

The epistemological challenge to metanormative realism: how best to understand it, and how to cope with it

David Enoch

Published online: 20 January 2009
© Springer Science+Business Media B.V. 2009

Abstract Metaethical—or, more generally, metanormative—realism faces a serious epistemological challenge. Realists owe us—very roughly speaking—an account of how it is that we can have epistemic access to the normative truths about which they are realists. This much is, it seems, uncontroversial among metaethicists, myself included. But this is as far as the agreement goes, for it is not clear—nor uncontroversial—how best to understand the challenge, what the best realist way of coping with it is, and how successful this attempt is. In this paper I try, first, to present the challenge in its strongest version, and second, to show how realists—indeed, robust realists—can cope with it. The strongest version of the challenge is, I argue, that of explaining the correlation between our normative beliefs and the independent normative truths. And I suggest an evolutionary explanation (of a pre-established harmony kind) as a way of solving it.

Keywords Moral realism · Moral epistemology

1 Introduction

Metaethical—or, more generally, metanormative—realism faces a serious epistemological challenge,¹ which can be put—very roughly, and in a way to be refined later on—thus: Realists owe us an account of how it is that we can have epistemic access to the normative truths about which they are realists. If normative beliefs are about an independent order of normative facts, how is it that we know anything

¹ For one *locus classicus*, see Mackie (1977, p. 38).

D. Enoch (✉)
The Philosophy Department, The Hebrew University in Jerusalem, Mt. Scopus Campus,
Jerusalem 91905, Israel
e-mail: denoch@mscc.huji.ac.il

about them? How are these beliefs even justified, when they are? If realists resort to the highly suspicious faculty of rational intuition, how do they suggest to best understand its nature? What evidence can they present that it exists?

That realism faces this challenge is, it seems, uncontroversial among metaethicists, myself included. But this is as far as the agreement goes, for it is neither clear nor uncontroversial how best to understand the challenge, what the best realist attempt of coping with it is, and how successful this attempt is. In this paper I try, first, to present the challenge in its strongest version, and second, to show how realists—indeed, robust realists—can cope with it. I do this, first, by criticizing (in Sect. 2) ways of articulating the epistemological challenge to realism that are found in the literature. I do not claim that these understandings of the challenge are theoretically pointless. But I do claim that they do not present the best—most challenging, and most challenging specifically for realists—version of the epistemological challenge. This purely negative discussion will result in a better grasp of the conditions a characterization of the epistemological challenge has to meet if it is to be theoretically useful in this way. With these desiderata in mind, I then present (in Sect. 3) a *very* strong version of the epistemological challenge to realism, a version which has not, I believe, been adequately addressed in the literature, and which is the one I proceed to clarify and—eventually—respond to. I then argue that realists need a solution to the epistemological challenge *thus understood*. This interim conclusion is, I want to emphasize, independent of my suggested solution to the challenge. In other words, even if I am wrong about how to solve the challenge, my conclusion in Sect. 3 regarding the best way to understand the challenge stands. We realists shouldn't kid ourselves—it is a solution to this version of the challenge that we should be after. Before proceeding to suggest a solution to the epistemological challenge, I dedicate Sect. 4 to a discussion of Sharon Street's recent "Darwinian Dilemma" challenge to realism. I argue that it is a particular instance of the epistemological challenge as articulated in the preceding section, and that noticing this fact is the first step towards adequately addressing it. In Sect. 5 I finally present my suggested solution to the epistemological challenge properly understood. The solution I suggest is somewhat modest in its ambition: I argue that the explanatory problem realists (and indeed, robust realists) face is at least manageable, and that the theoretical price realists may end up paying because of the epistemological challenge is much lower than it may initially be thought. I conclude that given the (purported) other, unrelated advantages of Robust realism (which I do not discuss here), it may very well still be the view to endorse, some possible remaining epistemological price notwithstanding. In the final Sect. 6 I briefly comment on the generality of the epistemological challenge (properly understood) and of my solution to it.

Now, the term "realism" is notoriously unclear and ambiguous, in the metaethical context as elsewhere. Let me note, then, that in what follows by "realism" I will for the most part mean Robust Realism, the view—which I characterize and argue for at length elsewhere²—according to which, somewhat roughly, there are irreducibly, non-naturalist normative truths, response-independent truths that are perfectly

² Enoch (2003, 2007).

objective and that are not reducible to—not even identical with—natural, not-obviously-normative truths. Notice that this is a realism of a fairly strong, uncompromising, kind, indeed probably no weaker a metanormative realism than almost any I know of in the contemporary literature.³ Restricting the discussion to just this kind of realism cannot be objectionable, then, and where relevant in what follows I comment on the applicability of the discussion to other, naturalist, forms of realism.

2 How *not* to understand the epistemological challenge

In this section I argue that the epistemological challenge to realism is not most helpfully put in terms of epistemic access, or epistemic justification of normative beliefs, or their reliability, or normative knowledge, or a particular instance of general skeptical worries, thus preparing the ground for the next section, where I present the version of this challenge that I will be most interested in.

2.1 Access

The worry is sometimes put in terms of epistemic access.⁴ If there are normative truths that are independent of us—of our motivations, our responses, our social practices, indeed our normative judgments themselves—how can we gain epistemic access to them? Indeed, even assuming for the sake of argument that we do have normative knowledge, *how* is it that we do? By what *means* can we gain such knowledge? And the thought is, of course, that if the realist has nothing to say about the *means* by which some privileged epistemic status (like perhaps knowledge) is secured for normative beliefs, she has to conclude that normative beliefs do not after all enjoy this privileged epistemic status.

I think we can rather safely postpone discussion of these worries to the following subsections, without saying much more on epistemic access. This is not just because one way of understanding talk of epistemic access is as an unofficial introduction to one of the other ways of stating the challenge, or because as they stand, worries about epistemic access are too metaphorical to be theoretically helpful (it isn't clear, after all, what "access" exactly means here). The more important reason why we can safely avoid further discussion of the worry put in terms of epistemic access is the following. In the following subsections, I discuss versions of the epistemological worry put in terms of justification, reliability, and knowledge. It is possible, of course, that my arguments there fail. But if they do not, what *remaining* epistemological worry could talk of epistemic access introduce? If in the next subsections I manage to convince you that there are no special problems with the justification of normative beliefs, with the reliability of normative beliefs, or with normative knowledge, it seems to me you should be epistemologically satisfied. I do not see how talk of epistemic access should make you worried again.

³ With the possible exception of Oddie's (2005) realism, which is committed—unlike my realism—to the causal efficacy of moral facts. I return to this point below.

⁴ See, for instance, Timmons (1990, p. 114).

However exactly talk of epistemic access is understood, then, it does not pose an epistemological challenge that is independent of those discussed below.

2.2 Justification

Many realists (and perhaps others as well) take the epistemological challenge to be one about the possibility—or perhaps actuality—of epistemic justification for our normative beliefs. And they reply in the obvious ways: by developing a theory of epistemic justification that allows for such justification of normative beliefs, or—more commonly—by showing that their favorite theory of epistemic justification in general nicely applies to the case of normative beliefs.⁵

It is, of course, an important metanormative (and also normative) question how normative beliefs (if indeed this is what they are) are to be justified, and so such intellectual endeavors are not without value. But still this is not a promising way of understanding specifically the epistemological challenge to realism, for the following three reasons.

First, on no theory of epistemic justification I am aware of do normative beliefs constitute an interesting particular instance of beliefs, an especially problematic class of beliefs (I'll have to qualify this point below). Whether you are a coherentist or a foundationalist (or perhaps hoping for some middle ground between them), whether you are an internalist or externalist about epistemic justification (or perhaps hoping for some middle ground between them), whether or not you think that epistemic justification is conceptually tied to epistemic responsibility, whether or not you like working with a conception of epistemic virtue—whatever your theory of epistemic justification, it is hard to see any *special* difficulties applying it to normative beliefs. Surely, for instance, the relations between beliefs that coherentists emphasize—consistency, explanatory coherence, etc.—can have as their relational normative beliefs as well.⁶ Similarly, though perhaps there are no attractive candidates for indubitable normative foundations, we have pretty much given up on foundations thus understood elsewhere as well, and once foundationalism is made appropriately modest in its requirements—settling, perhaps, for defeasible justificatory status for the foundations—there are after all attractive candidates for such foundations among our normative beliefs as well.⁷ Furthermore, it seems one can be as epistemically responsible (or irresponsible) in one's normative beliefs as in any other, it seems one can hold a normative belief—just like any other—against one's epistemic duties (if there are such duties), it seems one can exercise—or fail to exercise—any intellectual virtue in the forming, revising or holding of normative beliefs, and so on. If the problem is that of allowing for the epistemic justification of normative beliefs, then—given an acceptable general theory of epistemic

⁵ See, for instance, Brink (1989, Chap. 5), Sayre-McCord (1995), Scanlon (1998, pp. 64–72) and Shafer-Landau (2003, Chaps. 10–12). For an antirealist presentation of the challenge in terms of justification (and also in terms of knowledge), see, for instance, Waldron (1992, p. 175). For a skepticism-friendly discussion of moral epistemology that focuses primarily on moral knowledge and justification see Sinnott-Armstrong (2006).

⁶ See, for instance, Brink (1989, Chap. 5) and Sayre-McCord (1995).

⁷ See Audi (1995) and Shafer-Landau (2003, p. 258).

justification—it seems too easy to solve, at least in outline.⁸ And this suggests that the epistemological challenge to realism should not be understood as one about justification at all. Things would have been different, of course, if there was any reason to believe that some feature of normative beliefs made them an *especially* problematic class of beliefs when it comes to justification. But I know of no reason to believe that this is so.⁹

To see my second reason for thinking that the epistemological challenge to realism is not best thought of as a challenge about justification, remember that we are trying to articulate an epistemological challenge *to realism*, and so one that highlights a difficulty that metanormative realists face but that others do not, or at least a difficulty that is especially acute for realists. But the problem of accommodating epistemic justification for normative beliefs does not seem to be of this sort. Notice that nothing in the previous paragraph seemed to depend on a realist view of normativity. Of course, perhaps some metanormative views work better with some theories of justification—everything here depends on the details. And so perhaps it can after all be shown that the problem of accommodating the epistemic justification of normative beliefs is somewhat more acute for realism (or for some kind of realism) than in general. But I don't know of any way of filling in the details here that has this as a consequence. Until one is presented, there is no *specific* problem here for realism.

Third, let us not pretend that we have a theory of epistemic justification we are all happy with. Given this unhappy state of affairs, if a theory of epistemic justification rules out all normative beliefs as epistemically unjustified, this counts at least as heavily against that theory of justification as it does against the justificatory status of normative beliefs. For is there really a theory of epistemic justification in the truth of which you are (and should be) more confident than in the truth of such sentences as “I justifiably believe that torture is *prima facie* wrong”? When we try to determine which theory of epistemic justification to accept, it seems to me it counts rather heavily against a proposed theory if according to it no normative belief is justified.¹⁰ (It counts even more heavily against such a theory if according to it no a priori belief is justified, a point to which I return.) Now, perhaps given a theory of epistemic justification that entailed this surprising result, but otherwise scored sufficiently highly on the list of theoretical virtues, we would be justified in biting the bullet and acknowledging that no normative belief is ever justified.¹¹ But it seems to me highly unlikely that we have such a theory at hand or that one is likely to emerge. And this

⁸ At least with regard to justification, then, I agree with Sayre-McCord (1995, p. 138) when he says: “there is no distinctive epistemology of moral belief”.

⁹ It may be thought that the fact that normative beliefs are *controversial* in some fairly robust way could do the job here. But it is not at all clear that they *are* controversial in a way that many other beliefs aren't, nor is it clear what follows from such disagreement that we do find here. For my own views on this, see Enoch (forthcoming).

¹⁰ Compare: If the causal theory of knowledge entails that we have no mathematical knowledge, this counts heavily against that theory, perhaps *more* heavily against it than against mathematical knowledge.

¹¹ Compare: If the causal theory of knowledge entails that we have no mathematical knowledge, but otherwise scores sufficiently highly on the list of theoretical virtues, perhaps we would be justified in biting the bullet and acknowledging that we never have mathematical knowledge, indeed that such knowledge is impossible.

poses yet another difficulty for the attempt to understand the epistemological challenge to realism as one about epistemic justification.

Important questions about the justification of normative beliefs remain, of course, and I am not suggesting that there are no interesting philosophical challenges in filling in the details of a theory of justification and a metanormative theory so that the two cohere nicely with each other in a non-skeptical way. But for the reasons just specified, I suggest that we not understand the epistemological challenge to realism as one about justification.

2.3 Reliability

Reliability may be thought necessary for justification, or for knowledge, or perhaps epistemically relevant in some other way. So if there is a special problem about the reliability of normative beliefs, perhaps it is in terms of reliability that the epistemological challenge to realism is best understood.

I think there is something right in this suggestion, and I get back to it in Sect. 3. But let me here set aside what is not a good way of understanding the challenge. The problem cannot be that normative beliefs *cannot* be reliable, or that classifying them either as reliable or as unreliable constitutes a category mistake, or something of that sort.

A class of beliefs is reliable, I take it, if and only if a sufficiently large portion of it is true.¹² If so, wherever there is truth, there is the possibility of reliability. Now, it is not obvious, of course, that there *is* truth in ethics and in normativity, and some metaethicists think no such truth is to be had.¹³ But such a claim can certainly not be used as a premise in what is supposed to be an independent challenge to realism. If there are no normative truths, then realism is indeed defeated, but not by the epistemological challenge, but rather by whatever argument is supposed to support the claim that there are no normative truths. In our context, then, the antirealist cannot help himself to anything like the denial of normative truths. And if so, nor can he establish a worry about the reliability of normative beliefs.

There is no objection to the very attribution of reliability to some sets of normative beliefs, then, at least not one that is not question-begging in our context.

2.4 Knowledge

Isn't the epistemological worry, though, primarily about knowledge? Isn't the worry that realists have no plausible way of accounting for moral or more broadly

¹² You may think that more is needed for reliability, and in particular that reliability contains some modal feature, so that a set of beliefs is reliable only if, roughly speaking, it is not an accident that sufficiently many of them are true. At least one natural way of understanding such a requirement leads to what I take to be the proper understanding of the epistemological challenge, in Sect. 3.

¹³ I have in mind error-theorists, who think that moral or normative discourse is systematically mistaken, and so that no normative judgment is (non-trivially) true, and expressivists who think (or at least used to think, before quasi-realism became the fashionable way of being an expressivist) that normative judgments are not truth-evaluable at all.

normative knowledge? Isn't the worry that in order to allow for the kind of access or relation to the purportedly independent truths the realist is going to have to resort to that awful trick, the mysterious faculty of rational intuition? If she does, it seems her realism should be rejected for this very reason, and if she does not, she cannot accommodate normative knowledge, and so must settle for a very skeptical kind of realism. The epistemological challenge thus comes down to—unsurprisingly, perhaps—one about the impossibility of knowledge (on realist assumptions).

But this too is not, I think, a promising way of stating the challenge. To see why, let me assume (for now) that knowledge is justified true belief of a special kind (the qualification needed in order to deal with Gettier cases). If this is so, all that the realist has to do in order to accommodate normative knowledge is to account for normative truths, normative beliefs, justification of normative beliefs, and that extra bit needed to deal with Gettier. Very well, then, let's see: For the reason mentioned in the previous subsection, in the context of an attempt to articulate an independent challenge to realism, we are to assume that there are normative truths, and furthermore that at least some normative judgments express beliefs (rather than, say, some conative attitude). And we already dealt with justification in Sect. 2.2. All that remains, then, is to accommodate with regard to normative beliefs the necessary anti-Gettier clause, whatever it is. But then surely the realist is entitled at least to a wait-and-see attitude. "Given that I've given you already so much, what reason do we have to believe that I won't be able to give you *that* too? When you epistemologists have settled on an anti-Gettier clause, let me know, and we'll take it from there." Things would have been different had we had some reason to think that any plausible anti-Gettier clause is going to disallow normative knowledge altogether. But I see no reason so to think. Indeed, we can draw the distinction between Gettier-cases and non-Gettier-cases *within* the set of normative beliefs. Furthermore, the anti-luck intuition with which Gettier is often identified—that what is needed for knowledge is an unluckily justified true belief—seems to apply in the normative realm as it does anywhere else (compare a sophisticated true normative conclusion reached by a sloppy thinker and by a careful one).¹⁴

True, some suggested anti-Gettier clauses are hard to apply with regard to normative beliefs (I have in mind causal and counterfactual ones, which—assuming that normative truths are causally inert and that some of them are necessary—raise problems for normative knowledge). And indeed, it is no longer uncontroversial that knowledge can be analyzed as justified true belief together with some anti-Gettier clause or another. But here I want to rely again on a point already made with regard to justification: 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.¹⁵

¹⁴ But see footnote 48 below.

¹⁵ This is especially true if the scope for which the relevant account of knowledge causes trouble is even wider than the normative—if it includes, for instance, the a priori in general. For this dialectical point, mostly in the context of discussing the causal theory of knowledge, see Huemer (2005, p. 123), Liggins (2006), and somewhat more generally, Lewis (1986, p. 109).

2.5 Particular instances of skeptical worries

But perhaps all of this has been misguided, as I have been assuming that *some* of our beliefs are justified, that some of our beliefs constitute knowledge, etc. In an epistemological discussion, it may be thought, such assumptions cannot be taken for granted. Perhaps, in other words, the epistemological challenge to metanormative realism should be understood as a particular instance of a general skeptical worry, or perhaps a family of particular instances of general skeptical worries?

It is, of course, a worthwhile epistemological project to find ways of coping even with the most general of skeptical worries. But at least absent some further story, it is hard to see why worries about evil demons, or brains in vats, or criteria, or regresses, or fallibility, should prove to be especially problematic with regard to *normative* knowledge (or justification, or warrant, or whatever), or indeed why—if they cannot be satisfactorily addressed—such worries are especially problematic to metanormative realism compared to competing metanormative views.¹⁶ And though I think we should take such skeptical worries very seriously, they do not give the realist any reason to worry that is not shared by everyone else. They do not, in other words, increase the plausibility of any competing metanormative theory, and so are not even the beginning of an objection to realism.

If it could be shown that some general skeptical argument is *especially* worrying when it comes to normative beliefs, or perhaps with regard to normative beliefs realistically understood, this would of course be of interest in our context. But I don't know of any attempt—let alone successful attempt—to show any such thing.

2.6 What we want from the epistemological challenge

Here, then, are the desiderata for a suggested understanding of the epistemological challenge to metanormative realism: We want it to be peculiar to *normative* beliefs; we want it to pose a special, or at least especially hard, problem *for realism*; we do not want it to beg the question against the realist (by, for instance, assuming that there are no normative truths); and we want to avoid as much as possible relying on a highly contested epistemological notion (like that of knowledge or justification). Let me turn to the challenge that does, I think, satisfy all these desiderata.

3 How to understand the epistemological challenge: explaining correlations

It will prove helpful, I think, to introduce the challenge with a couple of analogies. After presenting the challenge, I then comment on some implications of understanding the epistemological challenge along these lines.

¹⁶ Much of Shafer-Landau's (2003, e.g. p. 235) epistemological discussion in the context of a general defense of metaethical realism is really an attempt to deal with particular instances of general skeptical worries. And he repeatedly emphasizes this very point—that the challenge to realism is a particular instance of general skeptical worries—as a part of his defense of realism (see, e.g. p. 239). But I take this not as evidence that realism can cope with the epistemological challenge, but rather as evidence that this is not the best way to put the epistemological challenge.

3.1 The challenge

Suppose that Josh has many beliefs about a distant village in Nepal. And suppose that very often his beliefs about the village are true. Indeed, a very high proportion of his beliefs about this village are true, and he believes many of the truths about this village. In other words, there is a striking correlation between Josh's beliefs about that village and the truths about that village. Such a striking correlation calls for explanation. And in such a case there is no mystery about how such an explanation would go—we would probably look for a causal route from the Nepalese village to Josh (he was there, saw all there is to see and remembers all there is to remember, he read texts that were written by people who were there, etc.). The reason we are so confident that there is such an explanation is precisely that the striking correlation is so striking—absent some such explanation, the correlation would be just too miraculous to believe.

Utilizing such an example, Field (1989, pp. 25–30) suggests the following problem for mathematical Platonism: Mathematicians are remarkably good when it comes to their mathematical beliefs. Almost always, when mathematicians believe a mathematical proposition p , it is indeed true that p , and when they disbelieve p (or at least when they believe not- p) it is indeed false that p . There is, in other words, a striking correlation between mathematicians' mathematical beliefs (at least up to a certain level of complexity) and the mathematical truths. Such a striking correlation calls for explanation. But it doesn't seem that mathematical Platonists are in a position to offer any such explanation. The mathematical objects they believe in are abstract, and so causally inert, and so they cannot be causally responsible for mathematicians' beliefs; the mathematical truths Platonists believe in are supposed to be independent of mathematicians and their beliefs, and so mathematicians' beliefs aren't causally (or constitutively) responsible for the mathematical truths. Nor does there seem to be some third factor that is causally responsible for both. What we have here, then, is a striking correlation between two factors that Platonists cannot explain in any of the standard ways of explaining such a correlation—by invoking a causal (or constitutive) connection from the first factor to the second, or from the second to the first, or from some third factor to both. But without such an explanation, the striking correlation may just be too implausible to believe, and, Field concludes, so is mathematical Platonism.

Notice how elegant this way of stating the challenge is: There is no hidden assumption about the nature of knowledge, or of epistemic justification, or anything of the sort. There is just a striking correlation, the need to explain it, and the apparent unavailability of any explanation to the challenged view in the philosophy of mathematics.

Returning to normativity, then, here is the version of the epistemological challenge I suggest that we focus on: 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 robust realism is inconsistent with a response-dependence view of normativity, and so 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.¹⁷ Nor does there seem to be some third-factor explanation available to the realist. And so the metanormative realist is committed to an unexplained striking correlation, and this may just be too much to believe.¹⁸

This way of understanding the epistemological challenge to metanormative realism nicely satisfies the desiderata mentioned in Sect. 2.6 above: It is not an instance of the most general skeptical worries, because in many other domains, most notably with regard to beliefs about middle-sized physical objects, this challenge can be coped with rather easily (what explains the correlation between these truths and our beliefs about them is presumably the causal-perceptual relations between us and the physical objects around us); it poses a special problem to realism, because other metanormative views are immune to it (a point I return to below); it does not beg any question against the realist; and it elegantly avoids relying on contestable accounts of knowledge, justification, or any other epistemological notion.

3.2 Some implications

When thus understood, to which metanormative views does the epistemological challenge apply? The initial challenge—the need to explain the supposed correlation between normative truths and our normative beliefs—seems to me to be perfectly general. But some metanormative views can explain such correlation with relative ease. In particular, metanormative views according to which normative truths are causally efficacious—perhaps because reducible to or identical with natural facts¹⁹—can (in principle, at least) explain the correlation in the obvious way, by claiming that the normative truths are causally responsible for our normative beliefs. Indeed, epistemological worries of some sort (though perhaps not precisely the challenge as I suggest we understand it) about non-naturalist realism can be seen as a part of the motivation for naturalist realist views. I say that naturalist versions of realism can *in principle* explain the correlation between the normative truths and our normative beliefs, because whether or not such an explanation can be supplied depends on the details of the relevant naturalist account, and nothing here is obvious. It's just that there is no general reason to suspect that a naturalist realist *cannot* come up with such an explanation, whereas there does seem to be such a reason in the case of the robust realist.

¹⁷ There is, of course, a huge body of literature on the question whether normative—and, in particular, moral—truths explain anything worth explaining. I do not wish to enter this debate here. I take it that whatever your view on moral or normative *explanations*, the suggestion that normative truths *cause* our normative beliefs is at least a bit of a stretch. I return to this point below.

¹⁸ I think—but I am not sure—that Nagel (1997, pp. 76, 131) flirts with the claim that such correlations are to be taken as brute, and perhaps necessarily so.

¹⁹ Oddie is not a reductionist or a naturalist, but he thinks values are causally efficacious, and indeed his motivation for this (surprising, I think) claim is that this is his way of coping with some version of the epistemological challenge. See Oddie (2005, p. 181).

In a very different way, response-dependence views of normativity are also not threatened by the challenge above, because they can explain the correlation by invoking a causal or constitutive connection from our relevant dispositions or judgments to the normative truths.²⁰

This suggested understanding of the epistemological challenge, then, is primarily about reliability. But it's not about the possibility of reliable normative beliefs. Rather, it starts from the claim that such reliability is highly unlikely to be brute, unexplainable. Given the apparent unavailability of any such explanation to the robust metanormative realist, then, the challenge concludes that such realism is highly implausible. Now, this way of putting things may give rise to the following worry. Perhaps, it may be conceded, the realist has a problem with the *reliability* of our normative beliefs. But their epistemic *justification* can survive, because reliability is not necessary for justification. The suggested challenge, in other words, is externalist in the sense of the internalist–externalist debate with regard to epistemic justification, it is about matters of fact in the objective world. But justification, so the thought goes, is internal, supervening on what is cognitively accessible to the thinker by reflection alone.²¹ And so perhaps the realist should just acknowledge the loss of reliability while insisting that she can account for justified normative beliefs, and perhaps also for normative knowledge.

True, the challenge does start from an externalist perspective, asking about the explanation of our reliability in normative matters or of the correlation between the normative truths and our normative beliefs. But in order to appreciate its full strength it is important to see that the challenge can be “internalized”, and this even assuming that internalists are right about epistemic justification.²² It can, in other words, be shown to defeat (or perhaps undermine) even justification internalistically understood. Perhaps, if internalists are right about epistemic justification, I can justifiably form a belief using an unreliable belief-forming method. But even if this is so, I can't justifiably form a belief using what *I know* is an unreliable method. Perhaps, for instance, I can be justified in forming color-beliefs based on color-perception even in circumstances in which my color-perception is not reliable, and without having any beliefs *about* the reliability of color-perception in general or in these circumstances in particular. But in circumstances in which I *know* (or *justifiably believe*) that my color-perception is unreliable (say, because of misleading lighting) surely I lose whatever justification I may have had for my color-perception-based beliefs. Perhaps, in other words, reliability is not necessary for initial justification, but any such justification is at the very least defeated (or perhaps even undermined) when one *knows* (or justifiably believes) that one's belief

²⁰ *Ideal*-response-dependence views may not be in a better shape regarding the epistemological challenge compared to robustly realist views. See Dancy (1986, pp. 178–179).

How powerful is this epistemological challenge against expressivist, or noncognitivist metanormative views? I think that when push comes to shove, such views *are* response-dependence views, perhaps of a special kind. If so, the discussion in the text applies. But it is now common for expressivists to insist that their view is not committed to any kind of response-dependence. If such claims can be defended, then perhaps the challenge in the text forcefully applies to expressivist views as well.

²¹ This characterization of internalism comes from Pryor (2001, Sect. 3).

²² For a closely related point, see Joyce (2006, pp. 179–182).

or the method that formed it is in fact unreliable. Getting back to our challenge, then: If a brute correlation is too much to believe, and if no explanation is available to the realist, then the realist must conclude that there is after all no correlation between his normative beliefs and the normative truths. And knowing that, any (internalistically understood) justification he may have had for his normative beliefs is defeated (or undermined). And if so, he can then maintain his commitment to realism only at the price of a rather thoroughgoing skepticism about the normative. And while this is a possible position to have, stakes have certainly been raised: If the only way to be a realist is to deny epistemic justification for any normative belief (at least of the epistemologically informed, those who can run the argument in this paragraph in their heads), then antirealism gains significant ground.

This is also the reason why the realist cannot avoid the need to explain the correlation by denying that it exists: Knowing that there is no correlation between even his own normative beliefs and the normative truths, he can no longer hold these normative beliefs justifiably (and so, perhaps nor can he hold them as beliefs at all). The price of denying the correlation is skepticism. And it is without a doubt a serious price.

This version of the epistemological challenge also nicely explains the role of the resort to the faculty of rational intuition in this context—both the need for it, and why it will not help. A realist may feel that her only way of explaining the correlation between the normative truths and her normative judgments depends on there being some quasi-perceptual faculty that puts us in touch with the normative truths in roughly the same way our perceptual faculties put us in touch with empirical truths about nearby mid-sized objects. But a resort to such a faculty won't help the robust realist: For either this quasi-perceptual faculty is causal (like perception), putting us in causal relations with the normative truths, or it isn't. If it is, then the normative truths cannot be causally inert, as on Robust Realism they must be. And if this faculty is not causal, then for everything thus far said it is very hard to see how it can help in explaining the correlation that needs explaining.

No realist (to the best of my knowledge) has ever addressed the challenge thus understood.²³ For instance, in Brink's (1989, Chap. 5), Scanlon's (1998, pp. 64–72) and Shafer-Landau's (2003, Chaps. 10–12) rather elaborate epistemological discussions—all conducted in the general context of defending some kind of moral realism—nothing like this challenge is even mentioned, let alone addressed.²⁴ And the same is true of a fairly recent collection of papers on moral epistemology.²⁵ This

²³ Wedgwood (2006) mentions something very close to this challenge, but by the time he gets to his solution, it is no longer, I think, this challenge he is addressing, and in conversation he has made it clear that he is indeed interested in a different challenge (that of explaining how it is that we are justified in our moral beliefs). And while Huemer (2005, pp. 123–127) does seem to raise the challenge in the text, it is not clear to me how his sketch of an epistemology of the a priori—offered as a response to the challenge—copes with the challenge, at least as I understand it. For anti- (or non-) realist hints at this challenge, see Wright (1988, p. 25, footnote 36), Timmons (1990) (though Timmons raises it as a worry only for the combination of (naturalist, new-wave) realism and a coherentist theory of justification), and Gibbard (2003, p. 258) (though Gibbard raises it as a worry for his quasi-realism rather than for realism).

²⁴ For a similar accusation against (fellow-) realists, see Wedgwood (2006, p. 62, especially footnote 2).

²⁵ Sinnott-Armstrong and Timmons (1995).

may give rise to the worry that “my” epistemological challenge is simply a *different* epistemological challenge from more traditional ones, that it does not capture the traditional worry(ies). But I don’t think that this is so, for the following reason. My version of the epistemological challenge to realism makes other versions almost redundant, in the following sense: If a realist can cope with my version of the challenge, it’s not clear that she needs to worry about others. And if she cannot deal with my version of the challenge, she should find no comfort in her ability to cope with the other versions. The first of these conditionals is partly supported by the arguments in Sect. 2, showing the limited force of other ways of understanding the epistemological challenge to realism. Furthermore, given a plausible explanation of the correlation between our normative beliefs and the normative truths, it seems we are justified in believing that *there is* such a correlation, and this seems to suffice for our being justified in holding on to many of our normative beliefs, regardless of whatever else we can say about their justification. In the other direction, the second conditional—if the realist cannot cope with this challenge, she is not off the epistemological hook—has been established by the earlier point in this section, namely, that without a solution to this epistemological worry realism entails skepticism, and this regardless of what else the relevant realist can say about knowledge, or justification, or some such. But if an epistemological challenge is such that without solving it the realist is still in epistemological trouble (regardless of her replies to other epistemological worries), and that after solving it the realist is (almost completely) off the epistemological hook, then, I submit, it is a very good understanding of what underlies traditional epistemological challenges to realism.

For these reasons, then, the challenge as stated in Sect. 3.1 is how I suggest that we understand the epistemological challenge to metanormative realism, and especially to Robust realism. And we realists should not kid ourselves: Without a solution to this *strongest* version of the epistemological challenge we are not off the epistemological hook.

4 Interlude: Street’s Darwinian Dilemma

As already mentioned, no one (as far as I know) has articulated the epistemological challenge to metanormative realism in quite this (strongest) way. But in a recent paper Street (2006) comes very close. And it will prove useful, I think, to discuss some of her claims in detail.²⁶

Realists, Street argues, are committed to the response-independence of some normative truths. And if they are to avoid large-scale skepticism about our normative judgments, realists must think that often enough we get things right, that our normative judgments are at least often enough in line with the independent normative truths. But our normative judgments have been shaped to a large extent by evolutionary pressures. So realists are committed to a rather strong correlation between the independent normative truths and the normative judgments you can

²⁶ For a precursor, see Gibbard (2003, pp. 263–265). And for a related argument, see Joyce (2006, Chap. 6).

expect evolutionarily successful creatures to make, the normative judgments that were (roughly speaking) selected for. But how can the realist explain such a correlation? In answering this question, realists face a dilemma. If the realist fails to supply such an explanation, she is committed to what looks as unbelievable as a miracle—the fact that as sheer coincidence evolutionary forces pushed us in the direction of the independent normative truths. A commitment to such an amazing fluke is unacceptable. On the other horn of the dilemma, the only realist-friendly explanation of the correlation seems to be a tracking account of some sort, according to which our normative judgments and the evolutionary pressures shaping them causally track the independent normative truths. And Street rejects such an explanation on scientific grounds (alternative explanations of the emergence of our normative judgments are simply much better). Either way, then, the realist is in deep trouble.

Let me not pause here to attempt an evaluation of the scientific claims and evolutionary speculations²⁷ Street relies on, for—as Street herself concedes (2006, p. 155)—there is nothing essentially Darwinian about her Darwinian Dilemma. Replace any other (non-tracking) causal explanation of why we make the normative judgments that we do in fact make, and the realist will again find herself up against the problem of explaining strong correlations analogous to the one Street draws attention to. So any such causal story will do in order to raise the challenge of explaining a correlation of this kind, and the realist still needs a way out.²⁸

Once this is noticed, it is clear, I think, that Street's Darwinian Dilemma can be seen as a particular instance of the most general epistemological challenge to realism, properly understood. The general challenge—so I suggested above—is that of coming up with an explanation of a correlation between our relevant beliefs and the relevant truths. If we plug in the further premise that what explains our having the normative beliefs we do in fact have are broadly speaking evolutionary explanations, we get, in essence, Street's Darwinian Dilemma. So Street's Darwinian Dilemma is a particular instance of the general epistemological challenge to realism. For *naturalist* realist views, Street's dilemma may very well be an especially interesting or problematic particular instance of the general epistemological challenge.²⁹ For *robust* realists, though, it is quite safe to assume that Street's dilemma is *merely* a particular instance of the general epistemological challenge.

Noting the relation between Street's Darwinian Dilemma and the epistemological challenge in its strongest version is, I want next to argue, more than just an interesting observation. Street's way of thinking about the epistemological challenge is also a good way to think about how to solve it. I move, then, to the presentation of my suggested solution to the epistemological challenge (and Street's Darwinian Dilemma along with it).

²⁷ For Street's own characterization using this term, see Street (2006, pp. 112–113).

²⁸ Perhaps this is why Thomas Nagel—a robust realist (in some moods, at least)—suspects (1986, p. 145) that realism is incompatible with the availability of a purely naturalistic explanation of our normative judgments.

²⁹ I think that it is in this spirit that Copp (forthcoming) replies to Street.

5 How to cope with the epistemological challenge

In order to present my realist way of coping with the epistemological challenge, I start with two methodological remarks (in 5.1 and 5.2). I then proceed to present the solution (in 5.3).

5.1 Scoring points in an explanatory game

The challenge, it should be clear, is an explanatory one. Brute, unexplained and even unexplainable, correlations are not, after all, impossible, and so even if the realist has nothing to offer by way of an explanation of the correlation between the normative truths and our normative beliefs (or—in Street’s version—the normative truths and the normative beliefs which are likely to have been selected for) realism is not yet *refuted*. It’s just that we should opt for the best metaethical theory, the one that—perhaps among other desiderata—best explains whatever needs explaining. And so, if the realist cannot explain a striking correlation such as the one the epistemological challenge focuses on, realism loses plausibility points, perhaps to the point of unacceptability.

But then, once it is kept in mind that the game is an explanatory one and that the winner is going to be determined on plausibility grounds, the challenge to the realist—though still serious—seems at least manageable. For explanatory games are *holistic*³⁰ and *comparative* in nature. We are to opt for the theory that *best* explains *whatever* needs explaining, or perhaps—if we can restrict ourselves now to just choice among competing metanormative theories—whatever needs explaining in the vicinity of normative discourse and practice. And once *this* is noticed, the following three points are worth making.

First, the correlation that calls for explanation is not *all that* striking. We are not, after all, *that* good in forming and revising our normative beliefs.³¹ Furthermore, how strong the correlation is between the normative truths and our normative judgments depends on—among other factors—who “we” are. All humans who ever existed? All humans now existing? All well-educated, adult, humans? All those living in liberal democracies? All those with a college education? All philosophers? Just you and me? Just me?³²

The reason the realist could not plausibly deny the correlation to be explained, remember, was that doing so would result in skepticism. But then the correlation the realist must acknowledge—and the one it seems independently plausible to believe in—can be fairly weak, so long as it is strong enough to block the inference to skepticism. In other words, so long as the realist accepts that she and the like of her are at least somewhat more likely to get things right rather than wrong with regard to normative issues (or a specified subset of them), she can rationally and

³⁰ This is a point Shafer-Landau (2003, p. 234) emphasizes in our context.

³¹ Or, in Street’s version—the correlation between (what we take to be) the normative truths and normative judgments that are likely to have been selected for is not that strong. Consider the false but seemingly evolutionarily useful “The interests of others with whom we have no privileged genetic of reciprocal relations do not count at all.” and “The interests of non-human animals do not count at all”.

³² For sometimes “‘we’ means: ‘you and I, and I’m none too sure about you’.” (Lewis 1989, p. 84).

wholeheartedly if modestly hold on to her normative beliefs (or to those in the specified subset) and not deteriorate to skepticism. And here as elsewhere, the weaker the correlation to be explained, the less impressive the explanans has to be in order to shoulder the explanatory burden. (Compare: If the correlation between Josh's beliefs about the distant Nepalese village and the truths about it is rather weak, so that he's right more often than you would expect a random guesser to be, but not by much, then a very weak explanans would suffice to dispel the mystery—perhaps he just read a story in a magazine about the village. If the correlation is much stronger, though, a more impressive explanans will be needed.). I return to this point shortly.

Second, the explanandum is weaker also in another respect. Given a starting point of normative beliefs that are not too far-off, presumably some reasoning mechanisms (and perhaps some other mechanisms as well) can get us increasingly closer to the truth by eliminating inconsistencies, increasing overall coherence, eliminating arbitrary distinctions, drawing analogies, ruling out initially justified beliefs whose justificatory status has been defeated later on, etc. Now, it may of course be asked what explains the reliability of these reasoning mechanisms themselves.³³ But first, this question is much more general, because such reasoning processes occur everywhere, not just in revising and shaping our *normative* beliefs, so that this is not especially a problem for the metanormative realist. And second, evolutionary answers to this question—what explains the reliability of our reasoning processes—can rather easily (if speculatively) be thought of here: Precisely because they are applied everywhere, including in revising more transparently evolutionarily beneficial beliefs, the ability to employ reliable reasoning mechanisms may very well be itself evolutionarily beneficial, or the upshot of more basic reasoning abilities that are themselves evolutionarily helpful. It seems, then, that in our context we are entitled to assume that given good enough starting points our reasoning processes will in general get us closer to rather than further from the normative truths.³⁴ So what remains to be explained is not the full scope of the correlation between our normative beliefs and the normative truths, but just the correlation in starting points,³⁵ why, in other words, we are not too far off from the start, so that even reasonably good reasoning mechanisms can't get us to the kind of correlation between beliefs and truths we do in fact see.³⁶ Of course, the details here depend on the details of the relevant empirical science. My point here is then a rather modest one: It seems

³³ What I say in the concluding section below can be seen as a partial answer to this question.

³⁴ Think of the relevant reasoning processes as the analogues of proper conditionalization, and of the normative truths as the analogues of probability-truths. Then the point in the text is analogous to the following well-known result: So long as our assignments of prior probabilities are not too far off, and so long as we are exposed to sufficient evidence, proper conditionalization will get us increasingly closer to the probability-truths.

³⁵ A point acknowledged by Street (2006, p. 123), and emphasized also by Copp (forthcoming).

³⁶ Wedgwood (2006, pp. 79–80) mentions the starting-point problem for his reflective-equilibrium-like theory of justification, noticing that his suggested solution of (his version of) the epistemological challenge does not solve it. But again, the challenge he is primarily interested in is different from the one I am primarily interested in.

plausible to speculate that some reasoning (and perhaps other) mechanisms can get our normative beliefs closer to the normative truths, so long as the starting points are not too far off. To the extent that this speculation can be supported, what the epistemological challenge requires that realists explain is this weaker explanandum, namely, that the relevant starting points are not too far off.

Finally, the comparative nature of the explanatory game we are engaged in—and of explanatory games in general—dictates that even if realism is locally implausible, it may still be globally better than all competing metanormative theories. This is not, of course, the place to list all of the problems with antirealist views (or indeed with Street's favorite (constructivist) version thereof³⁷). But it is important to keep in mind that the realist's task is not to show that realism does better than competing metanormative theories *in every respect, with regard to every problem*. If the realist can show that the epistemological challenge is one on which realism doesn't lose *too many* plausibility points, and if realism compensates for this loss elsewhere (say, in supplying a much more plausible account than any alternative metanormative theory of normative deliberation³⁸), this should be enough. And this means that flukes are not as theoretically frightening as Street, for instance, seems to think they are—so long, that is, as belief in them is well-motivated by the advantages of realism elsewhere. Of course, it is still important to try and show that the flukes realists are committed to are not all that fluky—and this, indeed, is what I am about to do. But that some flukes remain should not be a cause for too much distress.

5.2 How correlations get explained

That said, I agree that the (somewhat) striking correlation between the purported independent normative truths and our (perhaps selected-for) normative beliefs does call for explanation, and that declaring it a fluke is something we realists should avoid if at all we can. How, then, can the realist explain this correlation? Indeed, how can correlations *in general* be explained?

The obvious ways (and the ones Street focuses on) are two: If the correlated factors are A and B, then (roughly speaking) either A-facts are somehow (causally, constitutively, or both) responsible for the B-facts, or the B-facts are responsible for the A-facts. And applied to our case: The correlation would be explained if either the normative truths were responsible for our normative beliefs, or our normative beliefs were responsible for the normative truths. But, as argued in Sect. 3.1 above, it is exactly these two kinds of explanation that are unavailable to the robust realist. So the realist has to look for another way of explaining correlations in general, and the one the epistemological challenge focuses on in particular.

The thing to look for is a third-factor explanation. For it is possible that the explanation of a correlation between the two factors A and B is in terms of a third

³⁷ For a general criticism of all response-dependence views that involve some idealization—of which Street's seems to be one—see Enoch (2005). And for a general critique of constructivism, see Enoch (manuscript).

³⁸ A line I pursue in Enoch (2003, 2007).

factor, C, that is (roughly speaking) responsible both for A-facts and for B-facts.³⁹ Pre-established-harmony explanations, for instance, are always of this sort. And the realist-friendly explanation of the correlation I am about to offer is exactly such a third-factor explanation, or indeed a (godless) pre-established-harmony type of explanation.

5.3 Survival is good—an evolutionary speculation

Here, finally, thinking about Street's Darwinian Dilemma can help the realist in coping with the epistemological challenge more generally.

Assume that survival or reproductive success (or whatever else evolution "aims" at) is at least somewhat good. Not, of course, that it is always good, or that its positive value is never outweighed by other considerations, or even that it is of ultimate or of intrinsic value, or anything of the sort. Furthermore, I am not asking you to assume that the evolutionary "aim" is of value *because* it is the evolutionary aim. All I will be relying on is the assumption that survival (or whatever) is actually by-and-large better than the alternative.⁴⁰

Selective forces have shaped our normative judgments and beliefs, with the "aim" of survival or reproductive success in mind (so to speak).⁴¹ But given that these are by-and-large good aims—aims that normative truths recommend—our normative beliefs have developed to be at least somewhat in line with the normative truths. Perhaps somewhat ironically—because Street thinks evolutionary considerations serve to ground the epistemological challenge to realism—evolutionary considerations can help the realist cope with the challenge. Given that the evolutionary "aim" is good, the fact that our normative beliefs have been shaped by selective forces renders it far *less* mysterious that our normative beliefs are somewhat in line with the normative truths. This is so, then, neither because the normative truths are a function of our normative beliefs, nor because our normative beliefs causally track the normative truths, but because our normative beliefs have

³⁹ Street (2006, p. 134) is explicit about there being only two types of possible explanations of such correlations—the two mentioned earlier in the text. But in her discussion of pain (Sect. 9) she may be implicitly flirting with a third-factor kind of explanation. I briefly return to this part of her discussion below.

Throughout Copp's (forthcoming) discussion of Street's Darwinian Dilemma, he speaks of "the tracking thesis" as something realists must endorse if they are to avoid the skeptical horn of the dilemma. But talk of tracking here is misleading. "Tracking" is naturally understood as a causal term. If so, the availability of third-factor explanations shows that the skeptical horn can be avoided without endorsing a tracking account. If "tracking" is understood as not necessarily causal (perhaps merely indicating some counterfactual relations), then Copp may be including under the same term causally-tracking and third-factor explanations. Indeed, at the end of the day it is not entirely clear to me whether Copp's own way out—that relies on his society-centered version of moral realism—is an instance of the former or the latter. His talk of "a close relative of the tracking account" seems to suggest that perhaps his solution too can be understood as a third-factor explanation of sorts, though of course not the one I am about to offer.

⁴⁰ These clarifications should make it clear that the assumption I will be relying on is significantly weaker than those considered by Gibbard (2003, p. 264) and Timmons (1990, pp. 107–108). If I am right in what follows, the weaker assumption suffices for my purposes.

⁴¹ Of course, evolution has neither a mind, nor an aim in mind. Talk of the evolutionary "aim" in the text is meant as shorthand for the usual respectable, non-teleological, evolutionary way of putting things.

been shaped by selective pressures towards ends that are in fact—and quite independently—of value. The connection between evolutionary forces and value—the fact that survival is good—is what explains the correlation between the response-independent normative truths and our selected-for normative beliefs. The fact that (roughly speaking) survival is good pre-establishes the harmony between the normative truths and our normative beliefs.

The causal influence of selective forces only directly “pushes” us in the direction of having *evolutionarily beneficial* beliefs, not necessarily true ones. But here as elsewhere, the two may be systematically related.⁴² For we are the kind of creatures whose actions seem to be closely related to their normative beliefs about how they should act, or how it would be good to act, or what consequences it would be good to bring about. Our mental and motivational setup seems to include a mechanism roughly like Gibbard’s (1990, Chap. 4) “normative governance” mechanism. And this completes the explanatory story needed here: Survival (or whatever) is good; so behaving in ways that promote it is (pro tanto) good; but one efficient way of pushing us in the direction of acting in those ways is by pushing us to believe that it is good to act in those ways. And in fact, as we have just seen, it *is* good so to act. So the normative beliefs this mechanism pushes us to have will tend to be true.

Now, unlike the paradigmatic pre-established harmony kind of explanation—of the mind-body harmony, for instance—my explanation of the correlation between the normative truths and our normative beliefs is not causal in both directions. Perhaps the evolutionary “aim”—whatever exactly it is—causally shapes our normative beliefs, but the fact that the evolutionary “aim” is of value does not *causally* shape the normative truths. It seems more appropriate to say that it is related to them in some constitutive way, a way the details of which depend on the details of your favorite first-order, normative theory. For instance, the fact that survival is good is plausibly related in coherence relations to many (though perhaps not all) other normative truths, like that pain is pro-tanto bad, that some close relationships are good to have, etc. Relations of this nature between that survival is good and many other normative truths—though not causal—still allow me, I think, to explain the correlation between the normative truths and our normative beliefs in a way that resembles pre-established accounts sufficiently to merit the name.

To see this more clearly, think about the following example: What explains the correlation (if indeed there is one) between giving rise to strongly affectionate feelings and having a poor sense of time? Here’s one possible answer: Young children are cute. Being a young child explains—indeed, perhaps causes—having a poor sense of time. And of course, being cute is closely though perhaps not causally related to giving rise to strongly affectionate feelings. The fact that young children are cute, then, pre-establishes the harmony between giving rise to strongly affectionate feelings and having a poor sense of time.

Needless to say, I have no idea whether this explanation actually works (or whether the phenomenon it is supposed to explain is actually a real phenomenon). All that is crucial for me is that it *could* work, and that its *structure* is exactly similar to that of the explanation I am really after. Selective forces have causally shaped our

⁴² Discussions with Pete Graham and Sigrún Svavarsdóttir helped me with this paragraph.

normative beliefs; that survival is good is (non-causally but closely) related to many normative truths; and so that survival (or whatever the evolutionary “aim” is) is good explains the correlation between our normative beliefs and the normative truths. Similarly, being a young child causally inclines one to have a poor sense of time; cuteness is (perhaps non-causally but closely) related to giving rise to strong affectionate feelings; and so that young children are cute explains the correlation between giving rise to strong affectionate feelings and having a poor sense of time.

Now, given how qualified we should be about attributing goodness (or perhaps some other value) to survival and reproductive success, this explanation could not explain a very strong correlation between the normative truths and our normative beliefs. But as noted above the correlation in need of explanation is not that strong anyway. And given the effect of reasonably reliable reasoning mechanisms starting from reasonably good starting-points, there is even less that calls for explanation.⁴³ With these points in mind, it becomes clear that the fact that survival is good can play its role in this explanatory story even when it is understood rather loosely, as indeed it should be. Notice, for instance, that the children-are-cute explanation works well even if not *all* young children are cute, even if puppies are cuter than young children, and so on. In other words, the claim that children are cute can be understood rather loosely and still play its role in the explanation of the relevant correlation. Similarly, then, the fact that the assumption I relied on—that the evolutionary “aim” is good—is weak in the ways described above does not prevent it from playing an analogous role in explaining the correlation that the epistemological challenge calls upon the realist to explain.

How satisfactory, then, is this explanation of the correlation between the normative truths and our normative beliefs?⁴⁴

One natural worry is that the normative assumption my explanation relies on—that the evolutionary “aim” is good—may be false. I have already partly answered this concern, by emphasizing how loosely this normative assumption can be understood while still successfully playing its role in the relevant explanation. But I want to note here another important point⁴⁵: All the explanation really requires is that *our* survival or reproductive success is good, or that those of our ancestors are. After all, it is only the selective forces that acted on us and them that have played a role in shaping our normative beliefs, about whose reliability we are asking. Even if there are some organisms whose survival or reproductive success (or whatever) is of no even pro tanto value, then, this is consistent with the normative assumption my explanation relies on. And when it comes to creatures like us and our fairly close

⁴³ In particular, perhaps such reasoning mechanisms can explain—given reasonably good starting points—how “beings like us would be good judges of ultimate worth, even in those cases where ultimate worth comes apart from maximizing the long-run reproduction of one’s genes” (Gibbard 2003, p. 265), as indeed Gibbard himself suggests later on (2003, p. 266).

⁴⁴ Notice that the use to which I put (speculative) evolutionary considerations is not that of (directly) justifying our normative beliefs. Rather, it is that of explaining the correlation between the normative truths and our beliefs, thereby possibly defeating a defeater of the justification of our normative beliefs. To an extent, then, I need not disagree with Nagel when he (1997, p. 136) says: “This means that the evolutionary hypothesis is acceptable only if reason does not need its support. At most it may show why the existence of reason need not be biologically mysterious”.

⁴⁵ I am indebted here to comments from Ralph Wedgwood and Sigrún Svavarsdóttir.

ancestors, the claim that their survival and reproductive success is of value gains much plausibility, I think, from the observation that survival (or some such) is at the very least good *for* the creature surviving, or *for* a close group of relatives, or something of this kind. Again, this may not be true of creatures in general (some creatures just do not have interests, and so presumably nothing is good for them). But when it comes to us and to creatures like us, this claim seems very hard to deny. Furthermore, it seems almost undeniable that there are close (if not obvious) connections between being good *for* someone and being good. So it's very hard to deny that the survival (or some such) of creatures like us is good in the very loose sense needed for my explanation of the correlation between the normative truths and our normative beliefs to go through.

Another obvious challenge is that this explanation itself seems to invoke what may be thought of as a miracle. For isn't it an amazing fluke that whatever evolution "aims" at happens to also be good? And isn't this itself something that calls for explanation, an explanation that the realist is not in a position to offer?⁴⁶ One way of putting the point—following Copp (forthcoming)—is in terms of counterfactual robustness: On the story just told, had evolution "aimed" at something else, something that is not of value or perhaps is even of negative value, our normative beliefs would have been systematically mistaken. So doesn't it follow that the story just told, far from showing how the realist can avoid commitment to miraculous correlations, *relies* on a miracle?

It is not completely clear that the thought that a miracle remains here can be made fully coherent.⁴⁷ For what would have to be the case for this "miracle" not to occur? The evolutionary "aim" would have had to not be of any value. And how could that be? Fundamental normative truths are presumably necessary in a fairly strong sense, or at the very least so we are entitled to assume in the context of critically evaluating the epistemological challenge to Robust Realism. So the main way in which the evolutionary "aim" (which is actually of value) could have failed to be of value is if evolution had a very different "aim". But it's not clear what to make of this suggestion: For surely, it's not contingent that evolution has something to do, for instance, with survival and reproductive success rather than their opposites.

Nevertheless, I am willing to concede, at least for the sake of argument, that *some* miracle remains. But this remaining miracle does not place a particularly heavy burden on the shoulders of the metanormative realist. For first, of the striking correlation we started with (between numerous normative truths and numerous normative beliefs) we are now left only with the one-time "correlation" between whatever evolution "aims" at and the good. Given the rules of the explanatory game and how we can score points in it, it is hard to deny that significant progress has been made. Think about the analogous explanation again. That young children are cute may be considered a miracle. Perhaps it isn't—perhaps, in other words, we can

⁴⁶ Street (2006, p. 150) makes an analogous claim against the pain-explanation of the correlation she is interested in. And Copp (forthcoming) is explicit about his attempt to avoid a commitment to any such remaining miracles.

⁴⁷ I thank Ralph Wedgwood for relevant discussion here. For the claim in a very close context that an explanatory challenge remains, see Huemer (2005, p. 123).

come up with explanations of this fact (if it is a fact), indeed perhaps we can think of partly evolutionary explanations that are relevant here. And if we can, this will make the explanation of the correlation between giving rise to strongly affectionate feelings and having a poor sense of time even better. The important thing to note here, though, is that *even in the absence of such a further story*, even while still treating the cuteness of young children as a miracle, *still* explanatory progress has been made. Analogously, then, even if we have to think of the goodness of the evolutionary “aim” as a miracle, still explanatory progress has been made.

Second, how surprising is the miracle that remains? It is indeed true that had the causal forces shaping our intellectual and other normative faculties been very different, had they “aimed” at things that are of no value at all or that are of disvalue, we would have been systematically mistaken in our normative beliefs. And we are indeed epistemically lucky that this (presumably) isn’t the case, just as we are epistemically lucky that our other intellectual (and perceptual) faculties have been shaped to be reasonably reliable. We are, after all, epistemically lucky to have evolved in an environment in which having by-and-large true beliefs is presumably conducive to survival and reproductive success. So yes, some brute luck may remain. But it is the same kind of luck that is present everywhere else. And though in some moods I too feel the call to explain it, still the pressure to supply such an explanation is not particularly strong, nor does it pose a particular problem for the realist.

As already noted, then, the explanation suggested is not fully counterfactually robust. Still, it would be a mistake to suggest that it lacks *any* counterfactual robustness. For the explanation still works in a large number of fairly close (but non-actual) possible worlds. Had the selective forces worked only somewhat differently, or had the evolutionary “aim” been different but still of value, still the starting points of our normative beliefs would have been close enough to the truth for our normative beliefs to be (somewhat) correlated with the normative truths. The possible worlds in which—on the suggested story—there is no correlation between our normative beliefs and the normative truths are quite far: These are worlds in which, for instance, evolution “aims” (only or mostly) at suffering or humiliation, or worlds in which survival (or whatever) has absolutely no value. And though we can think about the fact that none of these worlds is the actual one as a miracle, it seems like a rather small miracle. Counterfactual robustness, then, comes in degrees. Perhaps more by way of counterfactual robustness can be hoped for. But *some* (significant) counterfactual robustness is satisfied by the suggested explanation of the correlation between the normative truths and our normative beliefs.⁴⁸

⁴⁸ In discussing knowledge in Sect. 2.4, I said that Gettier cases do not pose a special problem for normative beliefs. But at this point the related following worry may arise (for which I thank Mark Van Roojen): One common way of stating the problem Gettier was on to is in terms of luck: What we need for knowledge is *unluckily* justified true belief. And it may be thought that the counterfactual robustness of my explanation of the correlation between our normative beliefs and the normative truths—though perhaps it suffices to defeat the possible defeater of our justification in holding those beliefs—nevertheless does not suffice for Gettier. Perhaps, in other words, our normative beliefs, though correlated with the normative truths, are *luckily* so correlated, and so do not count as knowledge.

I am not sure how best to understand the anti-luck intuition, and so how its best version interacts with the considerations in the text here. But let me acknowledge that more discussion is needed here. And, of

What is the scope of the normative truths such that the correlation between them and our relevant normative beliefs is explained by the suggested story? All those normative truths that are closely enough related to the goodness of whatever it is evolution “aims” at. In order to say more about this scope, then, we need a first-order normative theory, a theory of (among other things) the relations between different normative concepts and judgments. But even without presenting such a theory, it seems to me clear that coherence relations between normative beliefs will take us quite far, and so that it is rather safe to speculate that the scope of the correlation is rather wide. And once the effect of reasoning mechanisms—including reasoning by analogy from other parts of the normative domain—is taken into account, the scope of the explained correlation becomes wider still, indeed perhaps maximally so.

In his closely related discussion, Gibbard (2003, p. 263) puts forward another, related, requirement for the sought-after explanation—it must explain “why, in such matters, I would tend to be *correct as such*.”, that is, without relying in this explanation on the substantive truths with regard to which our reliability is to be explained (for instance, that survival is good). But this requirement is either too strong to be plausible, or weak enough to be satisfied by the suggested explanation. If Gibbard requires that our reliability be explained completely independently of the relevant substantive truths, the requirement cannot be defended. Even with regard to perceptual beliefs, for instance, if facts about middle-sized objects were very different, we would not have been reliable. So the requirement must be understood in a weaker way, as saying (roughly) that the explanation shouldn’t rely on too many, too specific substantive truths of the relevant domain. But thus understood, it seems to me the explanation suggested here passes the test.

Let me not give the impression that this suggested way of coping with the epistemological challenge is ideal. Indeed, because of the (perhaps) remaining small miracle, perhaps the realist loses some plausibility points here. But not, it seems to me, too many. Given the rules of the explanatory game—and given the other, independent, advantages of realism (which I don’t discuss in this paper)—this conclusion may be enough for the realist to have satisfactorily addressed the epistemological challenge.

6 Generalizing

How general is the epistemological challenge to realism? And how general is my suggested way of coping with it?

As already noted, the need to explain the correlation between a specified class of truths and our relevant beliefs is quite global, since if we cannot explain this correlation in a given domain, we may conclude that no such correlation is likely to exist in that domain, and so that we are completely unreliable in our relevant beliefs,

Footnote 48 continued

course, even if there are here problems for normative knowledge, they don’t extend to the justification of normative beliefs.

a conclusion that will defeat—and perhaps even undermine—any justification for the relevant beliefs, and therefore also knowledge in the relevant domain.

The abstract challenge, then, is quite general, but it is not equally difficult to cope with across contexts. First, perhaps with regard to some classes of truths and beliefs, a brute correlation will not be too much to believe. If there are such cases, the problem with regard to them is, of course, solved. Second, and more importantly, the available explanations differ across contexts. In the paradigmatic case of, say, every-day beliefs about nearby middle-sized objects, the correlation gets explained in a rather obvious causal way: There are causal (perceptual) relations between the relevant truths (or facts, or properties, or whatever you think the causal relata are) and our (perceptual and other) beliefs.⁴⁹ And there may, of course, be cases where the explanation goes the other way, that is, where the relevant truths are determined by us and our responses.

This suggests the following conclusion: The epistemological challenge as understood in this paper is at its strongest when applied to discourses purportedly consisting of response-independent truths that are knowable—if at all—a priori.⁵⁰ In these cases, a causal explanation from truths to beliefs can't work (because of the apriority), and a causal or constitutive explanation from beliefs to truths cannot work (because of the response-independence). Perhaps this generalization should not come as a surprise, our discussion having started with the epistemological challenge to mathematical platonists and metanormative realists.

Does my suggested solution, then, generalize to all such discourses? The answer, I suspect, depends on the details.⁵¹ The suggestion my explanation utilized—that the evolutionary “aim” is good—may not have plausible analogues in other cases. And perhaps there are other third-factor explanations that are available in other domains, ones the analogues of which I could not think of with regard to normativity. So let me tentatively conclude with the following remark: A priori, response-independent domains for which something analogous to the explanation given in this paper can be made to work are off the epistemological hook. Those for which nothing like this story can explain the (supposed) correlation between the

⁴⁹ And there is, presumably, an evolutionary story backing up this causal story. See, for instance, Gibbard (2003, p. 255).

⁵⁰ Schechter (2006) notices this point. McGinn (1997, p. 58) notices that the epistemological challenge to realism is a particular instance of the general problem of a priori knowledge (but he thinks it cannot be solved by creatures like ourselves).

⁵¹ In his PhD dissertation (2006), and in work in progress based on it, Joshua Schechter discusses (among other things) the epistemological challenge to (what he calls) the objectivity of logic, and he understands the challenge in a way precisely analogous to the way I understand the epistemological challenge to metanormative realism in this paper. Furthermore, he pursues an evolutionary way of dealing with it. Indeed, my thinking about these matters was heavily influenced by my exchanges with him, and I am grateful to him for that. But it may be helpful here to note an important difference: Schechter discusses an evolutionary explanation that utilizes the evolutionary advantages of our employing reliable deductive rules (and so of forming true logical beliefs). I do not argue that having true normative beliefs is evolutionarily advantageous (for a sketch of an argument along these lines, see Huemer (2005, pp. 218–219)). My argument rather is a pre-established-harmony kind of explanation. This difference between these two explanations is an instance of the point made in the text: The details of the appropriate ways of dealing with the epistemological challenge (properly understood) may very well differ with context.

relevant truths and the relevant beliefs are still in trouble. Whether they can cope with it in some other way is a matter for further thought.

Acknowledgements I worked on this paper as a fellow of the Institute for Advanced Studies in Jerusalem. I am grateful for the Institute's support. I would also like to gratefully acknowledge the support of the Israel Science Fund. I presented this paper at the IAS in Jerusalem, and at the research workshop at the Bar Han University philosophy department, and I thank the participants for the helpful discussions that followed. For helpful conversations and comments on previous versions I am grateful to Erez Aloni, Hagit Benbaji, Pete Graham, David Heyd, Yair Levy, Ofer Malcai, Josh Schechter, Mark Schroeder, Yonatan Shemer, Sigrún Svavarsdóttir, Mark van Roojen, Ralph Wedgwood, and Ruti Weintraub.

References

- Audi, R. (1995). In W. Sinnott-Armstrong & M. Timmons (Eds.), *Intuitionism, pluralism, and the foundations of ethics* (pp. 101–136). Reprinted in his *Moral knowledge and ethical character* (pp. 32–65). Oxford: Oxford University Press (1997).
- Brink, D. O. (1989). *Moral realism and the foundations of ethics*. Cambridge: Cambridge University Press.
- Copp, D. (forthcoming). Darwinian skepticism about moral realism. In S. Landau (Ed.), *Oxford studies in metaethics* (Vol. 4).
- Cuneo, T. (2007). *The normative web: An argument for moral realism*. Oxford: Oxford University Press.
- Dancy, J. (1986). Two conception of moral realism. *Proceedings of the Aristotelian Society*, 60(Suppl), 167–188.
- Enoch, D. (2003). *An argument for robust metanormative realism*. NYU Dissertation, available at: <http://www.law.huji.ac.il/upload/Thesis.doc>.
- Enoch, D. (2005). Why idealize? *Ethics*, 115(4), 759–787.
- Enoch, D. (2007). An outline of an argument for robust metanormative realism. *Oxford Studies in Metaethics*, 2(2007), 21–50.
- Enoch, D. (forthcoming). How is moral disagreement a problem for realism? *The Journal of Ethics*.
- Enoch, D. (manuscript). Can there be a global, interesting, coherent constructivism about practical reason?
- Field, H. (1989). *Realism, mathematics and modality*. New York: Blackwell.
- Gibbard, A. (1990). *Wise choices, apt feelings*. Cambridge: Harvard University Press.
- Gibbard, A. (2003). *Thinking how to live*. Cambridge: Harvard University Press.
- Huemer, M. (2005). *Ethical intuitionism*. New York, NY: Palgrave MacMillan.
- Joyce, R. (2006). *The evolution of morality*. Cambridge: MIT Press.
- Lewis, D. (1986). *On the plurality of worlds*. Oxford: Blackwell.
- Lewis, D. (1989). Dispositional theories of value. *Proceedings of the Aristotelian Society*, 63(Suppl), 113–137; reprinted in his *Papers in ethics and social philosophy* (pp. 68–94). Cambridge: Cambridge University Press (2000).
- Liggins, D. (2006). Is there a good epistemological argument against platonism? *Analysis*, 66(2), 135–141.
- Mackie, J. (1977). *Ethics: Inventing right and wrong*. London: Penguin Books.
- McGinn, C. (1997). *Ethics, evil, and fiction*. Oxford: Clarendon Press.
- Nagel, T. (1986). *The view from nowhere*. Oxford: Oxford University Press.
- Nagel, T. (1997). *The last word*. Oxford: Oxford University Press.
- Oddie, G. (2005). *Value, reality and desire*. Oxford: Oxford University Press.
- Pryor, J. (2001). Highlights of recent epistemology. *British Journal for the Philosophy of Science*, 52, 1–30.
- Sayre-McCord, G. (1995). Coherentist epistemology and moral theory. In W. Sinnott-Armstrong & M. Timmons (Eds.), *Moral knowledge? New readings in moral epistemology* (pp. 137–189). Oxford: Oxford University Press.
- Scanlon, T. M. (1998). *What we owe to each other*. Cambridge, MA: Harvard University Press.
- Schechter, J. B. (2006). *Two challenges to the objectivity of logic*. Dissertation, New York University.
- Shafer-Landau, R. (2003). *Moral realism: A defence*. Oxford: Oxford University Press.

- Sinnott-Armstrong, W. (2006). *Moral skepticisms*. Oxford: Oxford University Press.
- Sinnott-Armstrong, W., & Timmons, M. (Eds.). (1995). *Moral knowledge? New readings in moral epistemology*. Oxford: Oxford University Press.
- Street, S. (2006). A Darwinian Dilemma for realist theories of value. *Philosophical Studies*, 127, 109–166.
- Timmons, M. (1990). On the epistemic status of considered moral judgments. *The Southern Journal of Philosophy*, 29(Suppl), 97–131.
- Waldron, J. (1992). The irrelevance of moral objectivity. In R. George (Ed.), *Natural law theory: Contemporary essays* (pp. 158–187). Oxford: Clarendon Press.
- Wedgwood, R. (2006). How we know what ought to be. *Proceedings of the Aristotelian Society*, 106(1), 61–84.
- Wright, C. (1988). Moral values, projection, and secondary qualities. *Proceedings of the Aristotelian Society*, 62(Suppl), 1–26.