

Forthcoming in Journal of Moral Philosophy

Enoch's Defense of Robust Meta-Ethical Realism

Gunnar Björnsson

Ragnar Francén Olinder

David Enoch's *Taking Morality Seriously* (Oxford University Press 2011) is the latest in a recent flurry of defenses of meta-ethical (and more broadly, meta-normative) non-naturalist realism. The position defended, which Enoch dubs "Robust Realism" (henceforth RR) says that there are *objective* and *absolute irreducibly normative* and more specifically *moral* truths and facts, which our normative and moral judgments represent. Throughout the book Enoch displays an impressive philosophical breadth and creativity, making a number of novel and intriguing contributions to the field. Many of these are familiar from Enoch's earlier publications, but the book ties together the various strands of his thinking in an admirably clear way.

The positive part of Enoch's case for RR consists of two original arguments, presented in chapters 2 and 3 respectively, and the "auxiliary chapters" 4 and 5. The fact that Enoch presents two novel positive arguments for a contested philosophical view is itself worthy of note, as is the ambition of the argument: to spell out ways in which Robust Realism is the only view that can take morality seriously.

In chapters 6 through chapter 9, Enoch plays defense, addressing the main objections to non-naturalist realism: metaphysical concerns about the queerness and supervenience (on natural facts) of non-natural normative facts; epistemic worries about

how we can access or gain knowledge about normative facts if they are non-causal; challenges from moral disagreement; concerns that realist views cannot account for the practical, normative or motivational aspects of morality. Enoch's general strategy in the defensive parts is to distinguish and clarify different versions of the objections in question, and to argue that each version either misses its target, fails to present a viable objection, or merely minimally diminishes the plausibility of RR. Enoch does an impressive job of dissecting the objections, and this alone is a valuable contribution to the literature. At the same time, the approach has its limits. Ambitious scope sometimes brings a lack of in-depth discussion and we suspect that Enoch sometimes fails to identify or take sufficiently seriously the strongest versions of the objections (for some examples, see our discussion of chapters 7 and 8 below). It should be noted, though, that Enoch is often refreshingly explicit about (what he considers to be) the limitations of his arguments.

As is customary, our review will focus on what we take to be particularly problematic central passages of the book, specifically the case made for RR in chapter 2 and 3 and the challenges from epistemology and disagreement in chapter 7 and 8.

We start with chapter 2. Non-objectivist metaethical views (including non-cognitivist, subjectivist, relativist and response-dependence views) take a person's moral opinions to consist in or depend on her responses in terms of preferences or pro- and con-attitudes. The argument in chapter 2 aims to show that such theories have implausible moral implications regarding conflicts where two persons disagree about how to act. Suppose that A and B have decided to spend the evening together but prefer different activities. A prefers playing tennis, B going to the movies. Suppose further that this

conflict is completely preference-based – it is due solely to the fact that they simply want different things. In such conflicts, Enoch contends, IMPARTIALITY holds: some impartial solution is morally called for, such as flipping a coin, or taking turns (letting the other choose next time); insisting on one's own preferred alternative is morally wrong. It is to act as if the other person's interests are not equally morally important.

How is this relevant to metaethics? On non-objectivist metaethical views, moral conflicts about how to act are preference-based. Thus, IMPARTIALITY should apply to moral conflicts. But, Enoch argues, this is implausible. Suppose that C and D have decided to perform medical experiments on animals and that there are two methods: one avoids animal suffering, the other inflicts serious suffering but is slightly cheaper. C advocates the first method since animal suffering is morally bad, while D advocates the second method on the ground that animal suffering is morally irrelevant. Given that animal suffering is bad, C and D ought not flip a coin and it seems at least permissible for C to stand her ground and try to convince D. So non-objectivist metaethical views have the wrong moral implication here. On the other hand, metaethical views according to which moral judgments are beliefs about objective facts get things right. For IMPARTIALITY is typically not plausible for fact-oriented conflicts. If you and I are to defuse a bomb and you correctly think that cutting the red wire will do the job, while I wrongly think that cutting the blue wire will do the job (when it will actually result in a lethal explosion), it seems not only permissible but perhaps even obligatory for you to insist that we do it your way.

A first worry about this argument is that its first-order moral premise is not defended in a thorough (or any) normative ethical argumentation. Let us grant, however, that all Enoch needs as a premise for his argument (or an argument in the vicinity) is

that IMPARTIALITY is pre-theoretically plausible for preference-based but not for moral conflicts. It might then be held that non-objectivist metaethical views, in contrast to objectivist views, cannot explain this intuitive difference.

Our main worry, however, is that such explanations seem available.¹ Start with the following seemingly tautological idea: to think that an action is morally right is to think that it takes proper account of the interests and preferences of everyone concerned. Given this idea, we have a straightforward explanation of why IMPARTIALITY is plausible for mere clashes in equally strong preferences but implausible for clashes in preference based on conflicting moral judgments. In the case of mere preferences, for A to stand her ground would be to treat her own preferences as more important on the mere ground that they are her own. This, we commonly think, is not to take proper account of the preferences of those concerned. In the case of preferences grounded in moral judgments, however, thinking that C is correct in her moral judgment already involves thinking that acting accordingly *does* take the preferences involved into proper account. But if this is what one thinks, then it is natural to think that C's action of standing her ground might also be right, taking into proper account the preferences involved (assuming, say, that C's standing her ground might result in C and D acting on C's judgment while having no other significant effects and involving no greater degree of disrespect than would their acting on D's judgment).

This explanation of the contrast strikes us as eminently plausible, and compatible with a variety of substantial ethical theories. But it also seems perfectly compatible with sophisticated non-objectivist theories. For example, it seems compatible with non-

¹ Another potentially major concern is whether the moral asymmetry between preference-based and belief-based conflicts is sufficiently clear-cut. See Kate Manne and David Sobel's "Disagreeing about how to disagree" (ms) for a nice critical discussion.

cognitivist theories. Any such theory will have to say something about how the attitudes constituting moral judgments differ from non-moral preferences, but there are various possibilities here. One is that moral attitudes are attitudes the characteristic function of which is to facilitate cooperation by promoting or preventing kinds of action of the moral judge and by expressions of approval and disapproval of corresponding actions by others.² Such attitudes, one might expect, will involve sensitivity to considerations that tend to be important for cooperation: considerations of benefit, harm and respect, the adherence to or violation of expectations on which cooperation is based, and considerations of reciprocity. More specifically, if moral attitudes have this function, we should expect moral judgments to have the features needed for the above explanation of the contrast between preferences based on moral judgments and “mere” preferences. First, since it generally makes for better cooperation if we do not simply insist on our own preferences but instead try to seek a way to decide that everyone can accept, we can expect the moral attitude expressed by IMPARTIALITY for ordinary clashes of preferences. Second, however, since one crucial part of the function of moral attitudes and corresponding preferences is to guide us when there are clashes of preferences and interests, our moral attitudes and preferences must take into account such preferences and interests but cannot themselves be lightly treated as one preference among others.

This non-cognitivist explanation of the contrast strikes us as highly plausible on the assumption that a general functionalistic account of moral attitudes is otherwise feasible. Moreover, it can be adopted by speaker-relativists or contextualists. Suppose for example that one takes judgments of moral rightness (or wrongness) to say that

² See Björnsson, Gunnar and McPherson, Tristram, “Moral Attitudes for Non-Cognitivists: Solving the Specification Problem,” *Mind* (Forthcoming/2013). Cf., Gibbard, Allan, *Wise choices, apt feelings: a theory of normative judgement* (Oxford: Clarendon Press, 1990), and Mackie, John L., *Ethics: Inventing Right and Wrong* (London: Penguin Books, 1977), ch. 5.

actions conform (or do not conform) to basic standards for social coordination and collaboration that the speaker endorses or that are salient in moral contexts. Then the functionalist story above can explain why these standards will have the shape required by the explanation, calling for impartial solutions in ordinary clashes of preferences but resisting such solutions to some degree when moral preferences clash.

One might of course resist this explanation for various reasons, but at this stage it strikes us as plausible, and in line with plausible explanations of the shape of our moral intuitions.³ We thus remain unconvinced by Enoch's argument from IMPARTIALITY in chapter 2.

Chapter 3 starts in Harman's challenge to non-reductionist moral realists: we have reason to believe in the existence of some unobservable phenomenon if, and only if, it is needed in the best explanation of some observation; moral facts are not needed in the best explanation of our (putative) moral (or other) observations; so we have no reason to think they exist. The realist reply has commonly been to either (i) deny that moral facts are explanatorily dispensable or (ii) argue that explanatory dispensability doesn't matter. Enoch proposes a new and interesting version of the second reply: *explanatory* indispensability is not required since the assumption of normative facts is indispensable for practical deliberation. When we decide how to act based on deliberation, i.e., asking ourselves which course of action makes most sense, we implicitly commit ourselves to the existence of normative reasons for our way of acting. This is what differentiates *deciding based on deliberation* from mere *picking*. Thus, in deliberating we commit

³ It is unclear whether Enoch can resist functional explanations of our moral intuitions in general, as his explanation in chapter 7 of why these intuitions tend to be true appeals to the evolutionary benefit of believing that (means to) survival and reproduction are good.

ourselves to the existence of normative facts. That normative facts are indispensable for deliberation in this sense justifies belief in their existence.

If this argument is successful, it both answers the explanatory challenge and provides a positive argument for the existence of realist normative facts (sect 3.7). But there are (at least) two fundamental worries here. First, does deliberation really involve a commitment to realist normative facts? Second, even if deliberation commits us to realist normative facts, how does that provide reason to think that there are such facts? After all, it is not a given that the world lives up to any old commitment that we might have. We will pursue the first type of worry, but Enoch's attempts to address the second are relevant to understanding his view.

To address the worry that the world does not in general live up to assumptions to which we commit ourselves, Enoch employs a strategy developed together with Josh Schechter. He introduces the idea that some belief-forming methods are basic, in the sense that we cannot give an epistemic justification for using them. He then notices that we nevertheless seem epistemically justified in employing some such methods: methods like modus ponens, perception, memory, and inference to the best explanation. So what distinguishes these methods from unjustified basic methods, such as inference to the worst explanation? This, Enoch suggests, is where indispensability matters: what vindicates our reliance on a basic belief-forming method is that it is indispensable for a rationally non-optional project – i.e., one from which it is not a rationally acceptable option for us to disengage. For example, the explanatory project is rationally non-optional and to pursue the explanatory project we need to assume the (likely) existence of the facts that are indispensable to our best explanations. Enoch's point, now, is that this idea makes arguments from deliberative indispensability just as respectable: the

deliberative project is also rationally non-optional for us, and we commit ourselves to the existence of irreducibly normative facts in pursuing the deliberative project.

One might of course have general worries about arguments from indispensability, and specific worries about Enoch's defense of such arguments. For Enoch's defense of RR, however, this matters little unless deliberation does commit us to irreducible normative facts, or indeed to any kind of normative fact. Unfortunately, what Enoch has to say here is sparse and highly impressionistic, and thus unconvincing for those of us who do not share his impressions. Much of the work is done by a comparison between deliberation and mere picking, and Enoch rightly points out that in the latter case, unlike the former, one is seeking (perhaps vague) answers to one's deliberative questions (72). What one is committed to in doing so, it seems, is the assumption that there are correct or incorrect (or perhaps better or worse) answers to one's questions, and that one cannot just decide that one answer is correct or best. (We take this to be what Enoch has in mind when he says that we are "trying to discover, not create" answers when we deliberate (73).)

This seems right to us, but it is as yet unclear how this commits one to the existence of normative facts in any substantial sense. That one is committed to finding answers satisfying some constraints that are independent of one's current deliberation does not in itself mean that these constraints are inherently normative: it just means that one is treating them as constraints for one's deliberation, i.e. that one is disposed to adjust the deliberation when it seems to violate them. But Enoch has a couple of further things to say that might be thought to provide further substance.

First, he says that because "only normative truths can answer the normative questions I ask myself in deliberation, nothing less than a normative truth suffices for

deliberation (80)". It is not entirely clear to us, however, why it is rationally non-optional to ask questions in deliberation that requires answers in *normative* terms. Why can't one rationally chose to instead ask questions about what options satisfy some non-normative condition that one treats as a constraint on one's deliberation?

Second, even if it is right that only normative truths can answer the questions asked in deliberation (truths about what actions are best supported by reasons, say), this might say precious little absent an account of what it is to take something to be a normative truth: as far as this argument is concerned, it is compatible with minimalist conceptions of truth, and with various forms of judge-relativist, subjectivist, or non-cognitivist accounts of the content of moral or normative claims. This, we should stress, is something that Enoch recognizes, and in response to the possibility of non-cognitivist or subjectivist interpretations of deliberation he claims that "the deliberative project loses much of its initial appeal ... once normativity is viewed as dependent on our attitudes (81)". If the appeal of engaging in the project were essentially tied to its connection to judge-independent constraints, and if we have to think of these constraints as normative truths, it would be an indispensable assumption of deliberation that it is constrained by normative facts holding independently of our concern with those facts.

Our main worry about this argument is that non-cognitivist, expressivist or speaker-relativist accounts of deliberation seem able to provide the relevant attitude independence. As long as the conditions one treats as constraints on one's deliberation do not themselves involve the condition that one treats them as constraints (as in e.g. *choose so as to maximize total happiness as long as you care about everyone's happiness*), the constraints are not attitude dependent.⁴ What is necessarily dependent

⁴ Cf. Blackburn, Simon: *Ruling Passions* (Oxford: Oxford U. P., 1998), ch. 9.

on one's attitudes is whether one *treats* something as a constraint on deliberation: to treat something as a constraint for deliberation is in itself to have a certain attitude towards it. But this is true whatever the character of the constraints in question: whether they are irreducible non-natural conditions or subjectivist naturalistic conditions.

We recognize, of course, that some find the appeal of deliberation undermined by the thought that the constraints of deliberation are ultimately constraints *endorsed by us*, in some sense that can only be avoided within a realist framework. What is not clear to us, however, is whether this reaction is justified once that sense is spelled out and clearly distinguished from other senses.

For these reasons, it seems to us that more must be said about the nature of rationally non-optional deliberation to make it plausible that there is an effective indispensability argument for robust realism in the offing.

Before leaving Enoch's positive case for RR, let us say a few words about the role of chapters 4 and 5. In chapter 4 Enoch argues that the conclusion of chapter 2 – that some form of *metaethical* realism, though not necessarily of a robust (non-naturalist) sort, is correct – and the conclusion of chapter 3 – that robust *metanormative* realism is correct – together support, but do not entail, robust *metaethical* realism. We have no immediate objections to the reasoning in this chapter, other than our doubts about the arguments of chapter 2 and 3. The purpose of chapter 5 is to add to the arguments in chapter 2 and 3, where the main target were subjectivist views, by expanding on why the metaphysics of (realist) naturalism, fictionalism, error-theory, and quietist views will not do. Here we can only mention our largest concern about this chapter, namely that the case against naturalist reductions rests entirely on the intuition that normative properties are “just too different”.

In chapter 7 Enoch takes on the epistemological challenge against RR: if moral facts are both causally inert and independent of (not determined by) our moral attitudes, it is hard to see how we can access them. After rejecting a number of ways of construing this challenge, Enoch settles for this:

“[...] very often, when we accept a normative judgment *j*, it is indeed true that *j*; and very often when we do not accept a normative judgment *j* (or at least when we reject it), it is indeed false that *j*. So there is a correlation between (what the realist takes to be) normative truths and our normative judgments. What explains this correlation? On a robustly realist view of normativity, it can't be that our normative judgments are causally or constitutively responsible for the normative truths, because the normative truths are supposed to be independent of our normative judgments. And given that (at least basic) normative truths are causally inert, they are not causally responsible for our normative beliefs.” (159)

The problem, for robust realists, on this construal of the epistemological challenge, is that they are stuck with an unexplained correlation. Enoch's explanation of the correlation, much simplified, is that for many of the things that are in fact normatively good – such as human survival or reproductive success, and things that promote these – believing that they are good promotes survival and reproduction, which has led to the selection of such beliefs. This explains why we tend to have these beliefs, and since they are also true, we have an explanation of the correlation. (7.4.2)

We have no major objections to this explanation, but we doubt that the challenge, understood as the question for an explanation of the correlation between normative beliefs and normative truths, captures what worries people. First, this challenge seems

too easy: surely there should be *some* explanation or other of why we have the moral beliefs we do, and so, given the same normative assumptions that Enoch makes in his explanation, some explanation or other of the correlation in question. Second, it seems to us that the idea that our moral beliefs constitute knowledge is directly undermined by RR's implication that normative facts play no role whatsoever in explaining the shape of the mechanisms by which we form these beliefs. Because the explanation of why we find a moral judgment plausible is exactly the same whether or not the judgment is actually true, the fact that we find it plausible cannot *indicate*, in the sense relevant for knowledge and justification, that the judgment is true rather than false.

Though we have no ready comprehensive theory of knowledge or justification that explains precisely why causal connectedness is required, the requirement strikes us as an independently plausible constraint on such theories. However, it might seem that Enoch has already considered this sort of challenge and provided an adequate reply:

... given the unhappy situation of our understanding of the nature of knowledge, if an account of knowledge entails that we can never know that racist discrimination is unjust (and the like), this should be taken as a strong reason to reject that account of knowledge, rather than to reject normative knowledge. (157)

It might well be true that, all else being equal, we should be more confident in our knowledge that racist discrimination is unjust than in some causal condition on knowledge. But even if we accept this, our choice is not between moral knowledge and a causal condition on knowledge in general, but notably involves a third factor: RR's claim that moral facts are causally inefficacious. For those of us who have much stronger antecedent confidence in the causal condition than in RR (or analogous forms of realism in other domains), the problem for RR would seem to remain: if we have

independent reasons to reject skepticism, these reasons translate into reasons to reject RR.

Not everyone will find global moral or normative skepticism implausible enough to reject any metanormative theory that implies it. But those who do not might notice that Enoch's *positive* arguments for RR in chapters 2 and 3 (as well as his argument against error-theory in chapter 5) would seem to be undermined if RR implies such skepticism: since some of the premises for these arguments are normative, they would be rendered unknowable by RR. If we have no other way of knowing that RR holds, then, RR would itself be unknowable. Because of this, it might matter greatly for the overall success of Enoch's case for RR whether we think that he has adequately addressed epistemic worries about causal inertness.

In chapter 8, Enoch addresses a variety of arguments from disagreement that have or could be directed at RR. The chapter provides a useful overview of such arguments, along with reasonable although largely familiar strategies for dealing with many of them. Though we find much of the discussion convincing, we would like to note two closely related worries.

The first of the two related worries is a familiar one: it is unclear what metasemantic story is compatible with RR. Given causal inertness, moral properties cannot be properties identified by their causal role in shaping our moral thinking or moral practices, so standard externalist accounts of reference are ruled out. But deep and systematic moral disagreements seem to undermine standard internalist accounts of reference in terms of speakers' fundamental substantive criteria of application for normative concepts, as such criteria seem to differ between parties of such

disagreements. Apparently, a defender of RR needs some special story about how reference is fixed.

Enoch sketches an answer to this worry in the previous chapter (section 7.6), where, building on work by Ralph Wedgwood, he suggests that normative concepts have the referent that would render the characteristic practical role of these concepts “largely OK”. For instance, “‘has the strongest undefeated reason to’ would refer to the property that renders our proceeding from the judgment that we have the strongest undefeated reason to ϕ to an intention to ϕ largely OK (180)”. Unfortunately, the suggestion is hard to assess because, as Enoch readily admits, it is less well developed than other core parts of his view. For example, nothing is said about what it means to render inferences associated with a concept “largely OK”, and depending on how this fundamental locution is interpreted, the suggestion might be in strong tension with absolutist assumptions. In particular, one might think that insofar as the rendering-largely-OK relation is supposed to capture what is *meant* by a locution, it has to be sensitive, not only to the practical inferences we normally make on the basis of judgments involving the locution, but also inferences on which our judgments are based, i.e. to our considered assessments of when the concept properly applies. But such assessments are of course exactly what seem to differ systematically and in deep ways between individuals, cultures, and times. Suppose that two communities have a practice of proceeding from the judgment that someone ought to ϕ in circumstances C to an intention to ϕ in C, but recognize systematically different kinds of reasons for their ought-judgments: one community recognizes only a certain kind of consequentialist reasons, whereas the other recognizes only a certain kind of deontological reasons. If assignments of referents are supposed to make *all* their typical inferences largely OK—

inferences *to* ought-judgments as well as inferences *from* such judgments, it might seem prima facie plausible that their concepts have different referents. This, of course, is one sort of worry that has moved people in the direction of various forms of (appraiser) relativism, and it is one that we think remains to be addressed.

The second, related, worry is partly based on our dissatisfaction with Enoch's positive arguments for RR, as the occurrence of deep moral disagreement has the potential to undermine the most prominent alternative reasons to accept the absolutist, non-relativist, aspect of moral realism. In support of this aspect, realists often appeal to the strong sense many have that judges who make different moral judgments and accept different normative theories have real moral disagreements rather talk past one another. Moral judgments, it seems, cannot merely be concerned with whether actions conform or do not conform to the standard of the judge. Noncognitivist expressivists and some relativists have tried to explain the phenomenology of disagreement in terms of agreement or disagreement in practical implications, but even if such explanations would be workable, they would seem to be less straightforward than the explanation offered by the absolutist, appealing to mechanisms not operative in other areas of discourse. This, then, might seem to support absolutism.⁵

Cases of deep and systematic moral disagreement and causal inertness pose a problem for this argument, for it seems that even if proponents might be able to come up with some story about why we think that parties of such disagreement are not talking past one another, this story will be non-standard, appealing to mechanisms not operative

⁵ Brink, David O., *Moral realism and the foundations of ethics* (Cambridge: Cambridge University Press, 1989), ch. 2; Huemer, Michael, *Ethical intuitionism* (Basingstoke: Palgrave Macmillan, 2005), chs. 2-3; McNaughton, David, *Moral vision: An introduction to ethics* (Oxford: Blackwell, 1988) pp. 39-41; Sayre-McCord, Geoffrey, "Moral realism", In *The Oxford handbook of ethical theory*, ed. David Copp, 39-62. (Oxford: Oxford University Press, 2006), p. 42; Shafer-Landau, Russ, *Moral realism: A defence* (Oxford: Oxford University Press, 2003), chs. 2-3; Streiffer, Robert, *Moral relativism and reasons for action* (London: Routledge, 2003), ch. 1.

in other areas. The worry is not necessarily that absolutists will need a non-standard metasemantic account: perhaps people do, deep down, accept the same fundamental criteria for the application of moral concepts, perhaps because of non-obvious commitments they have qua agents. Rather, the problem is that standard theories about reference-fixing would render sameness of reference *non-obvious*, and so unlikely to explain in normal ways why we think that there is substantial disagreement in cases of deep and systematic differences in moral judgment. Whether RR requires a non-standard metasemantic theory or not, then, it requires a nonstandard account of why we take people to be concerned with the same thing, appealing to mechanisms not operative in other areas of discourse. This seems to undermine RR's advantage over noncognitivism or relativism in accounting for intuitions of moral disagreement.⁶ This would not be a serious worry for absolutism if Enoch's argument in chapter 2 were successful, but to the extent that there are doubts about that argument, we think that the absolutist realist should address this worry head on.

In this review, we have discussed what strikes us as comparatively weak points in Enoch's overall argument. But while we are critical, we think that the book's many strengths, not least the fact that it attempts to provide new kinds of positive arguments for non-naturalist realism, makes it an important and highly engaging contribution to the field.⁷

⁶ For a development of this worry, see Gunnar Björnsson's "Do 'objectivist' features of moral discourse and thinking support moral objectivism?". *The Journal of Ethics* 16 (2012), 367-93.

⁷ We thank David Enoch for helpful comments on a draft of this review.