

The Philosophical Quarterly

Volume 59 • Number 235 • April 2009

CONTENTS

ARTICLES

Must There Be a Top Level? Expressivism and Moral Certitude	Einar Duenger Bohn Krister Bykvist and Jonas Olson	193 202
When is Parsimony a Virtue?	Michael Huemer	216
Luck and History-Sensitive Compatibilism	Neil Levy	237
A Science of Intention	Michael McDermott	252
On Essentially Conflicting Desires	Patricia Marino	274
Fixing Perceptual Belief	Gerald Vision	292

DISCUSSIONS

A Defence of the Explanatory Argument for Physicalism	Jared Bates	315
Extended Simples and Qualitative Heterogeneity	Kris McDaniel	325
Elephants, Microscopes and Free Beauty: Reply to Davies	Hans Maes	332

CRITICAL STUDIES

A Philosopher of Philosophy	P.M.S. Hacker	337
Velleman on Love and Ideals of Rational Humanity	Edward Harcourt	349

BOOK REVIEWS		357
--------------	--	-----

EXPRESSIVISM AND MORAL CERTITUDE

BY KRISTER BYKVIST AND JONAS OLSON

Michael Smith has recently argued that non-cognitivists are unable to accommodate crucial structural features of moral belief, and in particular that non-cognitivists have trouble accounting for subjects' certitude with respect to their moral beliefs. James Lenman and Michael Ridge have independently constructed 'ecumenical' versions of non-cognitivism, intended to block this objection. We argue that these responses do not work. If ecumenical non-cognitivism, a hybrid view which incorporates both non-cognitivist and cognitivist elements, fails to meet Smith's challenge, it is unlikely that 'purer' and more familiar versions of non-cognitivism will succeed.

I. INTRODUCTION

Michael Smith has recently argued that non-cognitivists are unable to accommodate crucial structural features of moral belief.¹ In particular, Smith argues that non-cognitivists have trouble accounting for subjects' certitude with respect to their moral beliefs. In response, James Lenman and Michael Ridge have independently elaborated 'ecumenical' versions of non-cognitivism, which are supposed to take care of Smith's challenge.² In this paper, we argue that the ecumenical non-cognitivists' responses are unsuccessful. This conclusion has significant repercussions. If ecumenical non-cognitivism, a hybrid view which incorporates both non-cognitivist and

¹ M. Smith, 'Evaluation, Uncertainty, and Motivation', *Ethical Theory and Moral Practice*, 5 (2002), pp. 305–20, repr. in his *Ethics and the A Priori* (Cambridge UP, 2004). Page references in the text refer to original publication. The argument was broached a decade earlier in a text book by Lars Bergström: see Bergström, *Grundbok i värde teori* (Stockholm: Thales, 1990), pp. 35ff.

² J. Lenman, 'Non-Cognitivism and the Dimensions of Evaluative Judgement', *Brown Electronic Article Review Service*, <http://www.brown.edu/Departments/Philosophy/bears/homepage.html>, posted 15 March 2003; M. Ridge, 'Ecumenical Expressivism: the Best of Both Worlds?', in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, Vol. II (Oxford UP, 2007), pp. 51–77. A note on terminology: Ridge talks about ecumenical *expressivism* rather than *non-cognitivism*. Lenman calls his favoured position 'impure non-cognitivism', but it fits the general characterization of ecumenical expressivism to be given in §III below. However, Lenman does not accept every detail of Ridge's specific version of ecumenical expressivism, to be outlined in §V below.

cognitivist elements, fails to meet Smith's challenge, it is unlikely that 'purer' and more familiar versions of non-cognitivism will succeed.

We shall proceed as follows. In §II we explain why Smith's argument is a challenge for non-ecumenical versions of non-cognitivism. In §III we explain how ecumenical non-cognitivism is supposed to solve the problem. We consider a simple version of ecumenical non-cognitivism and reject it on the ground that it fails to account for moral as opposed to merely empirical certitude. In §IV we discuss Lenman's 'impure non-cognitivism' and conclude that it suffers from a similar shortcoming. In §V we explain Ridge's attempt to avoid the problem which faces the simple version of ecumenical non-cognitivism and Lenman's impure non-cognitivism. §VI surveys some initial difficulties with Ridge's representation of moral certitude. In §VII we develop the criticism further, and we finish by posing a dilemma for ecumenical expressivism.

II. THE PROBLEM FOR NON-ECUMENICAL NON-COGNITIVISM

According to non-ecumenical versions of non-cognitivism, to make a moral judgement is to express a non-cognitive attitude, such as desire, rather than a cognitive attitude or belief. Still, everyday moral discourse suggests that moral judgements are in several respects akin to ordinary beliefs: we talk about moral beliefs (as in 'George believes that euthanasia is wrong'), about moral beliefs being more or less epistemically justified (as in 'George has no good grounds for his belief that euthanasia is wrong'), and about moral knowledge (as in 'George knows that murder is wrong'). To be sure, moral judgements also display 'desire-like' features. Most notably, people tend to be motivated to act in accordance with the moral judgements they endorse. This is what prompts Smith to say that '[moral] judgement has a decidedly Janus-faced character' (p. 305). Smith's challenge is a serious one for *quasi*-realists like Blackburn and Gibbard, who claim to be able to say everything the realist says without being committed to fully fledged realism.³ As we shall show, it is undeniable that what Smith calls 'certitude' and 'importance' are common sense features of moral belief, and if *quasi*-realists have no way to account for them there will be a significant crack in their fortress.

The gist of Smith's argument is that desires have too little structure to accommodate all the structural belief-like features which common sense finds in moral judgements. Smith focuses on the features he calls *certitude*, *robustness* and *importance*. Certitude is the degree of certainty with which a

³ See Blackburn, *Ruling Passions* (Oxford: Clarendon Press, 1998); Gibbard, *Thinking How to Live* (Harvard UP, 2003).

subject believes some proposition p . Robustness concerns how sensitive the belief that p is to new information; the more prone believers are to reassess their belief that p in the light of new information, the less robust their belief that p is. Whereas certitude and robustness are features of both moral and non-moral beliefs, importance is a characteristic of moral beliefs only. More accurately, importance is a feature that the believer ascribes to the action or state of affairs in question. To illustrate, most people believe both that one ought to save innocent lives if one can, and that one ought not to break promises. Most people also believe that if one can save an innocent life by breaking a promise, or by breaking several promises, one ought to do so. This means that most people believe that the duty to save innocent lives if one can is of greater importance than the duty not to break promises. Similarly, many people believe that both pleasure and autonomy have intrinsic value but that autonomy has greater intrinsic value than pleasure.

Robustness might not be a worry for non-cognitivists, since it is uncontroversial that non-cognitive attitudes, like cognitive attitudes, can be more or less sensitive to new information and reflection: just as we can be more or less certain that p in the light of new information and reflection, we can come to desire that p with greater or weaker strength in the light of new information and reflection. The problem for non-cognitivists is to account for both certitude and importance and to do so correctly. In particular, certitude and importance must be allowed to vary independently in the sense that one can believe with a high degree of certitude that some consideration is of moderate moral importance as well as one can believe with a low degree of certitude that some consideration is of great moral importance. To borrow an example from Smith (p. 308), one might believe that leading an autonomous life is of great intrinsic value, but one's certitude in holding this belief might be low. Someone might also believe with a great degree of certitude that experiencing pleasure is of moderate intrinsic value. Such a believer invests 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 pleasure is of moderate intrinsic value.

It might seem that the obvious way for non-cognitivists to account for importance is in terms of strength of desire (or whatever the relevant non-cognitive attitude expressed in moral judgements is). But then it is not clear how to account for certitude. Needless to say, if certitude is accounted for in terms of degree of strength of desire, importance cannot be accounted for in the same manner, since, as we have just remarked, certitude and importance can vary independently. As Smith puts it (p. 317),

The problem that looms, then, is that non-cognitivists seem not to have the resources to accommodate all three features of evaluative judgement that we commonsensically

ascribe to them. They can accommodate either importance and robustness, or certitude and robustness, but not all of importance, certitude, and robustness.

We may note at this point that Smith's challenge spells trouble for theories which associate moral judgements with intention-like states, such as plans. Since it is unclear whether intentions or plans come in degrees, it seems *prima facie* difficult to account for certitude and importance. The challenge is thus relevant to Gibbard's recent account in *Thinking How to Live*, according to which to judge that one ought all things considered to do something is to plan to do it. Smith's challenge is thus relevant to a wide family of expressivist views.

A possible non-cognitivist rejoinder, which Smith considers, is to represent importance in terms of first-order attitudes and certitude in terms of second-order attitudes. Such a sophisticated (but still non-ecumenical) 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 I judge that I have reason to keep my promises, I express a desire to desire that I keep my promises. The strength of the first-order desire could be taken to represent importance. So the stronger the first-order desire to keep my promises which I desire to have, the more important I consider it to keep promises. The strength of the second-order desire, on the other hand, could represent my degree of certitude that keeping my promises is right. So the more strongly I desire to desire to keep my promises, the more certain I am that it would be right for me to keep my promises. This proposal, which we shall call the *attitudinal ladder proposal*, faces several objections.

First, the attitudinal ladder proposal is plagued by arbitrariness (see Smith, p. 318): 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 round? Secondly, unlike degree of belief, desire-strength does not come in a neat interval with a clearly defined minimum and maximum. Certitude can vary from 0 (complete uncertainty) to 1 (complete certainty), 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 argue, this problem recurs for the ecumenical expressivist accounts 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 more strongly I desire to desire to ϕ , the more certain I am that I have reason to ϕ . But it seems possible that there are cases in which I desire to desire to ϕ without being at all certain that I

have reason to ϕ . Suppose an evil Kantian 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 have reason actually to keep all your promises.

III. ECUMENICAL NON-COGNITIVISM TO THE RESCUE?

It has recently become a popular view that the best form of non-cognitivism is *ecumenical*. As the label suggests, ecumenical non-cognitivism is a hybrid view which 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.⁴ 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 for moral sentences.⁵

More specifically, the idea is that a moral judgement concerning, e.g., the rightness of an action expresses (i) a general desire for actions in so far as they have a certain property, and (ii) 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 (Ridge, p. 55).

According to the simplest version of ecumenical non-cognitivism, the property in question is provided by the first-order moral view endorsed by the speaker. For example, a utilitarian who endorses the judgement that sticking to a vegetarian diet is morally right expresses approval of actions in so far as they tend to maximize overall well-being, and a belief that sticking to a vegetarian diet has *that property*, i.e., the property of tending to maximize overall well-being. 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: it is not the case that ‘Sticking to a vegetarian diet is morally right’ is true iff the belief that sticking to a vegetarian diet tends to maximize overall well-being is true.

⁴ As Lenman has reminded us, the idea that moral judgements have both non-cognitive and cognitive meaning is not new. R.M. Hare argued that the primary meaning of moral judgements is prescriptive, while their secondary meaning is descriptive. See Hare, *The Language of Morals* (Oxford: Clarendon Press, 1952), ch. 7. The similarities between Ridge’s ecumenical expressivism and Hare’s prescriptivism are discussed in Eriksson, *Moved by Morality* (Uppsala Philosophy Department, 2006), pp. 199–204.

⁵ Ridge, ‘Ecumenical Expressivism’, p. 54; Lenman, ‘Non-Cognitivism’, §2.

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 component* expressed in a moral judgement. So my certitude that sticking to a vegetarian diet is morally right is identified with the degree of my belief that sticking to a vegetarian diet tends to maximize overall well-being. Importance could then be represented by the strength of the desire component. Given that I believe that sticking to a vegetarian diet tends to maximize overall well-being, the more strongly I approve of actions in so far as they tend to maximize overall well-being, the more moral reason I think I have 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 a genuine belief. Another advantage is that it can allow for some cases of motivational maladies. I can judge that I ought to do something and yet still lack a desire to do it. It is true that my judgement will always express a desire, but it is a *general* desire to do actions in so far as they have a certain property, not a desire to do a particular action. Arguably, a general desire for actions can exist in the absence of a desire to do a particular action. I might, for instance, have a general desire to do some work today, but, in a state of listlessness, have no desire to do some particular kind of work.

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

Lenman and Ridge, who both seem sensitive to this objection, advocate a slightly more complex version of ecumenical non-cognitivism. Following Ridge we shall call it the 'ideal advisor' version. According to this version, a moral judgement concerning the rightness of actions expresses (1) a desire for actions in so far as they would be approved of by a certain sort of *ideal advisor*, and (2) a belief which makes direct reference back to the property of being approved of by that advisor.⁶ To judge that sticking to a vegetarian diet is morally right is on this view to express (1) a desire for actions in so far

⁶ Lenman, 'Non-Cognitivism', §§2, 4; Ridge, 'Ecumenical Expressivism', p. 57.

as they would be approved of by one's ideal advisor and (2) a 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 well-being, 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 spelt 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 I take there to be for me to keep my promises is represented by how strongly I think my ideal advisor would want me to keep my promises. This differs from the simpler view sketched above, according to which importance was identified with the motivational strength of the actual agent's desires. Lenman's and Ridge's view is thus an improvement on the simpler view, since it is possible to judge actions to be very important even though one's actual general motivation to act is weak (due to depression and the like). Just as depression and listlessness can sap my desire to do a particular action, they can sap my general desires to perform actions with certain properties. (It should be noted, however, that both the simple and the complex accounts rule out cases of motivational maladies where I judge that I ought to do some action but lack a general desire to perform actions in so far as they have a certain property.) While Lenman's and Ridge's treatments of importance are the same, they differ in their treatments of certitude. In the remainder of this paper, we shall focus on their accounts of certitude.

IV. THE SUPPOSED SOLUTION: LENMAN'S VERSION

Lenman's response to Smith 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* is wondering whether an *improved* version of himself would disapprove of ϕ ing. This story is consistent with non-cognitivism as long as 'improved' is given a non-cognitivist gloss, i.e., as long as the improved version of *S* is described in purely descriptive terms and 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; see the discussion of the difference between ecumenical cognitivism

and ecumenical non-cognitivism in the previous section). 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 he ought not to lie expresses a positive attitude to conforming with the desires of a *VS* (where '*V*' is a descriptive specification of a hypothetical version of *S* whom actual *S* endorses as improved, in the sense that actual *S* takes the desires of *VS* to be action-guiding), together with a belief that *VS* would desire that *S* does not lie.⁷ As Lenman stresses, the non-cognitivist should not insist on making any particular *V* a part of the analysis of moral concepts, since there is room for substantive debate about which patterns of motivation actual agents ought to endorse as action-guiding (§5). For illustrative purposes, we shall follow Lenman and say that a *VS* is simply a fully informed and clever version of *S*, where 'fully informed' and 'clever' are understood descriptively.

S's certitude that he ought not to lie is represented by the strength of his belief that *VS* 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 his moral judgements, and the strength of *S*'s general desire to act in accordance with what *VS* would desire that *S* does represents how much *S* cares (*de dicto*) about acting in accordance with his moral judgements. Since non-cognitivists are committed to motivational internalism, they must hold that any agent who makes a moral judgement has some desire to conform to the desires of his improved self. But since this desire need not be very strong, there is room for *akrasia*.

This last point is a virtue of Lenman's account. But as we shall now explain, the account leaves out an important dimension of moral uncertainty: it has no resources to represent uncertainty regarding whether being (dis-)approved by a certain descriptively specified improved agent is indicative of rightness (wrongness). This means that Lenman's account is unable to account for fundamental substantive moral uncertainty. To illustrate, suppose some person *P* believes that an improved version of himself ('*KP*') would display a Kantian pattern of motivation. Suppose further that *P* is perfectly certain that *KP* would desire that *P* should perform some action ψ . Suppose finally that *P*'s desire to ψ and his general desire to conform to *KP*'s desires about *P*'s actions are quite weak. On Lenman's account, this means that *P* is *akratic* or that he is not very concerned about acting in accordance with his moral judgements.

But there is another possible explanation of the weakness of *P*'s desire to ϕ and his general desire to conform to *KP*'s desires about *P*'s actions. *P* might believe that there is another improved version ('*UP*') of *P* who would

⁷ Lenman notes that Simon Blackburn has suggested a similar account. See Blackburn, *Ruling Passions*, pp. 261–9.

display a utilitarian pattern of motivation. Suppose that P is perfectly certain that UP would desire that P should *not* ϕ . Suppose also that P 's certitude that anything approved by UP is right is greater than his certitude that anything approved by KP is right. In other words, P has some degree of belief that the recommendations of a utilitarian ideal agent would indicate rightness and wrongness and some degree of belief that the recommendations of a Kantian ideal agent would indicate rightness and wrongness, the former degree of belief being greater than the latter. This difference in moral certitude is impossible to represent if P 's certitude that he ought (not) to ψ simply is his certitude that an improved version of P would desire that actual P should (not) ψ . (P 's certitude that KP would desire that P should ψ is equally as great as P 's certitude that UP would desire that P should not ψ .)

Lenman anticipates this difficulty, and acknowledges that a subject S can endorse several different (descriptively specified) improved versions V_1, \dots, V_n of himself. He suggests that in such cases, S 's certitude that the desires of some particular VS are morally authoritative is determined by S 's certainty that the desires of this particular VS cohere with the desire of a WS , where ' W ' is a yet another descriptive specification of an improved S . Lenman suggests that 'for a religious believer, V might be "fully informed" and W might be "fully conforming my will to God's commands"'.⁸

But this move does not help to represent fundamental substantive moral uncertainty, i.e., a subject's uncertainty as to whether anything that would be (dis)approved by a certain improved version of himself would be right (wrong), irrespective of whether this improved version displays a utilitarian or Kantian pattern of motivation, or conforms his motivational patterns fully to the commands of God, or whatever. 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 that this improved agent would approve of 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 that faces the simpler version discussed in the previous section. As we shall show in the next section, Ridge is aware of this problem and proposes a way to deal with it.

⁸ Lenman, 'Non-Cognitivism', §7. Cf. the discussions in Ridge, 'Certitude, Robustness, and Importance for Non-Cognitivists', *Brown Electronic Article Review Service*, <http://www.brown.edu/Departments/Philosophy/bears/homepage.html>, posted 15 March 2003. See also the discussion below.

V. RIDGE'S VERSION

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 (p. 71),

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 in so far as they would be approved of by the relevant sort of advisor.

The discussion in the previous section explains why something like clause (b) is called for. We pointed out 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 well-being 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 (cf. Ridge, pp. 71ff.). 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 ultimately indefensible.

But first, we should take notice of yet another clarification which Ridge makes and which will be relevant to the dilemma we shall pose for ecumenical expressivism in §VII 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 his moral judgements' (Ridge, p. 72). 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 his desire in (b) might stay the same. (In what follows, we use the terms 'desire' and 'pro-attitude' interchangeably.)

In the remaining two sections we shall argue that Ridge's treatment of certitude faces a host of problems. §VI briefly surveys some initial problems and §VII presents reasons for rejecting the ecumenical expressivist's account of certitude.

VI. INITIAL PROBLEMS AND CROSS-ATTITUDINAL COMPARISONS

According to Ridge's ideal advisor version of ecumenical expressivism, a subject S 's certitude that some action is morally right is a function of (a) the degree of S 's belief that this action would be approved of by the relevant sort of advisor, and (b) the relative strength of S 's desire for actions in so far 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 in §II 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 shall call *cross-attitudinal comparisons*. One 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 in so far as they maximize utility and the relative strength of S 's pro-attitude to actions in so far as they would be favoured by an ideal advisor, we need to be able to make comparisons in strength between on the one hand beliefs (such as that $2 + 2 = 4$) and on the other hand combinations of beliefs and desires. 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 argued, it is far from clear how to introduce such a scale.

VII. COMPENSATIONS AND A DILEMMA

Since moral certitude is a function of a belief and a desire, it seems that increases in the one component can 'compensate' for decreases in the other. This has odd implications. To give an example, suppose that at some time t_1 S believes that an ideal advisor would favour actions in so far as they maximize utility and that S has a general pro-attitude to actions in so far as

they would be favoured by this ideal advisor. Suppose next that at some later time t_2 the degree of S 's belief that the ideal advisor would favour actions in so far as they maximize utility decreases (but does not reach zero), while the relative strength of S 's pro-attitude to actions in so far as they would be favoured by the ideal advisor increases significantly. (Perhaps this happens because, fantastically, S meets the ideal advisor and is deeply impressed by his magnificent appearance and is also told by a reliable source, perhaps the ideal advisor himself, that it is not very likely that the ideal advisor would favour actions in so far as they maximize utility.)

If the increase in relative strength of S 's pro-attitude to actions in so far as they would be favoured by the ideal advisor is significant enough, S 's certitude that actions are right in so far as they maximize utility is greater at t_2 than at t_1 even though S 's belief that the ideal advisor would favour actions in so far as they maximize utility is (possibly much) weaker at t_2 than at t_1 . But this seems odd. Why should S 's certitude that actions are right in so far as they maximize utility increase just because the relative strength of S 's pro-attitude to actions in so far as they would be favoured by an ideal advisor increases?

In correspondence, Ridge has suggested that this objection begs the question. The ecumenical expressivist might bite the bullet and accept that it follows from his account that there are cases in which S 's certitude that some action is right increases even though the degree of S 's belief that this action would be approved of by the relevant ideal advisor decreases significantly. But the objection does not beg any question, since the point of it is simply to tease out a peculiar implication of ecumenical expressivism. It is undeniably a peculiar idea that S 's certitude that an action is right might increase just because the relative strength of S 's general pro-attitude to actions in so far as they would be approved of by the ideal advisor increases significantly, even though S 's certainty that the ideal advisor would approve of the particular action in question decreases significantly.

Perhaps the ecumenical expressivist can respond by modifying the function that represents certitude. To avoid the upshot that decreases in belief can be compensated by increases in relative strength of pro-attitude the belief component can be accorded more weight than the desire component. To rule out the possibility that massive increases in relative strength of pro-attitude can outweigh decreases in belief, the belief component can be accorded lexical weight to the effect that no increase in the desire component can compensate a decrease in the belief component.

But this manoeuvre seems *ad hoc*. The desire component was invoked in the representation of moral certitude in order to enable the expressivist to account not only for empirical uncertainty as to what some descriptively

specified ideal agent would (dis)favour but also for fundamental substantive moral uncertainty as to whether being (dis)favoured by such an ideal agent really is indicative of rightness (wrongness). It is difficult to see why the tiniest increase in *empirical* belief should entail greater *moral* certitude than any increase in relative desire-strength, however vast. Such a manoeuvre seems to downplay the significance of relative desire-strength in a way that does not sit well with the rationale for invoking it as a component of representations of moral certitude in the first place. The lesson to be drawn is that the attempt to represent moral certitude partly or wholly in terms of relative desire-strength is wrongheaded.

The final and most important objection to ecumenical non-cognitivism is that it faces a serious dilemma. What generates the first horn is that Ridge defines moral certitude partly in terms of relative desire-strength. The relative strength of *S*'s desire *d* is most naturally defined, by 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, 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 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 in so far as your ideal advisor would approve of them remains the same, which is likely, 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 my desire for eating ice-cream becomes stronger and the strengths of my other desires stay the same, my moral certitude has decreased.

It is equally obvious that examples can be given which work in the opposite direction. For instance, if I fall out of love with a person and the strengths of my desires concerning this person consequently diminish, or if the strength of my desire for ice-cream diminishes, the relative strength of my desire for actions in so far as my ideal advisor would approve of them increases. On Ridge's view, this means that my moral certitude increases.

But it is implausible that my moral certitude is determined in this direct manner on my falling out of love or on my 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 wipe out one's moral certitude. For Ridge, moral certitude depends in part on the strength of one's general desire to perform actions in so far as they would be approved by the ideal advisor. But as we pointed out in §III above, just as depression and listlessness can sap my desire to do a particular action they can sap my general desires to perform actions with certain properties.

VIII. CONCLUSION

We have argued that extant responses to Smith's argument that expressivists cannot account for certitude and importance with respect to moral judgements fail. Lenman's proposal can merely account for empirical uncertainty about what some non-morally specified ideal agent would (dis-)favour; it cannot account for fundamental substantive moral uncertainty. Ridge's ecumenical expressivist proposal is more complex, but fares no better. It is unclear how ecumenical expressivism is supposed to represent moral certitude and it can make little sense of what we called cross-attitudinal comparisons. More importantly still, the possibilities of increases and decreases in relative desire-strengths have absurd implications and ultimately confront the ecumenical expressivist with a dilemma: either increases in desires for morally irrelevant objects undermine moral certitude, or general motivational maladies undermine moral certitude.

We conclude that ecumenical versions of expressivism are unable to account for certitude and importance with respect to moral judgements. Since it is unlikely that non-ecumenical versions will succeed where ecumenical versions fail, this conclusion deals a serious blow to expressivism.⁹

Jesus College & Brasenose College, Oxford

⁹ Earlier versions have been presented at seminars at the Universities of Leeds and Oxford; at the Centre de Recherche en Ethique de l'Université de Montréal; at a Cumberland Lodge Weekend, organized by LSE, November 2007; and at the British Society for Ethical Theory conference in Edinburgh, July 2008. We are grateful to the participants of these meetings for helpful discussions. Thanks to Sven Danielsson for extensive correspondence and to Jimmy Lenman and Michael Ridge for their helpful written comments.