

Metaethical Implications

Cognitivism versus Non-Cognitivism

Introduction

So far, our discussion has almost entirely been focused on normative issues, about what is morally choiceworthy and what is appropriate to do in the face of moral uncertainty; apart from our discussion of metaethical nihilism in the previous chapter, we have not discussed issues of metaethics. In this chapter, however, we will show that moral uncertainty creates a challenge for another metaethical view, namely, non-cognitivism, according to which moral judgements are desires, or some other desire-like states, rather than beliefs. We will show that it is surprisingly difficult, though perhaps not impossible, for non-cognitivists to accommodate moral uncertainty.

Now, one could of course turn this argument on its head and say the fact that non-cognitivism cannot accommodate moral uncertainty (if it is a fact) shows that there is no such thing as moral uncertainty. This would be an incredible thing to say, however, since it seems so obvious that we can be uncertain about fundamental moral matters—just recall the intuitive examples we gave in Chapter 1. Furthermore, the leading non-cognitivists of today agree that it is important to accommodate fundamental moral uncertainty in a way that does not force them to give up on mundane facts such as that we can be more or less certain that an action is right (including the possibility that one can be fully certain that an action is right), and that we can be less certain that an action is right than that some non-moral proposition is true.¹ Indeed, they think it is important to ‘earn the right’ to other realist-sounding notions as well, such as ‘truth’, ‘fact’, and ‘evidence’,

¹ See, for instance, James Lenman, ‘Non-Cognitivism and the Dimensions of Evaluative Judgement’, Brown Electronic Article Review Service, 15 March 2003. <http://www.brown.edu/Departments/Philosophy/bears/homepage.html>; Michael Ridge, ‘Ecumenical Expressivism: The Best of Both Worlds?’, *Oxford Studies in Metaethics*, vol. 2 (2007), pp. 51–77; Simon Blackburn, ‘Dilemmas: Dithering, Plumping, and Grief’, in H. E. Mason (ed.), *Moral Dilemmas and Moral Theory*, Oxford: Oxford University Press, 1996, pp. 127–39.

since these notions permeate ordinary moral talk. So we do not think we are stacking the cards against non-cognitivism by assuming that fundamental moral uncertainty is a fact that needs to be accommodated by any plausible metaethical theory.

I. The Challenge for Non-cognitivism

According to a simple form of non-cognitivism, to make a moral judgement is to express a desire (a non-cognitive attitude) rather than a belief (a cognitive attitude). As we pointed out above, there is general agreement even among non-cognitivists that a metaethical theory needs to explain (and not explain away) fundamental moral uncertainty. Since moral uncertainty involves having some degree of certainty, less than full, in a moral judgement, one cannot accommodate moral uncertainty without accommodating degrees of certainty (call this *certitude*). So, non-cognitivists need to be able to give an account of degrees of moral certainty.

Another feature in need of explanation is the uncontroversial fact that we can ascribe degrees of value or normative importance to states of affairs or actions (call this *importance*). The challenge for non-cognitivism is that desires seem to have too little structure to account for both certitude and importance, where certitude is assumed to obey at least some of the axioms of probability theory. If certitude is identified with degrees of desire strength (so that the stronger the expressed desire is, the more certain you are about the moral judgement), then there is nothing left to explain importance. On the other hand, if importance is identified with degrees of desire strength (so that the stronger the expressed desire is, the more value or importance is ascribed to the act or state of affairs), then there is nothing left to explain certitude. Of course, this would not be a problem if certitude and importance always co-varied, but that is not true. One might, for example, invest *low* certitude in the belief that leading an autonomous life is of *great* intrinsic value and a *great* degree of certitude in the belief that experiencing bodily pleasure is of *moderate* intrinsic value. Similarly, one might have *low* certitude in the belief that one has a *strong* reason to save a stranger's two children at the cost of the life of one's own child and *high* certitude in the belief that one has a *weak* reason to satisfy one's whims.²

² This way of stating the problem for non-cognitivists is found in Michael Smith, 'Evaluation, Uncertainty, and Motivation', *Ethical Theory and Moral Practice*, vol. 5 no. 3

Obviously, the cognitivist, who identifies moral judgements with beliefs, has no problem capturing these cases. Degrees of moral judgements are simply degrees of beliefs and the degree of belief in a moral proposition can vary independently from the degree of moral importance ascribed to an action or a state of affairs. But for the non-cognitivist it is a real challenge.

The non-cognitivist could try to break out of this dilemma by identifying degrees of certitude with the degree to which one is insensitive to new information and reflection, so the less one is prone to change one's expressed desires in light of new information and reflection, the more certain one is about one's moral judgements. But this would be to conflate certitude with *robustness*. One can have high certitude in a moral judgement that is not robust. Just think of someone who vacillates between strong moral opinions depending on which newspaper she reads and which TV programme she watches. Or think about someone who starts off with very low credence in a moral judgement she has not seen much evidence for, e.g. that insects have moral rights. When more information comes in, she may retain the same low level of certainty in this judgement, i.e. the same high level of certainty in the judgement that insects do not have rights, even though the robustness of her judgement that they do not have rights increases.

The challenge to account for moral uncertainty also spells trouble for non-cognitivists who think that moral judgements express intention-like states, such as plans. Since it is doubtful that intentions or plans come in degrees, it seems *prima facie* difficult to account for certitude and importance.³ The challenge is thus relevant to Gibbard's recent account, according to which judging that one ought, all things considered, to do something is to plan to do it.⁴ The challenge is thus relevant to a wide family of non-cognitivist views.

A possible non-cognitivist rejoinder is to represent importance in terms of first-order attitudes and certitude in terms of *second-order* attitudes. This version of non-cognitivism says that a moral judgement that some action type, ϕ , is right expresses a second order desire—a desire to desire to ϕ . For instance, when Mary judges that she has a reason to keep her promises, she

(September 2002), pp. 305–20. In fact, the argument was broached a decade earlier in a textbook by Lars Bergström (*Grundbok i värde teori*, Stockholm: Thales, 1990, pp. 35f).

³ At least this seems so for *all-out* intentions that constitute decisions to act, which is the relevant notion of intention for an account of all things considered judgements of moral rightness. For the notion of partial intentions that function as inputs for decisions, see Richard Holton, 'Partial Belief, Partial Intention', *Mind*, vol. 117 (2008), pp. 27–58.

⁴ Allan Gibbard, *Thinking How to Live*, Cambridge, MA: Harvard University Press, 2003.

expresses a desire to desire that she keeps her promises. The strength of the first-order desire could be taken to represent importance. So the more strongly Mary desires to keep her promises, the more important she considers it to keep her promises. The strength of the second-order desire, on the other hand, represents her degree of certitude in the claim that that keeping her promises is right. So the more strongly she desires to desire to keep her promises, the more certain she is that it would be right for her to keep her promises. This proposal, which we might call the *attitudinal ladder proposal*, faces several objections.

First, the attitudinal ladder proposal is plagued by arbitrariness: what is the rationale for representing importance in terms of the strength of the relevant first-order desire and certitude in terms of the strength of the relevant second-order desire rather than the other way around?⁵

Second, unlike degree of belief, desire-strength does not come in a neat interval with a clearly defined minimum and maximum. Certitude can vary from complete uncertainty (credence 0) to complete certainty (credence 1), but there is no obvious analogy for desire strength. Complete indifference might be seen as the weakest possible desire, but what is it to *completely* desire (to desire) something? As we shall see, this problem recurs for other forms of non-cognitivism to be considered below.

A final problem for the attitudinal ladder proposal is that strengths of second-order attitudes and moral certitude may come apart. According to the attitudinal ladder proposal, the stronger Mary desires to desire to ϕ , the more certain she is that she has reason to ϕ . But it seems possible that there are cases in which Mary desires to desire to ϕ without being at all certain that she has reason to ϕ . Think of a case where an evil demon threatens to harm your family if you do not desire always to keep your promises. The demon does not care about whether you actually keep all your promises; he cares only about whether you desire to do so. In this case, you may well strongly desire to have the desire that you keep all your promises while you lack certainty that you actually have reason to keep all your promises.

II. Ecumenical Non-cognitivism

It has recently become popular to argue that *ecumenical* non-cognitivism is the version of non-cognitivism that is best equipped to meet the challenge

⁵ Smith, 'Evaluation, Uncertainty, and Motivation', p. 318.

of moral uncertainty and thus to accommodate both certitude and importance. As the label suggests, ecumenical non-cognitivism is a hybrid view that incorporates both cognitivist and non-cognitivist components. It inherits from traditional non-cognitivism the idea that moral judgements express desires and it inherits from cognitivism the view that moral judgements express beliefs.⁶ What makes it an ecumenical version of non-cognitivism rather than an ecumenical version of cognitivism is that the contents of the beliefs expressed in moral judgements do not provide truth conditions of moral judgements and sentences.⁷

More specifically, the idea is that a moral judgement concerning, e.g. the rightness of an action expresses (1) a general desire for actions insofar as they have a certain natural property; and (2) a belief that the particular action in question has that property. The belief component thus makes direct reference back to the property mentioned in the desire component.⁸

According to the simplest version of ecumenical non-cognitivism, what the property in question is depends on the first-order moral view endorsed by the speaker. To give a simple example, a utilitarian who endorses the judgement that sticking to a vegetarian diet is morally right expresses approval of actions insofar as they tend to maximize overall happiness and a belief that sticking to a vegetarian diet has *that property*, i.e. the property of tending to maximize overall happiness. The sentence 'Sticking to a vegetarian diet is morally right' lacks truth-value, since the expressed belief is not assumed to provide the truth-conditions for this sentence.

Since ecumenical non-cognitivists claim that moral judgements express both desires and beliefs, they seem to be in a better position to accommodate both importance and certitude. One obvious solution is to say that certitude is represented by the strength of the *belief* expressed by a moral judgement. So, one's certitude that sticking to a vegetarian diet is morally right is identified with the degree of one's belief that sticking to a vegetarian diet tends to maximize overall happiness. Importance could then be

⁶ As James Lenman reminded us, the idea that moral judgements have both non-cognitive and cognitive meaning is not new. R. M. Hare famously argued that the primary meaning of moral judgements is prescriptive, while their secondary meaning is descriptive. See *The Language of Morals*, Oxford: Clarendon Press, 1952, ch. 7. The similarities between Ridge's ecumenical expressivism and Hare's prescriptivism are discussed in John Eriksson, 'Moved by Morality: An Essay on the Practicality of Moral Thought and Talk', dissertation, Uppsala University, 2006, pp. 199–204.

⁷ Ridge, 'Ecumenical Expressivism', p. 54; Lenman, 'Non-cognitivism and the Dimensions of Evaluative Judgement', sect. 2.

⁸ Ridge, 'Ecumenical Expressivism', p. 55.

represented by the strength of the expressed desire. Given that one believes that sticking to a vegetarian diet tends to maximize overall happiness, the more strongly one approves of actions insofar as they tend to maximize overall happiness, the more moral reason one thinks one has to stick to a vegetarian diet.

One advantage of this account is that it does not have to translate degrees of certainty into degrees of desire, since certainty is here represented by degrees of belief. Another advantage is that it can allow for some cases of motivational maladies. One can judge that one ought to do something and yet still lack a desire to do it. It is true that one's judgement will always express a desire, but it is a *general* desire to do actions insofar as they have a certain property, not a desire to do a particular action. Arguably, a general desire to do an action of a certain type can exist in the absence of a desire to do a particular action of this type. One might, for instance, have a general desire to do some work today, but, in a state of listlessness, fail to have any desire to do some particular kind of work.

The most serious drawback of the account is that it seems unable to capture distinctively *moral* certitude. One's certainty that sticking to a vegetarian diet is morally right depends on one's empirical certainty that doing this tends to maximize overall happiness, but it also depends on one's *moral* certainty that maximizing overall happiness is a morally relevant feature. If one comes to doubt that sticking to a vegetarian diet tends to maximize overall happiness, one will also come to doubt that doing this is morally right. But it is also true that if one comes to doubt that maximizing overall happiness is a morally relevant feature, one will come to doubt that vegetarianism is the morally right option. The ecumenical non-cognitivist seems able to capture only one's *empirical* certainty, not one's distinctively *moral* certainty.

Lenman and Ridge, who both are sensitive to this objection, have advocated a slightly more complex version of ecumenical non-cognitivism. Following Ridge, we will call it the Ideal Advisor version. According to this version, a moral judgement concerning the rightness of an action expresses (1) a desire for actions insofar as they would be approved of by a certain sort of *ideal advisor* and (2) a belief that makes direct reference back to the property of being approved of by that advisor.⁹ To judge that sticking to a vegetarian diet is morally right is on this view to express (1) a desire for actions insofar as they would be approved of by one's ideal advisor and (2) a

⁹ Lenman, 'Non-Cognitivism and the Dimensions of Evaluative Judgement', sects 2, 4; Ridge, 'Ecumenical Expressivism', p. 57.

belief that sticking to a vegetarian diet would be approved of by one's ideal advisor. Different people might have different ideal advisors. Utilitarians, for instance, would think of the ideal advisor as someone who only approves of actions that would maximize overall wellbeing, whereas Kantians would think of the ideal advisor as someone who only approves of actions that are based on universalizable maxims. It is crucial, of course that the belief about one's ideal advisor can be spelled out in purely non-moral terms. Otherwise, this belief would not be a proper belief according to the non-cognitivist.

Common to both Lenman's and Ridge's Ideal Advisor theories is the idea that importance is represented by the motivational strength of the ideal advisor's desires. So, how much reason one takes there to be for one to keep one's promises is represented by how much one thinks one's ideal advisor would want one to keep one's promises. This differs from the simpler view sketched above, according to which importance was identified with the motivational strength of the actual agent's general desires for actions. While Lenman's and Ridge's treatments of importance are closely related, they differ in their treatments of certitude, as we will see.

III. Lenman's Version of Ecumenical Expressivism

Lenman's response to the challenge starts with the proposal that the non-cognitivist can give the following story about moral uncertainty: when a subject, *S*, wonders whether ϕ -ing is wrong, *S* wonders whether an *improved* version of *S* would disapprove of ϕ -ing. Certitude is then given by *S*'s degree of belief in what an Improved *S* recommends, while importance is given by the strength of the Improved *S*'s (dis)approval. This story is consistent with non-cognitivism as long as the improved version of *S* is described in purely descriptive terms and as long as the judgement that this version of *S* is improved expresses primarily a pro-attitude to this version of *S* (that is, as long as the truth condition of the judgement that some version of *S* is ideal is not provided by the content of a belief). Clearly, Actual *S* can be more or less certain that Improved *S* would disapprove of ϕ -ing.

On Lenman's view, then, *S*'s judgement that she ought not to lie expresses a positive attitude to conforming with the desires of Ideal *S* (where 'Ideal' is a descriptive specification of a hypothetical version of *S* that Actual *S* endorses as improved in the sense that Actual *S* takes the desires of Ideal *S* to be action-guiding), together with a belief that Ideal *S* would desire that *S*

does not lie.¹⁰ For illustrative purposes, let us follow Lenman and say that Ideal S is simply a fully informed and clever version of S, where 'fully informed' and 'clever' are understood descriptively.

Now, S's certitude that she ought not to lie is represented by the strength of her belief that Ideal S would desire that Actual S does not lie. The strength of S's desire not to lie represents how much S cares (*de re*) about acting in accordance with her moral judgements, and the strength of S's general desire to act in accordance with what Ideal S would desire that S does, represents how much S cares (*de dicto*) about acting in accordance with her moral judgements.

Even though Lenman's account can explain certitude in particular moral judgements, such as the judgement that one ought not to lie, it leaves out an important dimension of moral uncertainty: it has no resources to represent uncertainty regarding whether being (dis)approved of by a certain descriptively specified improved agent is indicative of rightness (or wrongness). This means that Lenman's account is unable to account for *fundamental* substantive moral uncertainty.

As Lenman recognizes and as we have hinted more than once, the non-cognitivist must, on pain of circularity, offer a purely descriptive specification of improved agents. But for any descriptive specification of improved agents, there is room for uncertainty as to whether anything of which the improved agent would approve really is right. Lenman's proposal can only account for uncertainty as to whether a descriptively specified ideal agent would (dis)approve of certain actions. But this is uncertainty about purely *empirical* matters of fact, it is not *moral* uncertainty. Lenman's Ideal Advisor version of ecumenical non-cognitivism thus faces a problem similar to the one faced by the simpler version discussed in the previous section. As we shall see in the next section, Ridge is aware of this problem and proposes a way to deal with it.

IV. Ridge's Version of Ecumenical Expressivism

Ridge's favoured version of ecumenical expressivism has several affinities with the account suggested by Lenman. One important difference, though, concerns the representation of certitude. According to Ridge,

¹⁰ Lenman notes that Simon Blackburn has suggested a similar account. See Blackburn, *Ruling Passions*, Oxford: Clarendon Press, 2001, pp. 261–9.

An agent's certainty that he should ϕ is represented by two factors: (a) his certainty (in the ordinary sense) that ϕ -ing would be approved of by the relevant sort of advisor; and (b) the relative strength of his pro-attitude in favour of actions insofar as they would be approved of by the relevant sort of advisor.¹¹

The discussion in the previous section explains why something like clause (b) is called for. We saw that certitude cannot be represented simply in terms of (a) since for any (non-moral) property F (e.g. the property of tending to maximize overall happiness or the property of being favoured by a descriptively specified ideal advisor), a subject can be perfectly certain that an object is F and less than certain that F is a right-making property, or in some other way indicative of rightness.¹² In other words, the addition of (b) is supposed to make Ridge's account succeed where Lenman's account failed, *viz.* in accounting for fundamental substantive moral uncertainty. However, to say that an addition like clause (b) is called for and that it is supposed to fill the gap in Lenman's account is by no means to say that it makes the ecumenical expressivist's account of certitude plausible. In fact, we shall argue in the following section that it makes the account indefensible.

But first, we should take notice of yet another clarification that Ridge makes and that will be relevant to the dilemma we will pose for ecumenical expressivism below. This is the assumption that the motivational strength in (b) is not *absolute* but *relative* to the strengths of the agent's other desires. The point of this assumption is to block the implausible implication that 'perfectly general motivational maladies (depression and listlessness, say) would count as undermining an agent's certainty in all of her moral judgements.'¹³ That is to say that even if the absolute strengths of each of an agent's desires are weakened by depression, listlessness, or the like, the relative strength of her desire in (b) might stay the same. (In what follows, we use the terms 'desire' and 'pro-attitude' interchangeably.)

V. Initial Problems and Cross-Attitudinal Comparisons

We know that according to Ridge's ideal advisor version of ecumenical expressivism, a subject S 's certitude that some action is morally right is a

¹¹ Ridge, 'Ecumenical Expressivism', p. 71.

¹² Cf. Ridge, 'Ecumenical Expressivism', pp. 71f.

¹³ Ridge, 'Ecumenical Expressivism', p. 72.

function of (a) the degree of *S*'s belief that that action would be approved of by the relevant sort of advisor, and (b) the relative strength of *S*'s desire for actions insofar as they would be approved of by the relevant sort of advisor.

Here we shall briefly raise some initial worries about how to interpret this function. We need a procedure that will aggregate degrees of beliefs and (relative) strength of pro-attitudes, so that degree of certitude can only vary from 0 ('complete uncertainty') to 1 ('complete certainty'). Such a procedure is feasible only on the assumption that there are minimum and maximum degrees of desire-strength. As we noted above, complete indifference might be seen as the natural minimum degree of desire-strength but there is no natural maximum degree of desire-strength.

This point is relevant to the possibility of what we might call *cross-attitudinal comparisons*. Intuitively, we should be able to make sense of comparisons of certitude between moral and non-moral beliefs. For instance, a subject *S* can be more certain that $2 + 2 = 4$ than that utilitarianism is true. But if *S*'s certitude that utilitarianism is true is a function of the degree of *S*'s belief that an ideal advisor would favour actions insofar as they maximize utility and the relative strength of *S*'s pro-attitude to actions insofar as they would be favoured by an ideal advisor, we need to be able to make comparisons in strength between beliefs (such as $2 + 2 = 4$) on the one hand, and combinations of beliefs and desires on the other hand. But what does it mean to say that a belief is stronger than the combination of a belief and a desire? Making sense of such comparisons seems to require a joint scale for beliefs and desires. But, as we have already seen, it is far from clear how to construct such a scale.

VI. A Dilemma

Moral certitude is supposed to be a function of a belief and a desire. But exactly what function? Ridge does not tell us, but one can show that it faces a serious dilemma. To uncover the first horn, recall once again that Ridge defines moral certitude partly in terms of relative desire-strength. The relative strength of *S*'s desire *D* is most naturally defined—in analogy with relative price—in terms of the ratio between the strength of *D* and a weighted average of the strengths of all of *S*'s other desires. This means that if the absolute strength of *D* remains the same while the absolute strengths of some other desires of *S*'s increase, then the relative strength of *D* decreases. Correspondingly, if the absolute strengths of *S*'s other desires

decrease while D 's absolute strength remains the same, the relative strength of D increases.

This feature of relative desire-strength generates absurd results. Suppose that you fall in love with a person you have known for many years and, as a result, the strengths of your desires concerning this person shoot up. If the strength of your desire for actions insofar as your ideal advisor would approve of them remains the same, which it is likely to do in most cases since a romantic endeavour need not affect moral commitments, then the relative strength of this desire decreases. But on Ridge's theory this implies that your moral certitude has decreased. Perhaps love can sometimes make you doubt morality, since 'in love everything is permissible', but it is surely not a necessary consequence of falling in love and feeling a strong desire for someone that your moral certitude thereby diminishes. Moral certitude cannot depend on the strength of non-moral desires in this way. Of course, the same point can be made using any non-moral desire, not just love. For instance, if one's desire for eating ice-cream becomes stronger and the strengths of one's other desires stay the same, one's moral certitude has decreased.

It is equally obvious that examples can be given that work in the opposite direction. For instance, if one falls out of love with a person and the strengths of one's desires concerning this person consequently diminish, or if the strength of one's desire for ice-cream diminishes, the relative strength of one's desire for actions, insofar as one's ideal advisor would approve of them, increases. On Ridge's view, this means that one's moral certitude increases. But it is implausible that one's moral certitude is determined in this direct manner on one's falling out of love or on one's desires for ice-cream.

Ridge could reply by defining moral certitude partly in terms of *absolute* rather than *relative* desire-strength, but then he is caught on the second horn of the dilemma. As he himself points out, defining moral certitude in terms of absolute desire-strength would have the unwelcome result that wide-ranging motivational maladies, such as depression and listlessness, will always wipe out one's moral certitude. Recall that for Ridge moral certitude depends in part on the strength of one's general desire to perform actions insofar they would be approved by the ideal advisor. But depression and listlessness can sap one's general desires to perform actions with certain properties without one's moral certitude being greatly affected.

VII. Sepielli's Account

Instead of going for an ecumenical version, Sepielli has recently defended a non-cognitivist account of moral certitude that draws heavily on Schroeder's recent treatment of expressivism.¹⁴ Whereas non-cognitivism is a theory about the nature of the mental state we are in when we make a moral judgment, expressivism is semantic theory, according to which the meaning of a moral statement is identified with (or determined by) the mental state it is conventionally used to express. Most non-cognitivists accept expressivism and are thus committed to the idea that the meaning of a moral statement is the *non-cognitive* mental state (e.g. a desire) it expresses. A notorious problem for expressivism is to account for the meaning of moral terms when they are embedded in complex sentences. This problem goes by various labels, such as the Frege–Geach problem, the problem of embedding, or the negation problem. On Sepielli's view, non-cognitivism can account for moral certitude only if expressivism has enough structure to solve the Frege–Geach problem. Sepielli's basic manoeuvre is to apply Schroeder's recent treatment of the Frege–Geach problem to the problem of moral certitude.

Let us follow Schroeder and focus on negation. Consider the sentence:

(1) Jon thinks that murdering is wrong.

There are various places where we can insert a negation in this sentence, e.g. the following:

(2) Jon thinks that murdering is not wrong.

Now, expressivists face at least two challenges. The first is to explain what kind of non-cognitive attitude (2) attributes to Jon. The second is to explain why this attitude is inconsistent with the attitude attributed to Jon in (1). Advocates of traditional expressivism might want to say that (1) attributes to Jon a negative attitude to murdering while (2) attributes to Jon an attitude of toleration to murdering. But then it remains to be explained why a negative attitude to murdering and toleration of murdering are inconsistent attitudes. Since the two attitudes have the same content (murdering), the alleged inconsistency between them cannot be explained in terms of their

¹⁴ Sepielli, 'Normative Uncertainty for Non-Cognitivists'.

content. According to Schroeder's diagnosis, traditional expressivism has too little structure to meet the second challenge, so his fix is to add more structure. He does so by introducing the attitude of *being for*. The idea is that to think that an action is wrong, right, etc., is to take the attitude of being for some other attitude to the action in question. To illustrate, sentence (1) should be understood as

(1') Jon is for blaming for murder

and (2) as

(2') Jon is for not blaming for murder.

It is easy to see that the content of the attitude attributed to Jon in (1') is inconsistent with the content of the attitude attributed to Jon in (2').¹⁵

Sepielli's response to this challenge is to claim that degrees of being for are for the non-cognitivist what degrees of belief are for the cognitivist. According to this 'being for' account of normative certitude, being highly certain that murder is wrong is to be strongly *for* blaming for murder. Sepielli identifies degrees of normative importance with degrees of blaming, so having some degree of certitude that there is strong reason not to murder comes out as being, to some degree, for strongly blaming for murdering. As can be readily seen, the degree of being for and the degree of blaming can vary independently. For example, one can be strongly for weakly blaming for not paying taxes, which would capture the case when one is very certain that not paying taxes is a minor wrong. Similarly, one can be weakly for strongly blaming for eating factory-farmed chicken, which would capture the case where one is not so confident that eating factory-farmed chicken is a major wrong.

Another challenge for the non-cognitivist, which we introduced above in our discussion of the ecumenical account, is to solve the 'normalization problem'. This is the problem of finding natural minimum and maximum levels of normative certitude. To elaborate, certitude varies from 0 to 1. To be *fully certain* that *P* is to believe, to degree 1, that *P*, and to have 0 degree of belief that not-*P*. To be *minimally certain* that *P* is to believe, to degree 0, that *P*, and to have 1 degree of belief that not-*P*. Finding minimum and

¹⁵ It is perhaps not obvious that this makes the attitudes in (1') and (2') inconsistent. However, for an argument that expressivists can legitimately assume that being for is an 'inconsistency-transmitting' attitude, see Mark Schroeder, *Being For*, Oxford: Oxford University Press, 2008, pp. 42–3, 59–60.

maximum levels of normative certitude is not a problem for cognitivists, since they understand normative certitude straightforwardly in terms of degree of belief, which is taken to vary between 0 and 1.¹⁶ But there is no obvious analogue for non-cognitive, desire-like, attitudes. As suggested above, complete indifference might be seen as the weakest possible desire, but what is it to completely desire something or to desire something to degree 1? In what sense can non-cognitive attitudes be said to vary between 0 and 1? Before we come to Sepielli's solution of the normalization problem, we shall see that the being for account is vulnerable to several problems that can also be pressed against other forms of non-cognitivism.

VIII. Problems for the *Being For* Account of Normative Certitude

Sepielli says that he accepts many of our objections to Lenman and Ridge,¹⁷ but he overlooks the fact that several of these objections have force against the *being for* account of normative certitude too.

Gradability

Sepielli does not say much about the nature of the attitude of being for, but in order not to betray non-cognitivism he must at least maintain that it is a non-cognitive attitude. To accommodate degrees of normative certitude he must maintain that it is a gradable attitude. But this needs to be argued for since it is not obvious that all non-cognitive attitudes come in degrees. Many do, of course, such as desires or wishes, but, as pointed out above, more problematic cases are those of *intentions* or *plans*. It is debatable, of course, whether these are non-cognitive attitudes, but let us assume with expressivists such as Allan Gibbard that they are.¹⁸ It is far from obvious that it makes sense to talk about degrees of intentions or plans, or stronger or weaker

¹⁶ To be clear, it is, of course, not an uncontroversial matter how to understand degrees of belief. It is not even uncontroversial that belief does vary in degrees. All we mean to say here is that accounting for degrees of normative certitude is not a special problem for normative cognitivists, whereas it is for normative non-cognitivists. For a recent defense of degrees of belief, see John Eriksson and Alan Hájek, 'What are Degrees of Belief?', *Studia Logica*, vol. 86, no. 2 (July 2007), pp. 185–215.

¹⁷ Sepielli, 'Normative Uncertainty for Non-Cognitivists', p. 194.

¹⁸ Gibbard, *Thinking How to Live*.

intentions or plans. Even if the idea of degrees of being for does make sense, this needs to be argued for. But Sepielli offers no such arguments.

Cross-Attitudinal Comparisons

A plausible account of normative certitude should allow comparisons between normative and non-normative certitude. For instance, one can be more certain that $2 + 2 = 4$ than that it is right to maximize overall happiness. We think that the root of this problem is the lack of a natural maximum level for desire-like attitudes. But even on the assumption that Sepielli has solved the normalization problem for being for and shown that being for varies between 0 and 1, just like belief does, it is not clear that his account can make sense of comparisons between normative and non-normative certitude. Remember that on pain of betraying non-cognitivism, Sepielli must hold that being for is a non-cognitive attitude, i.e. a kind of attitude different from belief. What, then, does it mean to say that one's belief that $2 + 2 = 4$ is stronger than one's attitude of being for blaming actions that do not maximize overall happiness? That does not seem to be a meaningful statement. Even if both degrees of desires and degrees of beliefs are bounded (so that each attitude has a natural minimum degree 0 and a maximum degree 1), they seem too different to be meaningfully calibrated. But intuitively, we *can* make sense of comparisons between normative and non-normative certitude, and for cognitivists this is no problem at all.¹⁹

Motivational Maladies

Since being for is a non-cognitive attitude, it is sensitive to general changes in a subject's psychology. Falling into a state of depression or listlessness generally has a negative impact on motivational attitudes. It is likely, then, that becoming depressed or listless makes one less for blaming and praising for various actions. Suppose that as a result of falling into a state of general listlessness, a person becomes less for praising for charity work

¹⁹ There is a proposal that avoids this problem in Schroeder, *Being For*. Schroeder suggests a non-cognitivist view of belief, according to which belief, too, reduces to the attitude of being for. Roughly, believing that *P* is being for proceeding as if *P*. But as Schroeder notes, this moves has many costs and it is likely that many expressivists will be reluctant to make it.

than she used to be. On Sepielli's view, this means that the listless person's certitude that charity work is right has decreased. But this is an implausible implication. Falling into states of depression or listlessness does not entail that one is less certain about non-normative matters, such motivational maladies do not affect the certitude that $2 + 2 = 4$ or that the CIA planned the murder of JFK. So why should falling into states of depression or listlessness entail that one is less certain about normative matters? Of course, one's interest in being moral may wane when one falls into a depression, but moral interest and moral confidence are different things. It is noteworthy that Lenman's and Ridge's expressivist accounts do take this on board and thus are not vulnerable to the problem of motivational maladies. So this is an aspect in which Sepielli's account is in worse shape than previous accounts.

The Wrong Kind of Reasons

Another serious problem is that the attitude of being for and the degrees to which one is for bearing some relation, such as blaming, to some action can vary independently of moral certitude. For example, a utilitarian might be certain that murder is wrong but not be for blaming for murder since he thinks the attitude of blaming for murder is suboptimal in terms of overall happiness. Or to take another example, one might be very much for blaming omissions to keep one's promises not because one has a high degree of certitude that keeping promises is right but because an evil demon has threatened to torture one's family unless one is for blaming these omissions. That is, the being for account is vulnerable to a version of the notorious wrong kind of reason problem, which has been much discussed recently in other areas of normative theory. In this context, the problem is that one can be for bearing some relation to an action without this having any bearing on one's normative certitude; the reasons for being for bearing some relation (such as blaming) to an action are, in some cases, of the wrong kind to capture normative certitude. There is as yet no general solution to the wrong kind of reason problem that has won general acceptance and it is not easy to see what the solution would be in this particular context.²⁰

²⁰ For samples of the recent debate on the wrong kind of reason problem, see e.g. Wlodek Rabinowicz and Toni Rønnow-Rasmussen, "The Strike of the Demon: On Fitting Pro-Attitudes

IX. Normalization of *Being For*

In trying to solve the normalization problem for *Being For*, Sepielli begins by noting that the normalization axiom is often formulated as $\text{prob}(\Omega) = 1$, where 'prob' signifies subjective probability and ' Ω ' signifies a "universal set" whose members are all possible events.²¹ Informally, this means that the subjective probability that some event or other will occur is 1. Whatever degree of belief one has that some particular event will occur, this degree cannot be greater than the degree to which one believes that some event or other will occur. Sepielli suggests an analogous normalization axiom for the attitude of being for. For example, our certainty that *A* is wrong or *A* is not wrong should be understood as the degree of being for (blaming for *A* or not blaming for *A*) = 1.

But this normalization procedure has some awkward consequences. Being for is supposed to be a practical, action-guiding attitude:

[N]ormative thought is tied to action, in the broadest possible sense. When you are for something [...], then other things being equal, this is what you do. So understood [being for] is a motivating state and hence naturally understood as akin to desire, rather than belief.²²

What are you motivated to do when you are for (blaming for *A* or not blaming for *A*)? Sepielli might respond that you are motivated to do precisely that, i.e. to blame for *A* or not to blame for *A*. This might seem a curious motivating state, but he can insist that it is simply what it is to be certain that *A* is wrong or that *A* is not wrong.

But there is a more worrying implication. Whenever one is less than fully for blaming for *A*, one must be more for (blaming for *A* or not blaming for *A*) than for blaming for *A*, and thus, in some sense, prefer (blaming for *A* or not blaming for *A*) to blaming for *A*.

To revert to our earlier example, suppose that you have a high degree of certitude that keeping promises is right, but you are less than fully certain.

and Value', *Ethics*, vol. 114, no. 3 (April 2004), pp. 391–423; Sven Danielsson and Jonas Olson, 'Brentano and the Buck-Passers', *Mind*, vol. 116, no. 463 (July 2007), pp. 511–22; Gerard Lang, 'The Right Kind of Solution to the Wrong Kind of Reason Problem', *Utilitas*, vol. 20, no. 4 (December 2008), pp. 472–89; Mark Schroeder, 'Value and the Right Kind of Reason', *Oxford Studies in Metaethics*, vol. 5 (2010), pp. 25–55.

²¹ Sepielli, 'Normative Uncertainty for Non-Cognitivists', p. 202.

²² Schroeder, *Being For*, p. 84.

On the being for account, this means that you are strongly—but less than fully—for blaming for breaking promises. Since it seems pretty obvious that if one is more for *A* than for *B*, then one prefers *A* to *B*, Sepielli would then have to say that you must prefer (blaming for breaking promises or not blaming for breaking promises) to blaming for breaking promises, since you must be fully for (blaming for breaking promises or not blaming for breaking promises) and you are less than fully for blaming for breaking promises. But this seems very implausible. First, it is very difficult to even understand what it means to have this kind of preference, at least if preferences are supposed to be action-guiding. How can this preference ever guide one's actions? Second, even if we can somehow make sense of what it means to have this preference, it still seems absurd to have it. As an analogy, consider your attitude towards being happy. You are for being happy, which, plausibly, implies that you prefer being happy to not being happy. But you are not fully for being happy; perhaps you prefer freedom to happiness. Wouldn't it be absurd for you to prefer (being happy or not being happy) to being happy?²³

We conclude that the being for account of normative certitude is not promising. First, it is vulnerable to several objections we have leveled at previous accounts. Second, its solution to the normalization problem has implausible consequences.

Conclusion

The general conclusion of this discussion is that while cognitivism has an easy time making sense of moral uncertainty, non-cognitivism is still struggling to find a plausible account of moral certitude, which does not conflate certitude with importance or robustness. Lenman's ecumenical account cannot accommodate fundamental moral uncertainty, only moral uncertainty that depends on empirical uncertainty. Ridge's version avoids this problem but has instead the absurd implication that moral credence depends on the strength of non-moral desires. Sepielli's

²³ Another way to express the worry about this preference is to say that it violates a famous principle of preference logic, often called *disjunction interpolation*, which is very compelling, at least when it is applied to contradictory pairs of alternatives. If *X* is weakly preferred to not-*X*, then *X* is weakly preferred to (*X* or not-*X*) and (*X* or not-*X*) is weakly preferred to not-*X*. In other words, if *X* is weakly preferred to not-*X*, then (*X* or not-*X*) cannot be ranked above *X* or below not-*X*. For more on this principle, see Sven Ove Hansson, *The Structure of Values and Norms*, Cambridge: Cambridge University Press, 2001, sects 6.6, 7.7.

non-ecumenical account invokes a notion of being for a tautology, which has many problematic implications. Both Ridge and Sepielli have problems with cross-attitudinal comparisons. It is not at all clear that it makes sense to compare strengths of beliefs with strengths of desires (or combinations of desires and beliefs).

Of course, we do not pretend to have shown that there cannot be a plausible non-cognitivist account of fundamental moral uncertainty, but it is at least clear that there are considerable obstacles to overcome.