶ Notice just how well this diagnosis fits with the spirit of Parfit’s assessment of his apparent disagreement with Bernard Williams: “Though Williams and I used the same normative words, we used them in different senses. We were not really, as we assumed, disagreeing.” Parfit (2011, p.448)

abolishing private property'. While a sophisticated nonnaturalist might say that this sentence expresses the proposition *that our abolishing private property promotes being just*, where 'just' is understood as picking something nonnatural out, a naturalist might claim that the term 'just' also demands an end-relational semantic treatment, such that it too contributes only a natural constituent to the propositions that sentences containing it express, whatever that constituent might be. This could potentially rule out the availability of nonnatural analyses by showing that a natural, end-relational semantic treatment for the meanings of *all* evaluative terms (e.g. 'right', 'rational', etc.) is the best available for each one.²⁷

There are two things to say to this response. The first is that even if it were true that every evaluative term contributes a naturalistic entity to the propositions that sentences containing them express in contexts of use, a naturalist would still have to establish the truth of Availability, Saliency, and Occupation, in order to show that nonnaturalist analyses are unavailable. The second thing to say is that if every evaluative term can be given an end-relational semantic treatment, this would still not show that there are no nonnatural entities in moral contexts. Rather, all this would show that there are no *privileged* words for expressing them.

Moreover, fulfilling the impossibly tall order of showing that moral contexts don't contain nonnatural propositions would not rehabilitate the *semantic* argument for reductive naturalism from above. Recall, the semantic argument for reductive naturalism starts with a claim about the simplicity of an end-relational semantics of an evaluative term (in this case 'good'), and concludes with a claim about support for reductive naturalism. But if the response above were correct, then any considerations of semantic simplicity would be otiose. This is because if it could be shown that no context contains even a single nonnatural proposition, then the case for reductive naturalism will already have been made, well before we ever start to worry about the semantics for evaluative terms. The final word will have come via a *metaphysical* argument for the truth reductive naturalism, not from a semantic one.

²⁷ Finlay anticipates this kind of retreat to other evaluative terms (e.g. 'important', 'matters') on behalf of nonnaturalists, and argues that all of them are indeed plausibly analyzed end-relationally and naturalistically. (Finlay 2014, p.252-253)

But it matters that end-relational semanticists with reductive naturalistic sympathies succeed via an argument from semantics instead of an argument from metaphysics, because reductive naturalism was sold to us in part on the basis of its semantic credentials. But if it isn't true that reductive naturalism has any claim over nonreductive nonnaturalism to such credentials, then we ought not to prefer reductive naturalism on semantic grounds. The debate between reductive naturalism and nonreductive nonnaturalism will have ultimately been decided, as some have recently argued it should have been decided, in the metaphysical and not the semantic arena.²⁸

Reductive naturalists might at this point try shifting dialectical direction, by reminding us at this point that naturalism is a simpler thesis than metaphysical nonnaturalism, because it entails the existence of fewer ontological kinds. Accordingly, they might then say, metaphysical naturalism ought to be our default hypothesis about the ontological nature of our world. But since much of the surface grammar of our language is such that this simple metaphysical picture of our universe is obscured, naturalists might admit, we are mistakenly led to suspect that there may be more kinds in our ontology. By showing, however, that all sentences can be uniformly analyzed in terms of naturalistic probabilistic relations between ends, and explaining away the distracting, seemingly nonnaturalist commitments of our language, naturalists can claim that they have done all the work that needs to be done to make the world safe again for metaphysical naturalism. They could concede that metaphysical naturalism may not follow from the simplicity of an end-relational semantics for 'good', but point out that metaphysical nonnaturalism doesn't follow, either. And if the semantics doesn't point us away from naturalism, we ought not give up this default position.

Interestingly, this response has brought us dialectically full circle. As we saw at the beginning of this paper, the very first thing nonnaturalists *concede* is the metaphysical extravagance of their view. To remind nonnaturalists that our default metaphysical picture of the world should be the simplest one (e.g. naturalism) doesn't advance the dialectic with them. And then to point out that the correct semantics for moral discourse doesn't give us any reason to reject naturalism is likely to raise eyebrows, because been a

²⁸ See McKeown-Green, Pettigrove, and Webster (2012)

nonnaturalists have not been primarily led to deny naturalism on the basis of its semantic credentials, or lack thereof.²⁹ While it may have been common to find nonnaturalists embracing nonnaturalism from semantic considerations during the introduction of Moore's Open Question Argument at the turn of the 20th century, many contemporary nonnaturalists tend to cite as the basis for their view the powerful *metaphysical* intuition that the moral is *just too different* than the natural to be identical with or reduce to it.

Thus, nonnaturalists are unlikely to find this response persuasive, since it doesn't shift the dialectical burden back to nonnaturalists in any way. But if nonnaturalists in the Moorean tradition do stop and think about the implausibility of their pre-theoretic semantic commitments, then I hope to have shown that they don't, after all, have anything to worry about, because they can avail themselves of a sophisticated nonnaturalism about the semantics of moral discourse.

Conclusion

One of my primary aims in this paper was to show that an end-relational semantics for 'good' doesn't support reductive naturalism to quite the extent that it may seem to support it. An end-relational semantics cannot guarantee that sentences containing the term 'good' yield thoroughly naturalistic analyses, because nonnatural entities might play roles in determining the availability (contra Availability) and saliency (contra Saliency) of end-propositions in context. Moreover, there is nothing in the underlying logical form of the term that precludes it from taking a nonnatural end-proposition as a value for contextual completion (contra Occupation). And so there is nothing about an end-relational analysis that guarantees that sentences containing 'good' won't express propositions with at least one nonnatural constituent.

One exciting result along the way in this paper was the development of a sophisticated nonnaturalism about the semantics of moral discourse. It is worth stressing

²⁹ One might complain that this claim is overstated, since there are contemporary nonnaturalists who do indeed take semantic considerations quite seriously, including FitzPatrick (2008, p.179). But the winds are noticeably shifting, in at least the writings of some of the most high-profile nonnaturalists. It is telling that Parfit does not enlist the Open Question Argument in the battery of objections that he raises against naturalists in Volume 2 of *On What Matters*. And in a representative passage, Enoch writes of it that the "...Open Question Argument fails miserably when understood as trying to make [a] metaphysical point..." Enoch (2011, Ch.5 fn.1)

the significance of this development, not because I am sure that a nonnaturalist end-relational semantics for moral discourse is right – something that we could only learn from its details, but because the possibility of this view suggests that the score between naturalists and nonnaturalists might be *even* with respect to semantics, appearances notwithstanding. So, if there is one thing that I hope can be taken away from this paper, it's that we're going to have to appeal to the virtues and vices of naturalism and nonnaturalism along some *other*, non-semantic dimension, in order to settle which family of views is superior.

References

- Cuneo, Terence (2007). *The Normative Web: An Argument for Moral Realism*. Oxford University Press.
- Dowell, J. L. (2012). Contextualist Solutions to Three Puzzles about Practical Conditionals. In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics*, volume 7. Oxford.
- Enoch, David (2011). *Taking Morality Seriously: A Defense of Robust Realism*. Oxford University Press.
- Finlay, Stephen. (2014) *Confusion of Tongues: A Theory of Normativity*. Oxford University Press
- Fitzpatrick, William (2008). Robust ethical realism, non-naturalism, and normativity. *Oxford Studies in Metaethics* 3:159-205.
- Foot, Philippa (1967). The problem of abortion and the doctrine of double effect. *Oxford Review* 5:5-15.
- Huemer, Michael (2005). *Ethical Intuitionism*. Palgrave Macmillan.
- Mackie, J. L. (1977). *Ethics: Inventing Right and Wrong*. Penguin.
- McKeown-Green, Jonathan ; Pettigrove, Glen & Webster, Aness (forthcoming). *Conjuring Ethics from Words*. *Noûs*.
- Moore, G. E. (1903). *Principia Ethica*. Dover Publications.
- Parfit, Derek (2011). *On What Matters*: Two-volume set. OUP Oxford.
- Phillips, David (forthcoming). Sympathy for the Error Theorist: Parfit and Mackie. *Ethical Theory and Moral Practice* :1-8.
- Schroeder, Mark (forthcoming). What matters about metaethics? In Peter Singer (ed.), *Does Anything Really Matter? Responses to Parfit*.
- Shafer-Landau, Russ (2003). *Moral Realism: A Defence*. Oxford University Press.
- Stalnaker, Robert (1973). Presuppositions. *Journal of Philosophical Logic* 2 (4): 447 - 457.
- Thomson, Judith Jarvis (1997). The right and the good. *Journal of Philosophy* 94 (6):273-298.
- Ziff, Paul (1960). *Semantic Analysis*. Ithaca, N.Y., Cornell University Press.