

Streumer on non-cognitivism and reductivism about normative judgement

Daan Evers, University of Groningen

Penultimate draft, forthcoming in the Journal of Moral Philosophy

1. Introduction

In his challenging book, Bart Streumer defends an error theory about all normative judgements. A successful defence requires defeating all other metaethical views, including non-cognitivism and reductive realism. In this paper, I explore some avenues of response. I will focus on non-cognitivism, but briefly reflect on the implications for reductivism at the end.

2. The Symmetry Objection to Non-Cognitivism

Streumer thinks that we take the following claim to be true of normative judgements:

- (A) When two people make conflicting normative judgements, at most one of these judgements is correct.

So, according to Streumer, we take conflicts between normative judgements to be *asymmetrical*: we do not think that conflicting normative judgements can be equally correct. Streumer believes that (A) is a central feature of normative discourse as it currently is: any account incompatible with its truth is not an account of normative discourse in the first place.

Streumer argues that non-cognitivists cannot explain why we take (A) to be true of normative discourse in the *way* we take it to be. Non-cognitivists believe that normative judgements, like

Euthanasia is wrong

are non-cognitive attitudes, such as an attitude of disapproval of euthanasia. If Streumer is right, then non-cognitivism is a revisionary theory of normative judgement.

Streumer does not tell us much about the content of (A), but I think this is deliberate. He does not want to *start* with an interpretation that guarantees that non-cognitivists cannot account for it. Rather, he wants to consider various things they might say about it, in order to see whether those things intuitively match whatever we implicitly take (A) to mean, or what we take to explain our commitment to it. If nothing is a good enough match, then that is evidence that our commitment to (A) is incompatible with non-cognitivism.

We do learn that the notion (or standard) of correctness is more demanding than whatever is appropriate for the assessment of conflicting judgements within a single perspective. Even correct use of 'tasty' is governed by a principle according to which the same thing cannot be simultaneously tasty and disgusting, but this norm is perspective-internal: *I* cannot appropriately judge the same thing to be both tasty and disgusting, and so at most one of my judgements can be correct. But Streumer formulates (A) in terms of *two* people making conflicting normative judgements. So when it comes to normative language, (A) extends beyond the perspective of an individual. When I judge that euthanasia is permissible, I consider someone else's conflicting judgement as mistaken; if I judge that vegemite is tasty, I don't consider someone else's conflicting judgement as mistaken. This means that what we are looking for is an explanation of the *contrast* between our reactions to different kinds of (apparently) conflicting judgements.¹

There are certainly cases where a contrast exists, although how deep it is and how far it extends is not clear. I suspect it is a matter of degree. Some empirical research suggests that whether people take conflicting moral judgements to be

¹ One possible explanation of the sense that (apparently) conflicting judgements of taste can both be correct is that such judgements are representations of the responses of individuals or groups. If so, then appeals to judgements of taste in order to pump intuitions against non-cognitivist accounts of normative judgements lose some of their force. If these respective classes of judgements are of a very different kind, then intuitions to the effect that normative judgements aren't like judgements of taste may not tell us much about whether normative judgements are non-cognitive states.

mistaken depends on background assumptions about the parties to the disagreement (Sarkissian et al. 2011). Other research suggests that people consider some but not all moral questions to have uniquely correct answers (Goodwin & Darley 2008). This is important, as Streumer's case against non-cognitivism depends on the assumption that we take (A) to be true of all normative language.

But let us suppose that we do accept (A) with respect to all normative judgements. Why does Streumer think that non-cognitivists cannot account for it? It seems that Streumer does two things: first, he considers a possible non-cognitivist interpretation of the *content* of (A), and argues that it fails. Second, he considers non-cognitivist explanations of our *acceptance* of (A), and argues that they fail. These explanations of our endorsement of (A) do not clearly provide an interpretation of its content. I will focus on some of the more promising suggestions.

First, Streumer considers a non-cognitivist account of the *content* of (A), along the lines of Blackburn's treatment of normative conditionals. According to this account, (A) expresses the following attitude:

(A*) Disapproval of two people approving and disapproving of a single thing.

Streumer doubts that this is really what (A) amounts to. For suppose that Bob likes peanut butter and Kate dislikes it. Suppose further that they belong to a community in which everyone has the following attitude:

(L*) A dislike of two people liking and disliking a single thing.

Finally, suppose that the attitude described by (L*) can be expressed by saying that

(L) When two people have conflicting likes or dislikes, at most one of these likes or dislikes is correct.

Streumer thinks that

'it is hard not to feel that *something* about the way we endorse (A) is missing from the way this community endorses (L)' (p. 74).

I take it that what Streumer means is that it is hard not to feel that our commitment to (A) consists in something other than an attitude of dislike towards two people approving and disapproving of a single thing.

Of course, if this is what he means, then a critic of the argument might say that there *is* a difference in the way Bob and Kate's community endorses (L) and the way we endorse (A). For the alternative community's endorsement of (L) consists in an attitude of *dislike*, whereas our community's endorsement of (A) consists of *disapproval*.²

But perhaps we feel that there is also some difference between endorsing (A) and disapproving of two people approving and disapproving of a single thing. We might test this as follows (as Streumer comes close to doing near the end of his discussion): suppose that Bob and Kate's community consists of people who disapprove of people liking and disliking a single thing, and suppose that their disapproval can be expressed by means of (L). Streumer would probably say that it is hard not to feel that there is some difference between our endorsement of (A) and this community's endorsement of (L).

How to assess this verdict? What would be evidence for Streumer's view is the possibility of disapproving of an action while lacking the sense that those who approve of the action are mistaken. I find this hard to determine introspectively, so long as we keep in mind the distinction between dislike and disapproval. We can of course distance ourselves from our states of disapproval and doubt that they are uniquely

² A reviewer for this journal finds it unobvious whether there is a distinction between liking and approval. It seems to me that these are not the same. Liking something is connected with enjoyment, whereas approval is connected with being in favour of something. As a result, one may like something without approving of it and vice versa (perhaps one likes licking one's plate in public, without approving of it). Of course, one's approval of a thing is sometimes based on the fact *that* one likes it. Perhaps one approves of blue paint on the wall *because* one likes it. (Note, though, that the object of liking and approval may not strictly be the same in this case: what one likes or enjoys is the blue paint on the wall, whereas what one approves of is its remaining there, or being applied to it.)

correct, but we can do this too with our normative judgements. So I think the argument is inconclusive.

Streumer also considers the possibility that people have false beliefs about (some aspects of) normative language. For instance, we may have false beliefs about the way we endorse (A), which is really no different from the way the imaginary community endorses (L). However, Streumer rightly notes that this requires explanation. Why do we feel that there is a difference if there is in fact no difference? According to Streumer, the most promising explanation for the non-cognitivist is to say that

- (1) The way we endorse (A) is different from the way Bob and Kate's community endorses (L)

itself expresses an attitude of approval or disapproval, such as:

- (1*) Disapproval of thinking that (A) expresses the attitude described by (A*).

Although Streumer is sceptical of this idea, he does not immediately dismiss it. I have no faith in it myself. (1) does not appear to be a normative claim, but the idea under consideration entails that it is. One cannot simply announce that words not known to play a normative role in other contexts suddenly take on a normative aspect when they are combined to express (1). Such a view is *ad hoc* without some principled defence *other* than that it saves non-cognitivism.

Nevertheless, there is a better approach involving false beliefs. The non-cognitivist might say that (A) is *itself* a false belief about the status of normative judgements. Of course, if Streumer is right that (A) characterizes the *nature* of normative judgements, then non-cognitivists cannot accept that (A) is a false metaphysical belief. But a plausible non-cognitivist explanation of our acceptance of

(A) might itself undermine the claim that it reveals the nature of normative judgement.³

The idea that (A) is a false belief in the status of normative judgements is compatible with another suggestion by Streumer on behalf of non-cognitivism. According to this suggestion, our commitment to (A) is explained by the function of normative language, which is to coordinate our attitudes:

'[Non-cognitivists] could say that we endorse (A) because this creates pressure on ourselves and others to have similar attitudes of approval and disapproval, which helps normative discourse to fulfil its function.' (p. 76)⁴

As before, Streumer objects by drawing an analogy. He imagines an alternative community that needs to coordinate its members' likes and dislikes. As a result, the function of their discourse about such likes and dislikes is to coordinate their attitudes. Everyone in the community therefore endorses

(L) When two people have conflicting likes or dislikes, at most one of these likes or dislikes is correct.

According to Streumer,

'This imagined community [...] endorses (L) in exactly the same way in which, according to this explanation, we endorse (A). But once again, it is hard not to

³ This can be so provided the explanation is compatible with non-cognitivism and there is independent evidence for this family of views.

⁴ Earlier, I indicated that the story about the function of normative language does not by itself provide an interpretation of the content of (A). Streumer seems to agree, as he also suggests that non-cognitivists can combine the claim about the function of normative language with the idea that this function explains why we have the attitude described by (A*) above (i.e. why we disapprove of two people approving and disapproving of a single thing). However, once we separate non-cognitivist readings of the content of (A) from explanations of our acceptance of it, the assumption that (A) is itself a state of disapproval is optional.

feel that *something* about the way we endorse (A) is missing from the way this community endorses (L).' (pp. 76-77)

I think that this is also inconclusive. First, our conscious commitment to the subjectivity of likes and dislikes could wrongly suggest that our commitment to (A) cannot be explained in the same way as the alternative community's commitment to (L). Second, it may seem as if members of Streumer's community consciously decide to coordinate attitudes by adopting (L). However, if we accept (A) because of the function of normative language, then we accept it as a result of language learning in early childhood, and there will be no conscious strategy. Third, we may have difficulty believing that our commitment to (A) is causally explained by our understanding of the function of normative language *precisely because we believe (A)*. If we do believe (A) - and (A) is a mistaken belief about the status of normative judgements - then this belief suggests to us that the function of normative judgement cannot fully explain our commitment to (A).

I don't think it is implausible that we have some false beliefs about the status of normative judgements. I also think that such beliefs might be explained by the function of normative language, or result from non-cognitive states or processes. This is of course an empirical hypothesis, but I am not the first to propose it. Immanuel Kant held that judgements of beauty come with an impression of universal validity, despite the fact that he took them to be grounded in a pleasurable feeling. Such impressions need not be due merely to the state that constitutes the judgement. For instance, the mental states expressed by *moral* sentences are typically embedded in a larger syndrome of cognitive, conative and affective states that might produce beliefs in objectivity. A moral judgement is typically backed by dispositions to emotional arousal and comes with higher-order dispositions to promote the judgement in others (who are themselves disposed to react very differently to our normative judgements than to our judgements of taste). To the extent that such mechanisms differ for different normative judgements, (A) may be less generally accepted or obvious to us.

Thorough-going quasi-realists may not accept that we have false beliefs in the objectivity of (some) normative judgements. Quasi-realism is the attempt to understand apparently objectivist trappings of normative language and thought in a

way compatible with non-cognitivism. One of its proponents, Simon Blackburn, would almost certainly suggest that the conviction that a moral proposition is objectively true is itself a first-order normative commitment, probably of the same kind as that expressed by counterfactuals like the following:

(2) Had we all approved of torture, then it would still be wrong

Although (2) appears to state a metaphysical claim about the independence of moral truths from our attitudes, Blackburn denies that there is such a reading:

'When philosophers ask whether moral truths are mind-dependent, the question appears to be asked at the meta-level. [...] But there is no such meta-level question. According to me, 'moral truths are mind-dependent' can *only* summarize a list like 'If there were no people (or people with different attitudes) then X...', where the dots are filled in by some moral claim about X. One can then only assess things on this list by contemplating the nearest possible world in which there are no people or people with different attitudes but X occurs. And then one gives a moral verdict on that situation.'
(1998, p. 311).

So Blackburn would probably say that there is *no* way to understand a belief in the objectivity of moral facts as anything other than a first-order normative conviction. If Blackburn is right about this, then non-cognitive states (or syndromes of such states) cannot give rise to (mistaken) beliefs in metaphysical standing, and I cannot appeal to this possibility in explaining our commitment to (A) on behalf of non-cognitivism.

However, even if Blackburn is right that there is a normative reading of claims like (2) (or claims more directly framed in terms of objectivity), then that would not show that we cannot have false metaphysical beliefs about morality. In order to rule that out, Blackburn must show not only that there *is* a normative reading of apparently metaphysical claims, but also that there is *no* non-normative reading. This is unlikely for at least two reasons. First, if we can understand (if only to reject) the idea of a relation of reference between objects in the world and our thoughts about it,

then we can surely construct a metaphysical notion of objectivity. Second, it takes a lot of argument to persuade both students and philosophers that (2) can only ever be a normative judgement. This is defeasible evidence that people can entertain metaphysical readings of such dependence claims. And if so, then they may have (perhaps inchoate) false beliefs about some normative judgements. I think that Streumer has not shown that such beliefs cannot result from non-cognitive processes or an implicit understanding of the function of normative language.

Whether the ascription of false beliefs about normative language is a cost depends on the extent to which it is plausible that non-cognitive processes can give rise to them.⁵ The hypothesis that they can has some, albeit defeasible and often indirect, empirical support. For instance, Nichols & Folds-Bennett (2003) found that a majority of young children think of certain actions as disgusting 'for real' as opposed to disgusting 'for some' (the examples used involve spitting in water before drinking it, and wiping one's nose on hand and sleeve). This is defeasible evidence that disgust can trigger beliefs in objectivity.⁶ Goodwin & Darley (2010) found that wrongs related to harm and injustice are considered more objective than certain other kinds of wrong. If reflection on such wrongs is correlated with comparatively strong emotional arousal, then that would lend indirect support for the hypothesis that non-cognitive processes can result in beliefs in objectivity. Some, albeit limited, evidence that the antecedent holds is provided by Greene et al. (2001) and Greene et al. (2004). These studies show that reflection on scenario's involving "up close and personal" harm is correlated with increased activity in areas of the brain associated with emotional arousal.⁷ Finally, Fisher et al. (2017) found that perception of an argumentative exchange as competitive or cooperative influences the extent to which participants rate

⁵ Thanks to an anonymous referee for forcing me to think about this issue.

⁶ Goodwin & Darley (2010), p. 167 note that we cannot be sure whether the children's judgements reflect ideas about universality or mind-independence, which is one reason why this research is only very defeasible evidence for my empirical hypothesis. However, even a belief in universally shared responses might explain a commitment to (A) for certain kinds of judgements.

⁷ Again, the support provided by these studies is very limited, since impersonal harm-based wrongs are not similarly correlated with increased emotional arousal, and are yet thought of as more objective than certain other moral issues. Still, impersonal scenario's involving harm may trigger emotional arousal that is nevertheless stronger than scenario's that do not involve harm (or injustice) at all.

the issue at stake as objective or subjective. A competitive exchange is one in which participants try to win the debate, or to outcompete the other, while a cooperative exchange is one in which participants try to learn as much as possible from each other. It seems that if people prepare for a competitive exchange about a moral issue, then they are more inclined to consider it to have an objectively correct answer than when they prepare for a cooperative exchange. This study is the most direct evidence (of which I am aware) that non-cognitive processes can lead to beliefs in objectivity.

3. Instrumental Normativity and Shared Standards of Correctness

So far, I have argued two things: first, that Streumer's argument against a non-cognitive interpretation of (A) is inconclusive, and that a belief in (A) might be explained by the function of normative language, or the way normative judgements are embedded in a larger syndrome of cognitive, conative and affective states. However, even if beliefs in the objectivity of *moral* judgement can be caused by, say, emotional arousal, not all normative judgements are backed by the same processes. For instance, instrumental and epistemic normative judgements are not plausibly related to emotional arousal.⁸ If we do accept (A) for such judgements, then non-cognitivists cannot appeal to this particular mechanism in order to explain it (although the function of normative language might still help). In this section, I will argue for two things: first, that our commitment to (A) is actually limited in the case of instrumental normative judgements, and second, that the extent to which we do accept (A) for instrumental normative judgements can be explained by reference to an assumption about shared standards of correctness. I will address epistemic normativity in the next section.

Streumer gives the following example of an instrumental normative judgement:

(3) If you want to become a doctor, then you ought to study medicine.

⁸ Nor clearly with something like Blackburn's staircase of emotional ascent in (1998), p. 9.

Judgements like these are clearly subject to at least *some* shared standards of correctness. In particular:

- (i) If the recommended action is not a means to (or necessary condition for) achieving the end, then the instrumental normative judgement is incorrect,

and

- (ii) If the action is significantly worse as a means than available alternatives, then the instrumental normative judgement is also incorrect.⁹

So, if we assume that others (by and large, and perhaps under idealized conditions) apply the same criteria for counting an action as a means to an end, and apply the same standards for what counts as a better means, then we can explain why we think that at most one of two conflicting instrumental normative judgements can be correct. For even if such judgements involve some kind of approval or recommendation and cannot be reduced to non-normative claims, there are shared criteria for the correctness of the approval or recommendation.¹⁰

Apparently, though, criterion (ii) for the correctness of instrumental normative judgements itself involves a normative notion. It says that claims like (3) are incorrect if the prescribed action is significantly *worse* as a means than relevant alternatives. What does it take for a means to be worse than other means?

The quality of the means is at least in part a matter of efficiency, which is (perhaps) explicable in terms of the degree to which taking the means raises the probability of the end's realization. That degree can then be compared to the degree to which alternatives raise the probability. This criterion of being a good means is not

⁹ We may also consider an instrumental ought judgement *false* if the recommended means are not significantly better than others, since 'ought' suggests that the means are *best*.

¹⁰ And so the suggestion that there are shared standards for sufficiently effective means does not entail that instrumental normative judgements are not normative, or cannot be understood along non-cognitivist lines.

something which people assume varies between people. Of course, whether it does in fact vary depends both on the nature of probability and the threshold for what counts as *sufficient* probability raising. However, ordinary people are unlikely to have a theory of probability which would bar them from assuming that judgements about the relative likelihood of outcomes are generally shared. Furthermore, it seems reasonable to assume that people expect others to endorse similar thresholds for *good* means, since people also assume similar thresholds for what counts as goodness in a knife or hammer in the light of some objective.

I also think that if we did discover that others presuppose different thresholds for sufficient probability raising, we would not clearly endorse the idea that at most one of the judgements can be correct. It is important to keep in mind that we are considering a case where people agree about the extent to which the action is conducive to the end. What they disagree about is what it takes for a means to be sufficiently effective to count as good. The situation here is analogous to judgements about attributive goodness. Two people might agree on the ease with which a knife cuts through bread, and yet one might consider the knife to be good and the other bad. These people have different standards for sufficient ease, and it is not clear that at most one of their judgements can be correct.

Generally, though, people endorse (or assume that people endorse) similar standards for sufficiently effective means. This means that the sense that at most one of two conflicting instrumental normative judgements can be correct can be explained by the fact that the standards to which such judgements are subject are shared between people, or assumed to be shared.

Might this picture be too simple? There are two reasons for doubt. First, you might think that there are many cases where there are no shared standards concerning the best means to realizing a certain end. For example, there is disagreement about the best means to living a happy life.¹¹ However, in this case, it seems that the disagreement boils down to a disagreement about what *constitutes* a happy life. Once we fix on some idea about what it takes, we apply, or assume that people apply, very similar criteria for sufficiently effective means.

¹¹ I owe this example to an anonymous referee for this journal.

A second reason to doubt the claims defended so far is this: being a good means to an end may often require more than that the means are efficient, more than that they raise the probability of the end beyond some threshold. For although instrumental ought judgements bracket *many* considerations relevant to all-things-considered ought judgements, they need not bracket all. Many instrumental oughts seem to presuppose that the means are not seriously detrimental to other values assumed to be important. If I want to drive a nail into the wall, and judge that I ought to use a hammer, my sense of correctness may depend on the assumption that using the hammer does not break my windows or destroy my shoes.¹² This appears to be so even if the instrumental ought is not *all* things considered (I may think that I ought not to be driving nails into the wall in the first place.) Does the relevance of additional assumptions about value jeopardize my explanation of the commitment to (A) for instrumental normative judgements?

This depends on our reactions to cases where we know that people's instrumental judgements are made in the light of different assumptions about what is to be avoided or achieved. Do we think that at most one of them can be correct? This seems questionable. If people are clearly presupposing different aims, and are otherwise correct about the probabilities, then we tolerate their instrumental judgements as correct. The reason for this is that instrumental normative judgements are *essentially relative*: it is in their nature to recommend actions on the basis of efficiency, but only *relative* to certain aims. If these aims differ, and people are otherwise correct about the probabilities, then their judgements are equally correct as ours. This means that we don't really experience instrumental normative judgements made in the light of different aims or values as conflicting, and so (A) does not apply. But I have argued that non-cognitivists can explain why we accept (A) for those instrumental judgements that we do take to be conflicting.

4. Epistemic Normativity and Shared Standards of Correctness

¹² This is clearly not a matter of (thresholds for) efficiency in realizing the end of driving a nail into the wall.

I have argued that non-cognitivists can explain why we tend to accept (A) for instrumental normative judgements that presuppose the same ends, and that we reject (A) for instrumental normative judgements made in the light of different ends. But what about epistemic normativity?

It seems plausible to me that epistemic normative judgements are akin to instrumental normative judgements, and that they presuppose an end that is conceptually related to having true beliefs.¹³ This idea might enable the non-cognitivist to explain why we think that at most one of two conflicting epistemic judgements can be correct in a way similar to that proposed for instrumental normative judgements more generally: epistemic normative judgements are subject to (presuppositions of) shared standards of correctness.

Suppose that parties to an epistemic disagreement assume that they are judging relative to the same end. In that case, they should also assume that disagreements about what we ought to believe boil down to disagreements about how the relevant body of evidence relates to the realization of that end. In such cases, one party will take the evidence to make it sufficiently likely that the proposition or theory is true, while the other thinks that it does not. If they also assume similar standards for sufficient likelihood, and assume that there is a fact of the matter about such likelihoods, then they will reasonably believe that at most one of their conflicting judgements can be correct.¹⁴

¹³ Something like this is defended in Finlay (2014), although he, unlike non-cognitivists, defends a reductive understanding of instrumental normative judgements. Nevertheless, non-cognitivists can also accept that the standards of correctness for epistemic normative judgements are in part set by the end of having true beliefs. The idea is that what one ought epistemically to believe is a matter of what it is best to believe in order to achieve the end of having a true belief with respect to the relevant issue.

¹⁴ A referee for this journal notes that Streumer might take judgements of likelihood to be normative judgements about the credences one ought to have. But, first, it is not clear from anything in his book that Streumer accepts a normative account of judgements about likelihoods, and, second, the correctness of such an account need not be apparent to ordinary speakers (they might still believe that there are non-normative facts about likelihoods). However, it has to be acknowledged that my suggestions in defence of non-cognitivism are less effective if judgements of probability were themselves normative judgements.

You might doubt that this can fully explain the sense that (A) is true in the case of epistemic oughts, for at least two reasons. The first is that our sense that at most one of two conflicting judgements is correct might survive the discovery that different parties hold different standards for sufficient likelihood. Second, you might think that there is not one uniquely rational response to a given body of evidence.¹⁵ I'll take these cases in turn.

Suppose we discover that someone has extremely low standards for what counts as sufficient likelihood of truth. As a result, she judges that it is reasonable for her to believe a proposition that we take to be very weakly supported by the evidence. Would we not reject her judgement as (uniquely) incorrect? I take the answer to this question to be Yes, but I think that can be explained as follows: someone who has very low standards for sufficient likelihood is not *really* interested in true beliefs. Since, as we are supposing, epistemic judgements are conceptually linked to some end related to it, we can reasonably reject her judgement as mistaken.

The second source of doubt is as follows: some philosophers believe that there is not one uniquely rational response to a given body of evidence (for discussion, see Kopec & Titelbaum 2016). If no single belief is rational given a body of evidence, then perhaps one person might, reasonably, judge that she ought to believe p , while another might, reasonably, judge that she ought not to believe p . Furthermore, this can be so even if they agree on the extent to which the evidence makes p likely. Would they not equally judge that at most one of their ought judgements is correct?

I think there are two cases to consider. One case is such that the parties to the disagreement *believe* that there is no uniquely rational response to the relevant body of evidence. It seems to me that this belief should reduce confidence in the judgement that at most one of the two conflicting judgements is correct. The second case is such that the parties do not believe that there is only one uniquely rational response to the relevant body of evidence. In that case, they might judge that at most one of their respective judgements can be correct, but this can be explained by the assumption that

¹⁵ This second difficulty was raised by an anonymous referee for this journal.

the evidence either does or does not make the contested belief sufficiently likely to be true.¹⁶

More could be said about epistemic normativity, but it does seem worth exploring the suggestion that we assume shared standards of correctness for epistemic normative judgements in those cases where we also believe that at most one of those judgements can be correct. This would help the non-cognitivist to explain why (A) seems true in the case of epistemic normative judgements.

5. Reducing Instrumental and Epistemic Normativity

If what I have argued in the previous two sections is correct, then reductivists about instrumental and epistemic normative judgements can also help themselves to it. Reductivists believe that instrumental and epistemic normative claims aren't *irreducibly* normative, which means that they can be reduced to non-normative claims about, say, the relation between actions and desire-satisfaction, or end-realization (a form of naturalism about such judgements). Streumer thinks that this is a mistake. He argues that if instrumental normative claims were conceptually reducible to claims about non-normative relations, then we should be able to deduce the relevant ought claims directly from claims couched in non-normative vocabulary, which he thinks is not the case.

This argument may not persuade those firmly convinced that instrumental normative claims are reducible to claims about non-normative relations. For if they can be so reduced, then the statement that they cannot be deduced from claims describing those non-normative relations in non-normative terms seems question-begging.

¹⁶ A referee for this journal claims that my explanation of our commitment to (A) may fail if contextualism about knowledge claims is true. Contextualism is the view that the standard of knowledge varies with context (see e.g. DeRose (1992)). However, the plausibility of contextualism depends on intuitions to the effect that what it takes to know that p in one context may be less demanding than what it takes to know that p in other contexts. If we do have such intuitions, then we should not really accept that at most one of two apparently conflicting knowledge ascriptions can be correct once it is revealed to us that the conflicting ascriptions are made in different contexts.

However, even if one accepts that instrumental normative claims are not conceptually equivalent to non-normative claims, and can't be *logically* derived from claims couched in non-normative vocabulary, one might still think that the fact that they describe *is* a non-normative relation.

At this point, Streumer would appeal to another argument against the idea that instrumental normative claims describe non-normative relations. According to Streumer, normative judgements are essentially such that

- (G) There are no descriptively specified conditions in which people's normative judgements are guaranteed to be correct.

According to Streumer, (G) characterizes the nature of normative discourse, in the sense that anything that does not satisfy (G) does not count as normative discourse in the first place.¹⁷ One might think that instrumental normative judgements do not satisfy (G), since there may appear to be descriptively specified conditions in which people's instrumental normative judgements are guaranteed to be correct. Take, for instance

- (4) If you want to become a doctor, then you ought to study medicine.

You might think that if a person considered all descriptive information relevant to becoming a doctor, then her instrumental normative judgements are guaranteed to be correct. But Streumer disagrees:

'[s]uppose that Bob is deeply attached to homeopathy. Suppose that whenever he is confronted with evidence against the effectiveness of homeopathy, he endorses ad hoc hypotheses to explain away this evidence. And suppose that, after considering all relevant descriptive information, he would think that if you want to be a doctor, you ought to study homeopathy. Is Bob's judgement then guaranteed to be correct?' (p. 113)

¹⁷ A judgement satisfies (G) just in case (G) is true of it, i.e. if it is true of the judgement that there are no descriptively specified conditions in which it is guaranteed to be correct.

Streumer thinks the answer to this question is clearly No.

What is surprising about the previous argument is that it also applies to judgements that are clearly non-normative. If the reason why there are no descriptively specified conditions in which normative judgements are guaranteed to be correct is that even a person who considers all relevant information can invoke *ad hoc* hypotheses, then the same is true of claims about the origin of species, the composition of the Earth, etc. What this means, I think, is that the reason why instrumental normative judgements fail to meet (G) is not a reason to think that the relation described by instrumental normative judgements is something other than descriptive, as uncontroversially descriptive claims meet (G) as well (and for the same reason as instrumental normative judgements).

If the foregoing is right, then Streumer's argument against the idea that judgements about epistemic reasons describe non-normative facts also fails. Here too, Streumer says that there are no descriptively specified conditions in which judgements about epistemic reasons are guaranteed to be correct. For even if someone has considered all relevant descriptive information, she may still arrive at apparently mistaken judgements about what there is reason to believe. For instance, by invoking *ad hoc* hypotheses, such a person may still judge that we have little reason to believe the theory of evolution. If I am right, then this is not significant evidence that judgements about epistemic reasons do not concern descriptive facts.

Furthermore, my explorations on behalf of non-cognitivism suggest a (somewhat) positive characterization of the conditions under which someone's instrumental normative judgements are guaranteed to be correct. They are guaranteed to be correct insofar as the person abides by my two earlier principles governing the correctness of instrumental normative judgements, namely:

- (i) If the recommended action is not a means to (or necessary condition for) achieving the end, then the instrumental normative judgement is incorrect,

and

- (ii) If the action is significantly worse as a means than available alternatives, then the instrumental normative judgement is also incorrect.

In my characterizations of what it is for a means to be worse than some other means, I did not appeal to irreducibly normative properties. For instance, I said that the quality of the means is a matter of efficiency and standards for being sufficiently effective in the realization of various ends. Similarly in the case of epistemic normative judgements. I appealed to judgements about likelihood, and considerations relating to the realization of ends like having true beliefs.

I also suggested that instrumental normative judgements are essentially relative to ends, such that disagreement intuitions disappear when it is made explicit that two apparently conflicting judgements are in fact made in light of different ends. If so, then this standard style of argument against reductivism about non-instrumental normative judgements cannot be used to refute reductivism about instrumental normative judgements. If I am right that epistemic normative judgements are a species of instrumental normative judgements, then the same holds for the latter.

6. Conclusion

According to Streumer, non-cognitivists cannot explain our commitment to (A), the thesis that when two people make conflicting normative judgements, at most one of them is correct. I think his arguments are inconclusive. I also think that our commitment to (A) is not equally strong for all normative judgements, and that the extent to which we accept it for instrumental and epistemic normative judgements can be explained by an assumption of shared standards of correctness. This means that there is hope both for non-cognitivism as a general theory of normative judgement, and reductivism about instrumental and epistemic normativity.

Acknowledgements

I would like to thank Hendrik Siebe, Bart Streumer, Christine Tiefensee, Anandi Hattiangadi and three anonymous referees for the Journal of Moral Philosophy for helpful feedback on this paper.

References

Blackburn, S. (1998), *Ruling Passions*, Oxford: Clarendon Press

DeRose, K. (1992), "Contextualism and Knowledge Attributions", *Philosophy and Phenomenological Research* 52(4): 913-929

Finlay, S. (2014), *Confusion of Tongues. A Theory of Normative Judgement*, Oxford: Oxford University Press

Fisher, M., Knobe, J., Strickland, B. & Keil, F. (2017), "The Influence of Social Interaction on Intuitions of Objectivity and Subjectivity", *Cognitive Science* 41, pp. 1119-1134

Goodwin, G. & Darley, J. (2010), "The Perceived Objectivity of Ethical Beliefs: Psychological Finding and Implications for Public Policy", *Review of Philosophical Psychology* 1, pp. 161-188

Goodwin, G. & Darley, J. (2008), *The Psychology of Meta-Ethics: Exploring Objectivism*, *Cognition* 106, pp. 1339 - 1366

Greene, J., Sommerville, R., Nystrom, L., Darley, J. & Cohen, J. (2001), "An fMRI Investigation of Emotional Engagement in Moral Judgment", *Science* 293(5537), pp. 2105-2108

Greene, J., Nystrom, L., Engell, A., Darley, J. & Cohen, J. (2004), "The Neural Bases of Cognitive Conflict and Control in Moral Judgment", *Neuron* 44(2), pp. 389-400

Kopec, M. & Titelbaum, M. (2016), "The Uniqueness Thesis", *Philosophy Compass* 11(4), pp. 189-200

Nichols, S. & Folds-Bennett, T. (2003), "Are Children Moral Objectivists? Children's Judgments about Moral and Response-Dependent Properties", *Cognition* 90(2), pp. 23-32

Sarkissian, H., Park, J., Tien, D., Cole Wright, J. & Knobe, J. (2011), "Folk Moral Relativism", *Mind and Language* 26(4), pp. 482-505

Streumer, B. (2017), *Unbelievable Errors*, Oxford: Oxford University Press