

**David Enoch, *Taking Morality Seriously: A Defense of Robust Realism*
Oxford: Oxford University Press, 2011, pp. 336. ISBN 978-0-19-957996-9,
\$75.00 Hbk.**

David Faraci

Turn over a copy of David Enoch's *Taking Morality Seriously: A Defense of Robust Realism*, and you will find the usual array of quotations extolling the book's virtues. The first is from Russ Shafer-Landau—arguably the philosopher most responsible for the current resurgence of Moorean-style non-naturalism—who writes that this “is the best book that has ever been written in defense of ethical realism.” The last is from Hartry Field, who claims that “on the scale of texts arguing for an obviously false conclusion, this actually ranks pretty high!”

These particular quotations nicely reflect the central take-home messages of this review: First, that this is a book that anyone with an interest in metaethics—not to mention a number of other philosophical disciplines—should read. Second, that it does less than one might hope when it comes to adjudicating the disputes between opposing metaethical camps. Of course, it is true of most metaethics—if not most philosophy—that it rarely persuades the opposition. But part of what makes this case distinctive is that for some opponents, at least, it seems Enoch would not merely fail to persuade them to become Robust Realists; he would ultimately bolster their confidence in their preferred views.

Enoch is convinced—or at least suspicious—that there can be no knock-down arguments for one metaethical theory over another. In light of this, he suggests that we should proceed by looking at the arguments for and objections to each of the available options and determine which, at the end of the day, has the most “plausibility points” (pp. 14-15). Enoch thinks he has found the winner: Robust Realism. On this view, moral and, more broadly, normative facts are *sui generis*, non-natural, irreducible, metaphysically autonomous facts about what we ought to do, what is valuable, what matters, etc. that exist, as it were, in Plato's Heaven.

The book can be split roughly in two. In Chapters 2-5, Enoch's goal is to motivate Robust Realism—to gain it plausibility points—by arguing that the view makes sense of everything a metanormative or metaethical view needs to make sense of—at least, on balance, better than the competition. In Chapters 6-9, Enoch turns to defense, arguing that a number of objections to Robust Realism are either entirely misguided or weak enough that the plausibility points lost are relatively insignificant.

Enoch begins, in Chapter 2, with an argument for the objectivity of moral truth. When people disagree about something, Enoch points out, there is a moral question as to whether they should compromise or stick to their guns. The right answer, he says, depends on the subject matter. When we disagree about things that

ultimately depend in some way on our attitudes—like, perhaps, where would be best to go for dinner—we should compromise. When we disagree about things where there is some objective fact of the matter—like whether the earth is flat—we should, or at least may, stick to our guns. Since, he thinks we will agree, it is right to stick to our guns where morality is concerned, this indicates that there is some objective fact of the matter about morality. So non-objective metaethical views—including, for Enoch, both obviously subjectivist views but also various forms of constructivism and non-cognitivism—are false.

In Chapter 3, Enoch offers an argument from “deliberative indispensability” for the existence of Robust normative facts. He begins with a general discussion of parsimony. He accepts that parsimony has theoretical value, and thus that we should worry about accepting theories that include redundant elements. But while generally our worries about parsimony are reflected in worries about explanatory indispensability, he argues, it is possible, even quite plausible, that theoretical elements can be indispensable in other ways. And so if, as Enoch believes, Robust normative facts are indispensable for deliberation—and deliberation is itself important—acceptance of such facts is hardly ontologically profligate.

Chapters 2 and 3 are the central elements of Enoch’s positive case for Robust Realism, so let us pause for some preliminary critical observations. In Chapter 5, Enoch mentions that he sees normative error theory as a sort of second-best metanormative option, because “error theorists do not kid themselves about the commitments of normative discourse, and they proceed to boldly follow the argument from these commitments and the belief that they are unsatisfied to its natural, if extreme, conclusion.” (p. 115).

Given this, it is rather surprising that Enoch’s argument for objectivity in Chapter 2 should leave error theory out in the cold. After all, Enoch’s argument relies centrally on a substantive moral premise—something the error theorist clearly cannot accept. Of course, this does not mean that Enoch’s argument fails as an argument for objectivity. But it does imply that the argument fails to fully capture the reasons objectivists have for being objectivists, at least insofar as error theorists are warranted in being objectivists. Given Enoch’s respect for error theory, it seems he should agree that there must be a more ecumenical argument in the offing.

Turning to Chapter 3: Enoch’s argument from deliberative indispensability is innovative and important. Indeed, this chapter is an excellent indication of the book’s breadth, and why it should be of interest even to those outside metaethics, given the broad applicability of the discussions of parsimony and indispensability. If nothing else, then, this chapter is useful in that it illuminates the possibility of non-explanatory indispensability. Arguably, the arguments in this chapter are also successful at countering the worry that Robust normative facts are just obviously theoretically extravagant. But elegant as the arguments are, the proof is in the

pudding: The question is not whether deliberative indispensability can, in theory, justify acceptance of otherwise redundant theoretical elements, but whether Robust normative facts really are indispensable.

Enoch says a bit in this chapter in support of such indispensability. It seems worth highlighting one aspect of this discussion because, whether or not there is something inherently problematic in what Enoch says, it is particularly striking given his other commitments. Enoch points out that just because a theoretical element is indispensable for a particular project, this does not mean that we are therefore justified in including that element in our ontology. Such “instrumental indispensability” for a project has ontological implications only insofar as the project itself it is “intrinsically indispensable.” Enoch suggests that deliberation is indispensable in this way at least in part because it is “rationally non-optional” (pp. 69-71).

This claim needs a lot of defense, but let us set that aside. The focus here is on how striking this claim is given the work Enoch has done elsewhere to undermine constitutivist views according to which normative truth is derivable from the demands of agency. Enoch’s position has been—very roughly—that such views fail because they necessarily leave open the question of why one ought to be an agent at all. But if one understands what it is to be an agent just as being confronted with the question of what one ought to do—as Christine Korsgaard sometimes seems to—then it seems that perhaps either (a) Enoch is hinting at a potential solution: being an agent is rationally non-optional or (b) Enoch runs up against his own argument: saying that deliberation is rationally non-optional is question-begging precisely because it is a normative claim, and thus just as the constitutivist cannot use it to ground her view, as this would be circular, Enoch arguably cannot appeal to a normative fact to explain why normative facts are indispensable. Now, certainly, no pretense should be made that this is a fully thought-out objection. But it is nevertheless worth noting if only because the link is so striking. Indeed, Enoch even comments on the “Kantian” nature of his argument (p. 76).

That aside, the bit about rational non-optimality is meant to show only that deliberation is “intrinsically indispensable.” It remains to be shown that Robust normative facts are in fact instrumentally indispensable for deliberation. The bulk of Enoch’s case for this claim comes in Chapter 5, with his rejection of naturalism, fictionalism, error theory and quietism—alternatives that might make good on normativity’s role in deliberation while “doing with less” metaphysically-speaking.

Before getting to these, Enoch uses Chapter 4 to combine his arguments from the previous two chapters in order to defend Robust Metaethical Realism. In brief, Enoch offers reasons for thinking that if one accepts Robust Metanormative Realism, as argued for in Chapter 3, then given the relationship between morality and normativity, it makes sense to further accept Robust Realism about the metaethical.

He goes on to consider and reject some general worries—such as about categoricity—that might lead one to be suspicious of Robust Realism about moral facts even if one accepts it about normative ones.

Before considering Enoch’s arguments in Chapter 5 in more detail, let us continue on to the second half of the book. By the end of Chapter 5, Enoch takes himself to have garnered a significant number of plausibility points for Robust Realism. But, of course, Robust Realism is also subject to a number of important objections. Were any of these successful, the loss in plausibility points might more than make up for what has been gained. So in Chapters 6-9, Enoch addresses some of the objections historically raised against his view.

In Chapter 6, Enoch tackles metaphysical objections. Most of his work here is in addressing the worry that non-naturalists cannot make sense of the supervenience of the moral on the non-moral. Enoch points out that this is actually two challenges: one, “general supervenience,” to explain why the moral supervenes on the non-moral at all, and the other, “specific supervenience,” to explain why the moral properties supervene on the particular non-moral properties they do (pp. 142-143). General supervenience, Enoch claims, can be dealt with by an appeal to conceptual necessity—it is part of our moral concepts that the moral supervenes on the non-moral.

Specific supervenience is more difficult, Enoch admits. He does a bit to try to soften the blow, but ultimately concludes that Robust Realism is indeed committed to some basic, substantive moral facts—that is, some relations between moral and non-moral properties—that have no explanation, that are brute. He acknowledges that this violates the widely-held “Hume’s Dictum” that there can be no “necessary relations between distinct existences” (p. 147). But, ultimately, it seems to Enoch that if he can have everything he wants just so long as he gives up this “piece of metaphysical dogma,” he is really not so badly off (p. 148).

In Chapter 7, Enoch turns to epistemological objections. There are many, and so Enoch first develops what he sees as the best version of the challenge: to explain correlations between normative beliefs and normative facts. As he points out, correlations are typically explained by showing that one correlate explains the other. Unfortunately, explanations of this kind for the correlation between normative beliefs and normative facts are, it seems, ruled out by the combination of Robust Realism’s commitment to some version of the thesis of attitude-, mind-, response-, or stance-independence, its commitment to the metaphysical autonomy of the normative, and the fact that Robust normative properties are almost certainly causally inefficacious.

Enoch’s innovation is to recognize a third form of explanation for correlations: The correlation between normative beliefs and normative facts might be explained by a third factor that generates a sort of “pre-established harmony”

between them (p. 168). He proposes that such an explanation might be had by looking to evolution. This is ironic, given that the current fashion is to use evolution to make the challenge. Enoch's suggestion, very roughly, is that if we grant that survival is good, then when evolutionary pressures explain our normative beliefs, they also ensure that those beliefs track something good, since it is survival—which, again, is good—that generates these pressures in the first place.

Next, in Chapter 8, Enoch considers arguments from moral disagreement. Here, he distinguishes and tackles an impressive list of such arguments—everything from the simplest, invalid argument that disagreement entails skepticism to more sophisticated abductive arguments. Enoch addresses each in turn, arguing of each formulation either that it fails outright, that it reduces to an objection addressed elsewhere in the book, or at least that there is enough hope of overcoming it to leave Robust Realism on solid ground. He concludes that arguments from actual disagreement have little to no force against Robust Realism, while arguments from idealized disagreement will lose the view some—but probably not very many—plausibility points.

In Chapter 9, Enoch addresses the purported connection between normativity and motivation. Enoch begins by laying out his theory of what it is to act for a reason, and distinguishes the idea of an “agent’s reason” for acting from motivational and normative reasons (p. 221). He then argues that Robust Realism is consistent with acting for a reason—something some have been concerned the view could not make sense of. Next, he addresses other issues about motivation, including the worry that Robust Realism is incompatible with normativity’s being “practically relevant,” the “why be moral?” question, and both judgement and existence internalisms about reasons (pp. 237-266). In each case, he argues that the worries about Robust Realism are likely illusory. He concludes that Robust Realism loses no plausibility points because of motivation. He then recaps his arguments throughout the book and tallies Robust Realism’s plausibility points in a final chapter.

By this point, the grounds for our first take-home message—that anyone with an interest in metaethics should read this book—should be fairly clear. There are, as one might expect, small things throughout to quibble with. For example, Enoch’s discussion in Chapter 7 of the forms of explanation one might offer for a correlation is somewhat lacking. Indeed, I argue elsewhere that Enoch misidentifies the form of his solution. But for the most part, the arguments are careful and even persuasive. Aside from anything else, insofar as this book is meant to show that Robust Realism is a metaethical contender that must be taken seriously, it passes with flying colors. The problems come when one considers the book as a positive argument for Robust Realism as the best view on the table.

Here is a very rough reconstruction of the line of argument in this book:

- (1) Moral beliefs concern something objective.
- (2) If Robust normative facts are indispensable for deliberation, we have some reason to believe in them, and thus to accept Robust Metanormative Realism.
- (3) Insofar as we have reason to accept Robust Metanormative Realism, we have reason to accept Robust Metaethical Realism.
- (4) Robust normative facts are indispensable for deliberation because alternatives to Robust Realism that are consistent with (1) and with normativity's role in deliberation fail.
- (5) None of the metaphysical, epistemological, semantic or psychological objections to Robust Realism are significantly damaging.
- (6) Therefore, we have most reason to accept Robust Realism in both metaethics and metanormative theory.

Now suppose you held any of the following metanormative or metaethical views: constitutivism, error theory, fictionalism, naturalism, quasi-realism, quietism. There is a good chance that you would accept (1) and (2). Depending on whether your view is metanormative or metaethical, you might also accept (3). It turns out, then, that you might well read the first one hundred pages—four chapters—of Enoch's book without batting an eyelash, with a few notable exceptions. Given the large percentage of ethicists and metaethicists who accept one of the views mentioned, then, it seems that the bulk of Enoch's work—insofar as he is pushing for Robust Realism over these alternatives—will be in defending (4) and (5).

The first bit of defense of (4) comes in Chapter 2—one of those notable exceptions within the first hundred pages—where Enoch addresses constitutivism and quasi-realism. Enoch thinks that these views are going to turn out to be inconsistent with morality's objectivity. The argument against quasi-realism is basically that a view on which the moral truth depends in some important way on our moral attitudes—like quasi-realism—cannot possibly be objective in the relevant sense. Now, that may well be true. But it is not clear how this argument moves the dialectic forward. After all, here is what a quasi-realist—of one kind, at least—thinks: Moral language expresses non-cognitive moral attitudes, but nevertheless the deflated moral truths this language concerns are appropriately objective. So to some extent, Enoch's objection seems to be little more than a restatement of the quasi-realist's view. Now, certainly, it would be uncharitable to say that Enoch's objection is just a restatement of the view. But the argument does have a bit of a foot-stomping

feel. It is little more than insistence that quasi-realists cannot do what they think they can do.

The rest of the views mentioned—error theory, fictionalism, naturalism and quietism—are addressed in Chapter 5. Let us focus on naturalism—though in fairness to Enoch, it should be noted that at the end of the book, when he lists the bits that he is least confident in, he is explicit that he wishes he had more to say against naturalist views. Enoch says a number of interesting things about how to talk about reduction so as to step on as few metaphysical toes as possible, about the objections to naturalism already present in the literature, and about why a particular argument from Mark Schroeder that is supposed to speak in favor of naturalistic reduction actually seems to speak against it. Ultimately, though, what Enoch relies on most heavily is the idea that moral and normative properties are “just so different” from naturalistic properties that the former could not possibly reduce to the latter (p. 80).

Now, suppose you are a committed naturalist of a particular sort. You read the first hundred pages of Enoch’s book and find that you agree with nearly everything he says. You agree that when we disagree about moral matters, we should stand our ground. You agree that normative truth of some kind is indispensable for deliberation. And you agree that we should treat morality and normativity roughly the same, metaphysically-speaking. Then you get to Chapter 5, where you read a careful and well-thought-out explanation of why Enoch rejects your view. But careful and clear as it may be, at the end it comes down to a just-too-different intuition, which you have been hearing about from every non-naturalist you have ever met.

At this point, if you are anything like Enoch, you will start counting plausibility points. You have gained a number of points relative to certain other views because of Enoch’s arguments in Chapters 2 and 3. So now you have a whole bunch of plausibility points that you may or may not have taken yourself to have before you started reading. In fact, you have roughly the same number of plausibility points as Robust Realism, and that is a lot more than some alternatives!

Now, you know that if you go Enoch’s Robustly Realist route you will start losing points, at least because you will have to accept non-naturalism. But luckily for you, you do not share his just-too-different intuition, so in sticking with your own view and rejecting that intuition, you do not lose many plausibility points—or at least so it will seem to you. And so if you have, or think you can develop, a view that can deal with all of the other objections Enoch mentions—open question arguments, moral twin earth, etc.—then it seems your view is looking at least as attractive as it was before you began reading, if not significantly more so.

Of course, the question is not just who has the best positive case, but who loses the most plausibility points to attacks. And you might not be as confident as you would like about your particular reductive story. So you turn to the second half of the book, which comprises Enoch’s defense of (5) in the loose reconstruction of his

overall argument above, to see if Robust Realism is, after all, more plausible than you thought it was.

There is something that is at once incredibly admirable and incredibly frustrating about the second half of Enoch's book: Enoch seems to try very, very hard to get the objections to his view as clear and strong as possible before responding to them. Arguably the best example of this comes in his discussion of epistemological objections in Chapter 7. There are many, many different epistemological challenges to realism in the literature. So Enoch offers a general challenge that bypasses a host of contentious epistemological assumptions built into the original objections as presented. After all, it is doubtful that most non-Robust Realists take their objections to succeed because of their particular epistemological views—the problem is not really about justification *per se* rather than, say, reliability. So by moving to a challenge to explain a correlation that, it seems, every non-skeptic is committed to, Enoch both strengthens our understanding of the real issue and makes his ultimate response that much more powerful.

But Enoch's hard work in these areas is also a source of frustration in at least two ways. The first is in comparing Enoch's reconstruction of others' objections with his development of his own. Contrast, for example, the thirty-two-page Chapter 8 in which Enoch attempts to distinguish all the ways in which one might use moral disagreement to make trouble for Robust Realism with the short section in Chapter 5—one quarter the length—in which he rejects naturalism. One has to wonder why, when Enoch is so careful to try to figure out what is behind the objections to his view, he is willing to play somewhat fast and loose with his own objections—saying very little, for example, about exactly why it is that moral facts are just too different to be natural.

The other source of frustration along these lines is that despite his hard work, it seems that Enoch's careful formulations of his opponents' objections sometimes fail to capture their real worries. Now, in all fairness, this is not entirely—probably not even largely—Enoch's fault. Indeed, the mere fact that he has to spend so much time and effort clarifying the objections against his view speaks to how poor, or at least how varied, the original formulations of these objections often are. Nevertheless, Enoch has asked us to join him in counting up more plausibility points for Robust Realism than for any alternative, and here he arguably falls short.

First, setting aside for the moment the worry about getting the objections right, it is not at all clear that Enoch's plausibility point losses are as low as he thinks. Consider again his arguments about epistemology. According to Enoch, his losses are minimal here because it turns out that if survival is good and the best way for our normative beliefs to promote survival is for them to be true and this can account for enough of our normative beliefs to satisfy the non-skeptic's worry, then the Robust

Realist has a way to explain the correlation between normative beliefs and normative facts.

Now, how many plausibility points this story loses him is up for debate. But two things should be noted: First, the plausibility points game is a comparative one, and it certainly seems as though the above explanation is comparatively less plausible, all else being equal, than, say, the explanation that there is a causal connection between our beliefs and the natural properties normative properties reduce to.

Second, one must be careful—which it is not always clear Enoch is—not to infer from Robust Realism’s mere ability to explain some X that it does not lose a significant number of plausibility points with respect to X. After all, the whole point of plausibility points was supposed to be that we are moving away from attempts to show that a view is utterly devastated by its complete inability to make sense of something and instead considering its relative plausibility with respect to other contenders. So surely how many plausibility points a view loses in its attempts to explain X depends not just on whether it can, but on how far it has to bend over backwards to get an explanation off the ground. And with the evolution-based story, one might very well worry that Robust Realism is going to get a hernia.

In any case, let us return to the concern that Enoch is not always successful in his attempts to understand what is going on in the objections he addresses. Perhaps the clearest case is in his discussion of the objection from supervenience in Chapter 6. Enoch spends a great deal of time explaining how Robust Realism can make sense of specific supervenience and defending his rejection of Hume’s Dictum. His discussion of how to explain general supervenience is comparatively short: He simply appeals to the fact that it is a conceptual truth.

But this seems too quick. Normative supervenience might well be a conceptual truth. But there remains a worry that Robust Realism is in tension with its being so. Out there in Plato’s Heaven there is a non-natural, abstract property of goodness. Somehow, here on Earth, we know that this property is related to the non-moral such that there can be no change in a thing’s goodness without a change in its non-moral properties. Not only do we know this, but we know it as a matter of conceptual necessity. Now this is decidedly bizarre. It is difficult to see how it could be that when this concept developed, it captured a necessary fact about this property—that it supervenes. And, for that matter, it seems we should wonder why it supervenes at all, why this non-natural property of goodness cannot just attach itself to whatever it feels like, so to speak.

The above is, of course, quite inchoate. And it may or may not be the concern had by all or most of those Enoch is responding to—though Tristram McPherson does raise related worries in a paper Enoch makes reference to. Nevertheless, it does seem

to be a genuine worry about Robust Realism's ability to account for supervenience. And it is a worry that Enoch does not address.

At the close of the book, Enoch invites the reader to share his conclusions:

[W]hen I look at the project as a whole, I find myself reasonably confident in it and its conclusions. Perhaps somewhat paradoxically, I find myself more confident in the general picture than in some of its details. Trying to tally plausibility points, then, I find myself thinking that Robust Realism does pretty well, better, in fact, than competing views. Don't you? [p. 271]

For those friendly to non-naturalism, it seems the answer would be a resounding yes. Enoch's book is full of interesting, important, clear and, frankly, fun arguments. And insofar as you share his inclinations, you are likely to find a good deal to vindicate your shared intuitions in the pages of this book. The trouble is that insofar as you do not—insofar as you lack the just-too-different intuition, or think an expressivist semantics can capture objectivity, or think that the supervenience objection is not quite what Enoch says—you will probably find more in this book to support the view you prefer than you will cause to become a Robust Realist. Then again, that is really not so bad. Because whoever Enoch's arguments ultimately support, they are, to echo Shafer-Landau, some of the best you will find.