



<http://www.diva-portal.org>

Preprint

This is the submitted version of a chapter published in *Moral Skepticism: New Essays*.

Citation for the original published chapter:

Tersman, F. (2017)

Moral Skepticism and the Benacerraf Challenge.

In: Diego Machuca (ed.), *Moral Skepticism: New Essays* New York: Routledge
Routledge Studies in Ethics and Moral Theory

N.B. When citing this work, cite the original published chapter.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-325227>

This is a pre-peer-reviewed version of the article. It will be published in final form in D. Machuca (ed.), *Moral Skepticism: New Essays* (Routledge Studies in Ethics and Moral Theory), New York: Routledge, 2017.

Moral Skepticism and the Benacerraf Challenge

Folke Tersman, *Uppsala University*

1. Introduction

Platonists in the philosophy of mathematics believe that mathematical beliefs posit objects (numbers, functions, sets, etc.) that are mind- and language-independent, causally inert, and exist beyond time and space. The Benacerraf challenge concerns how to reconcile that belief with the existence of mathematical knowledge. After all, if numbers belong to a realm that is entirely cut off from us in the indicated ways, how could we possibly have any access to the truth about their features and relations? Some suspect that there is no viable way of answering that question, which in turn is taken to be a reason to abandon Platonism and to adopt some competing view, such as, for example, a constructivist account of mathematical entities.

There are other forms of realism, in other philosophical domains, that make similar claims about the facts and properties they posit as those Platonists make about numbers. Those positions are therefore held to be vulnerable to analogous challenges. One candidate is ethical non-naturalism. That ethical non-naturalists, and moral realists in general, face difficulties in accounting for the existence of knowledge in the target area is the conclusion of several well-known arguments. This is the upshot of some so-called “evolutionary debunking arguments” (see, e.g., Joyce 2013 and Street 2006), but also of certain versions of the

argument from moral disagreement (see Tersman 2006, chapter 4, as well as of Gilbert Harman's contention that we have no reason to believe in objective moral facts because they are never assumed by the best explanation of anything observable (Harman 1977, chapter 1). However, while many of those more familiar arguments appeal to complex empirical assumptions, the point of the Benacerraf challenge is to show that the assumptions the target position makes about the facts it posits *by themselves* rule out the existence of knowledge. By applying that reasoning to ethics, the critics of non-naturalist moral realism may thus hope to avoid criticism of the kind the invocation of controversial empirical claims inevitably invites and to reach their conclusion in a more comfortable way.

The aim of this paper is to explore whether the Benacerraf challenge really does provide the critics of (non-naturalist) moral realism with such extra resources. Is there a compelling version of the Benacerraf challenge that pertains to ethical non-naturalism and that is distinct from the more familiar ones just mentioned? That's the central question in this paper. The upshot is that this question should be answered in the negative. The moral version of the Benacerraf challenge must (unlike the mathematical one) be supplemented by considerations that are external to it and does ultimately not constitute a distinct and separate objection that is independent of the more familiar ones.

2. Field's version

The Benacerraf challenge owes its name to Paul Benacerraf, who gave a first formulation of it in his seminal paper "Mathematical Truth" (Benacerraf 1973). Benacerraf's own version is widely seen as unsatisfying, however, because he explicitly presupposes a causal theory of knowledge whose plausibility can be questioned on independent grounds. I shall therefore focus on the more recent version that is due to Hartry Field.

Unlike Benacerraf, Field does not invoke the concept of knowledge. Instead, he takes the challenge to concern the “reliability” of our mathematical beliefs. According to Field, the problem with the claims Platonists make about the nature of mathematical entities is that they seem to exclude any viable explanation of their reliability.

Field assumes that to explain the reliability of our beliefs in an area is to explain an “actual correlation” between our beliefs in that area and the relevant facts. That is, with regards to mathematical beliefs, it is to explain that the following schema

If we accept ‘p’ then p

“(and a partial but hard to state converse of it) holds in nearly all instances, when ‘p’, in this case, is replaced by a mathematical sentence.” (Field 1989, 26).¹

In a sense, then, to explain the reliability of our mathematical beliefs is to explain why they are mostly true. Or, better, it is to explain why we have managed to arrive at beliefs in this area that are mostly true (since it is not the truth of their *contents* that is to be explained).² Field sometimes rather formulates himself a bit differently. What is to be explained (or shown by our account) is that, given that we have in fact arrived at mathematical beliefs that are mostly true, this is not a result of a “huge coincidence” or sheer luck.

Now, on Field’s construal of the argument the Benacerraf challenge potentially gives rise to, a central premise is the claim that the principled inability to provide such an explanation

¹ Field focuses on the beliefs held by *mathematicians* and not those that are held by people in general (“us”), but that makes no difference here. I shall continue to write about “our” mathematical beliefs.

² By a “belief”, I refer to a certain psychological state or an instance of a certain psychological attitude. Each such state has a content that is constituted by a proposition, or so I assume, and people are said to “share” beliefs to the extent that they have beliefs with the same contents.

tends to undermine the justification or epistemic permissibility of those beliefs.³ The underlying intuition is illustrated by Field's famous story about a remote village in Nepal. Suppose that we initially have various beliefs about this village, such as that an odd number of its inhabitants are born on a Friday, perhaps because we have been told so by someone who claims to have lived there. Suppose also that we subsequently learn that the person in question has never in fact been in any contact with the village, and that there is no alternative account of the reliability of our beliefs that holds any promise. In that case, it would be wise of us to drop the beliefs and to adopt an agnostic attitude towards claims about the goings-on in the village. The analogy is stressed in the following passage:

Consider the Benacerraf worry about (certain kinds of) platonism in the philosophy of mathematics. The worry is that the same reasons that would lead us to advise against having beliefs about the happenings in a remote village in Nepal, when one has reason to think that there is no possible explanation of the reliability of those beliefs, should equally lead us to advise against having beliefs about mathematical entities platonistically construed, given that it appears that there is no possible explanation of the reliability of those beliefs. (Field 2009, 289)

The quote suggests the following formulation of the epistemic idea that Field has in mind:

A If it is in principle impossible to explain the reliability of a set of beliefs they are unjustified and we should abandon them.⁴

³ Thus, at one juncture, he writes that "a principled inability to [explain the reliability of some set of belief] tends to undermine the justification of those beliefs" (1996, 377).

⁴ In some passages, Field formulates the antecedent of the principle so that it merely implies that we have *some* reason to give up the beliefs, such as when he writes that the absence of an explanation of the relevant kind "tends" to undermine the justification of the

The second step of the argument is the claim that Platonism excludes the possibility of a viable explanation of the reliability of our mathematical beliefs. Thus:

B Given Platonism, it *is* in principle impossible to explain the reliability of our mathematical beliefs.⁵

And from A and B we can derive:

C Given Platonism, we should abandon our mathematical beliefs.

How does **C** in turn create a problem for Platonism? It may seem tempting to go on as follows. The claim that no mathematical belief is epistemically permitted is highly counterintuitive. So the fact that Platonism implies that they aren't represents a theoretical cost. It might perhaps have advantages over its competitors that compensate for this cost. But in the absence of such advantages it is a reason to reject Platonism.

However, although this construal may square with how some commentators conceive of the argument it does not capture Field's conception. For Field *agrees* with the conclusion that we should drop our mathematical beliefs, in the sense that we should cease to take the

target beliefs. Field also offers formulations given which it is the "appearance" of the impossibility of providing such an explanation, or our possession of reasons for thinking that this is impossible, that are the crucial considerations rather than whether it *in fact* is impossible. (Field (1989, esp. 25-30, 230-9), Field (1996) and Field (2005).) These variations do not, however, have any bearing on the arguments I shall pursue.

⁵ Note, however, that, although Field is pessimistic about the possibility of explaining the reliability of our mathematical beliefs on Platonists grounds, he never really commits to **B** and thus also not to the soundness of the argument I take the Benacerraf challenge to "potentially" give rise to. This is why I distinguish between the Benacerraf challenge and the argument it generates.

contents of our mathematical beliefs to be “literally true”. In fact, that’s the central component of his “fictionalist” view on the nature of mathematics. Field’s point is rather that, by ceasing to take the contents of the beliefs to be true, we are no longer Platonists. Note that Platonism consists of two elements, a “conceptual” one and a “metaphysical” one. That is, to be a Platonist is to think that our mathematical beliefs posit entities that exist beyond time and space, etc. But it is also to *have* some of the beliefs Field thinks should be given up and to believe that those entities actually exist. What the Benacerraf challenge potentially brings out is a tension between those elements. The idea is that the Platonist, by making assumptions that rule out the possibility of a viable explanation of the reliability of our mathematical beliefs, commits a kind of epistemic suicide. For what follows from **C** is, in effect, that, given Platonism, we should not be Platonists.

This implication does not establish the falsity of Platonism, of course. After all, the fact that a position implies about itself that it is unjustified is not inconsistent with its truth. However, it does seem to undermine Platonism in the sense that it supports not only the hypothetical claim that *given* Platonism, we should not be Platonists but also the categorical claim that we should not be Platonists full stop. Of course, that conclusion does also not follow, strictly speaking. But if we are forced to concede that a position we have adopted is permissible only if it is false then we have surely ended up in an awkward place. In any case, on the second, alternative construal of the argument, it continues in the following way:

D Given Platonism, we should not be Platonists.

E Therefore, we should not be Platonists.

We may accordingly distinguish between two versions of the argument that the Benacerraf challenge potentially generates (something I hereafter will call “the BF-argument”). On the

first version, the alleged inability to explain the reliability of our mathematical beliefs on Platonist grounds provides a reason to think that Platonism is false, since it means that it conflicts with an important intuition. On the second version, it rather threatens Platonism by rendering it epistemically self-defeating (in the way indicated by **D**). This may plausibly be taken to undermine the justification with which the position is held even if it doesn't undermine it in the more straightforward sense of showing it to be false. We may call the first version the "intuitive version" and the second "the self-defeat version".

3. Platonism and Non-Naturalism

Whether a plausible and distinct argument that is parallel to the BF-argument can plausibly be mounted against ethical non-naturalism depends on a number of issues. It depends on how closely related the positions are. But it also, as I shall illustrate, depends on which responses that are available to an ethical non-naturalist and to a Platonist, respectively.

Ethical non-naturalists believe in the existence of mind- and language-independent moral facts and moral properties. What makes their position non-naturalist is that they deny that those properties and facts are reducible to natural ones. Which facts are natural? Answering this question in an enlightening way has turned out to be difficult, but some are sympathetic to Michael Smith's suggestion that it depends on whether the fact is "the subject matter of a natural or social science" (Smith 1994, 17). Ethical non-naturalism is also often associated with the thesis that the moral properties it posits are causally inert.⁶

Now, Platonism makes similar assumptions about mathematical facts. But, as the above account indicates, what is often supposed to give rise to the Benacerraf challenge is rather the

⁶ The causal inertness it assumed by David Enoch's version of non-naturalism. However, other non-naturalists, such as Russ Shafer-Landau seem less committed to it. See Enoch (2011) and Shafer-Landau (2003), respectively.

fact that Platonists in addition make those assumptions about the *objects* mathematical terms are supposed to denote, such as numbers. Those entities are also taken to not bear any causal or spatio-temporal relations with us. Ethical non-naturalists, by contrast, don't posit any objects of that type. The bearers of moral properties are straightforwardly natural phenomena, such as actions, political institutions, or people, and no assumption is made to the effect that we don't causally interact with those phenomena. It may perhaps be argued that this makes a difference, and that it makes non-naturalists less vulnerable to the BF-challenge. However, that is not a strategy that I shall pursue.

What I do want to argue, however, is that there are responses to the versions of the BF-argument that applies to morality that do not seem available to a Platonist. To block those responses, the moral BF-argument must be supplemented by considerations that are external to the argument.

The response in question is to develop so-called "third-factor" explanations of the reliability of our moral beliefs. Consider Enoch's version of this strategy (Enoch (2010, for related suggestions, see Brosnan (2011), Schafer (2010) and Skarsaune (2011)). As we have seen, to explain the reliability of our moral beliefs is to explain an alleged correlation between our moral beliefs and the facts in virtue of which they are supposed to be true. Now, if the facts belonging to one set (*a*-facts) are correlated with those that belong to another set (*b*-facts) we may try to explain this either by arguing that *a*-facts cause *b*-facts or vice versa. However, the idea behind the third-factor strategy is to instead point to an external factor that is supposed to be responsible for both types of facts.

Let us in what follows focus on "rightness beliefs"; i.e., beliefs to the effect that certain types of actions are morally right. Enoch tries to provide the sought-for explanation by invoking an evolutionary account according to which we are disposed to have the rightness beliefs we have because the actions they prescribe increase the agent's fitness. According to

the version of the evolutionary account Enoch has in mind, moreover, they often do so through helping the agent or her kin to survive. It is that fact—the fact that the actions in questions promote survival—which constitutes the third factor to which Enoch appeals.

How does the explanation work? An important element is a further assumption that Enoch invokes, namely the substantively moral claim that “survival or reproductive success [...] is at least somewhat good” (Enoch (2010, 430)). The claim is clarified as follows:

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 (430).

Given this assumption, we may reason as follows. On the evolutionary account, the factor that explains our rightness beliefs is the fact that the actions they ascribe rightness to tend to promote survival. But, given that survival is good, this factor also tends to make the actions right, at least if we make the additional assumption that the fact that an act promotes the good counts in favor of its rightness. Thus, the same factor that is supposed to be responsible for the existence of our rightness beliefs—that the actions they prescribe promote survival—is also, at least partially, responsible for their truth. Enoch thinks that his account suggests that many of the beliefs we are endowed with by evolution are at least “somewhat in line with the normative truths” (2010, 430).⁷ He concedes that the level of reliability we can squeeze out is

⁷ A central concept in Enoch’s account is that of a fact being “responsible” for some other fact. Enoch takes this term to cover both a causal relation and “a metaphysical relation of constitution”. It is the latter type of relation that is supposed to hold between the fact that an action promotes survival and the fact that it is right. Note, however, that, as Enoch rejects

not overwhelming, as he thinks it merely suggests that the moral beliefs we have been endowed with through evolution are “reasonably good starting-points” that are not “too far off”. But he thinks that it still gives non-naturalists what they need.

The salient feature of Enoch’s proposal is that it tries to reconcile the reliability of our moral beliefs with the evolutionary account claim by invoking a substantive moral assumption. The assumption in question posits a relation between the facts in virtue of which our moral beliefs are true and the (non-moral) facts to which the evolutionary account attributes them.

There is an obvious way of criticizing Enoch’s strategy, namely by denying that he is entitled to invoke such a bridge principle. After all, the viability of the response arguably presupposes that he has some reason to *accept* the bridge principle in question. And since the principle in question is a substantive moral claim, the assumption that there is such a reason is in direct conflict with the conclusion advocates of the moral BF-argument think non-naturalists are committed to, namely that no moral belief is justified. In other words, by invoking the pertinent bridge principle Enoch seems to blatantly beg the question.

However, in my view, this objection is not persuasive. For by requiring that a non-naturalist must be able to justify the reliability of moral beliefs in an entirely non question-begging way is to impose a too strong condition that risks collapsing into global skepticism. For example, consider perceptual beliefs. We do have an account of the reliability of such beliefs; an account that consists of appeals to the theory of evolution and to assumptions about how our sense organs work and how they interact with our environments, and so on. This account involves the same circularity as Enoch’s does. For the assumptions in which it consists constitute straightforward empirical claims whose justification depends on the credentials of the very beliefs whose reliability it tries to establish. So if we deem Enoch’s

naturalism, the relation is not that of identity. For more about the type of relation Enoch has in mind, see (2010, 431).

reasoning illegitimate on account of its circularity, we seem committed to the same verdict about our perceptual beliefs. Accordingly, in so far as the advocates of the BF-argument don't want their position to collapse into some sort of global skepticism, the circularity that is involved by Enoch's strategy should be deemed legitimate.

Enoch's response to the BF-argument could obviously be questioned in various other ways. However, we may imagine further instances of his strategy and the points I presently want to make do not presuppose the soundness of that particular version.

The first point is simply that the same type of response does not seem available to a Platonist. It is not available since Platonists are not in any position to argue that the natural facts that are causally responsible for our mathematical beliefs are also (metaphysically) responsible for mathematical facts. This is important in the present context, because (and this is my second point) an advocate of the moral BF-argument will, unlike an advocate of the mathematical one, need some way of ruling out responses of that sort.

How could they be ruled out? It is at this juncture where the more familiar arguments to the effect that non-naturalism entails skepticism have a potential role to play.

The point of departure of one critical strategy is the aforementioned observation that the degree of reliability of the target judgments that one might hope to secure through Enoch's account is rather low. Darwinian factors might have ensured that our rightness beliefs track types of behavior that promote our *own* chances of survival or of the survival of *our* kin. But that doesn't necessarily mean that they are similarly sensitive to facts about how those types of behavior affect the survival of *others*. That detracts from their reliability. For if the claim that survival is good is supposed to concern survival in general, the effects on others are also relevant to the rightness of the target behavior. Moreover, survival is presumably not the only thing that is good and whose existence can influence the rightness of that behavior. So the

forces of natural selection have at best ensured that our rightness beliefs are sensitive to a very limited set of the considerations that determine rightness.

This is a potential problem. For to provide a response to the skeptical challenges a realist must not only reconcile her position with the existence of knowledge about trivialities (such as that pain is bad or that survival is “somewhat” good) but also about more controversial issues, including, for example, what weight one is permitted to assign to one’s own survival when facing a situation where the survival of many others are also at stake. It is hard to see how one could reach justified conclusion about such issues merely on the basis of the account Enoch sketches and the level of reliability it allegedly ensures.

To strengthen the case for the significance of the account, Enoch points to the additional, benign effects of ordinary reasoning processes. He writes:

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. (2010, 428)

It is the *combined* effect of these factors—the fact that our initial beliefs about what is right and wrong depend on the extent to which they promote survival and the benign effects of subjecting them to a reflective equilibrium-style reasoning process—that are thought to be sufficient for the non-naturalist to avoid skeptical implications.

However, it is not clear how much “extra” reliability the influence of this type of process really can generate. Presumably, all inquirers are affected by the genetical dispositions we have inherited through evolution, although some may be more capable of distancing

themselves from them than others. Still, the history of moral inquiry illustrates that different individuals, when seeking coherence in the way Enoch suggests, end up with radically different moral outlooks. Some turn out as consequentialists whereas others become virtue theorists or particularists, and this is so even if we lack compelling (and independent) grounds for thinking that some of them are less well equipped with the relevant reasoning skills than others. Disagreement of that type provides a potential problem for strong claims about the reliability of the process Enoch describes in the quote. Accordingly, it also potentially undermines Enoch's idea that he, by appealing to it, can strengthen the case for the viability of his response to skeptical challenges. We seem thus to be back at square one.

Notice that my aim here is not to try to refute Enoch's response through the above remarks. Maybe it can be elaborated in ways that avoid the problems. The point is rather to illustrate that in fending off Enoch-style responses an advocate of the moral version of the BF-argument must invoke considerations that are external to the argument, such as disagreement. Thus, it is not capable of standing alone and does not constitute an independent objection to non-naturalism that is distinct from other challenges.

4. Safety and sensitivity

We get further reasons for this conclusion by considering another recent response to the BF-argument, namely a response that, unlike the Enochian one, *is* available to the Platonist.

The response in question has been articulated by Justin Clarke-Doane (Clarke-Doane 2015). Clarke-Doane's idea is that, to put the skeptical worries that might be raised by our inability to explain the reliability of a set of target beliefs to rest, it is sufficient to provide a compelling argument to the effect that, given the truth of the beliefs, they could not easily have been false (in which case they are said by Clarke-Doane to be "safe"). And his point is

that a Platonist *can* consistently construct such an argument, in spite of her peculiar assumptions about the nature of mathematical objects.

The argument proceeds as follows. One way in which our beliefs in an area could easily have been false is through not being “sensitive”, where they fail to be sensitive if they would have remained the same (had the same contents) even if the facts they represent had been relevantly different. “Sensitivity” denotes a type of modal dependency we expect from a bathroom scale. Suppose that the scale shows that a person weighs 65 kg but that we learn that it had shown this even if the person had been significantly lighter. This saps our reasons for trusting it, regardless of whether the person actually does weigh 65 kg. Now, according to Clarke-Doane, the sensitivity of our mathematical beliefs is secured by the metaphysical necessity of mathematical facts. For their necessity ensures (on one view on the truth conditions of counterfactuals) that it is vacuously true that if the mathematical facts had been different, then our beliefs had been correspondingly different. After all, if there is no possible world in which they are different, there is also no world in which they are different while our beliefs remain the same.

Given the sensitivity and actual truth of our mathematical beliefs, there is only one remaining way in which they could easily have been false. This is to argue that the beliefs could easily have been *different*, even if the facts had remained the same. For if we could easily have had different mathematical beliefs then we could easily have had false mathematical beliefs even supposing that they are in fact (necessarily) true. However, this alternative way of questioning the safety of our mathematical beliefs can also be ruled out, according to Clarke-Doane, at least in the cases of our “core” or most basic mathematical beliefs. It can be ruled out, more specifically, with reference to an evolutionary account of their origins. The idea seems to be that given such an account we would have had different beliefs only if our environment had been quite different, which means that it excludes the

existence of nearby possible worlds in which our (core) mathematical beliefs are different.

Clarke-Doane concludes that the necessity of mathematical facts and the evolutionary inevitability of our core mathematical beliefs explain their reliability in the sense of showing that, given their truth, they could not easily have been false. He also thinks that this is the only sense in which the inability to produce such an explanation undermines their justification. Moreover, since the explanation is consistent with the assumptions a Platonist makes about the nature of mathematical objects, it is clearly available to her. The same reasoning applies, according to Clarke-Doane, to ethical non-naturalism. For moral beliefs can also be given an evolutionary explanation and (some) moral facts also seem to hold with necessity. Neither the moral nor the mathematical version of the BF-argument therefore gets off the ground.

Let us grant Clarke-Doane the conclusion that Platonists and non-naturalists can consistently and plausibly argue that the target beliefs are both safe and sensitive on the basis of the considerations he adduces. In order for the BF-argument to survive this concession there has to be an alternative interpretation of “explain reliability” to the one Clarke-Doane relies on. This alternative interpretation must in turn meet two conditions. It must, first, be such that the argument Clarke-Doane has proposed does not qualify as an explanation in that sense, and, second, such that it is plausible to hold that the principled inability to explain the reliability of the target beliefs in that sense undermines their justification.

5. Explaining reliability

Is there an interpretation that meets those conditions? We may begin to explore this question by reconsidering why the inability to explain the reliability of a set of beliefs is supposed to undermine them. Notice in this context that there are two ways to undermine a belief’s justification. The most straightforward way is to offer some compelling reason to think that it is false. Another is to show that we don’t have sufficient reason to think that it is *true*. If a

consideration of the latter type does not also provide evidence for the falsity of the belief, it supports agnosticism with regards to its content rather than outright denial.

The inability to viably explain the reliability of a set of beliefs, such as the set that includes our beliefs about the village in Nepal, seems at best to undermine them in the second way. After all, the fact that we cannot construct such an explanation shows hardly that an odd number of the inhabitants are *not* born on a Friday. Let us therefore assume that in so far as the inability to explain the reliability of our beliefs in some area undermines them, it does so through ruling out that we have sufficient *positive* reasons for believing in their truth.

In order for a reconstruction of the BF-argument along those lines to work, however, the availability of an explanation of the relevant kind must presumably *be* (or provide) such a reason. Clearly, this does not hold for the explanation Clarke-Doane offers; i.e., the one that appeals to the necessity of mathematical and moral facts and the evolutionary inevitability of our beliefs about them. What a Platonist or non-naturalist might accomplish by citing those considerations is to show that, *if* the beliefs are mostly true, then they could not easily have been false. But that conclusion does not in turn give us any reason to conclude that they *are* in fact true. Consider an analogy. Suppose that two friends visited a certain city in the weekend and that we wonder if they met. Suppose also that we know that each had a very determinate plan about her visit. Then we may perhaps conclude that if they met this was not a coincidence and that it could not easily have been different. But we would still be in the dark as to whether they actually *did* meet.⁸ Clarke-Doane's account of the reliability of our mathematical beliefs does accordingly not serve the dialectical role the BF-argument assigns to such accounts, given the reconstruction of it that we are now pondering.

⁸ If we were better informed about the details of their plans, then we would perhaps be able to figure out if they met. But the point is that this further information is not required to reach the conclusion that *if* they met it could not easily have been otherwise.

Which types of explanations, then, *do* serve that role? I wrote earlier that to explain the reliability of our beliefs is to explain that we have arrived at beliefs in the relevant area which are mostly true. Presumably, this involves explaining why we *have* arrived at those beliefs in the first place. Some such explanations are neutral relative to the truth of the beliefs, in the sense that the truth of the assumptions they involve doesn't affect or increase their likelihood of being true. For example, the fact that we have formed them as a result of reading a book is thus neutral as long as we don't know anything about its trustworthiness. Other explanations *do* increase their likelihood, however, by indicating or providing grounds for thinking that they are true. Indeed, the best explanations of some of our beliefs even *presuppose* their truth, such as when we explain why we have the perceptual belief that there is a table in front of us by citing the fact that there *is* a table in front of us. On one suggestion, it is explanations of the second kind—explanations of why we have the target beliefs that *also* provide grounds for thinking that they are true—that we should refer to as “explanations of reliability”. This understanding of what constitutes an explanation of the reliability of a set of beliefs fits, unlike Clarke-Doane's, with the interpretation of the BF-argument we are presently considering. For if the best explanation of a set of beliefs provides a ground for thinking that they are true, we may argue that the inability to provide such an explanation undermines their justification by depriving them from that ground (premise **A**).

A problem with this suggestion is that the support it provides for **A** seems so weak. After all, the fact that *one* potential source of justification for a belief is ruled out may not exclude that there are other sources. Consider again our beliefs about the village in Nepal. The absence of a plausible explanation of their reliability (in the present sense) is consistent with the fact that they are related in such a way that each of them gets some conditional support from the others, so that, for example, our belief that an odd number of inhabitants are born on a Friday can be backed up with our belief that Sonam, Kiran and Yash were born on a Friday

and no one else who lives in the village. The same holds for our mathematical beliefs. Why isn't that type of internal coherence enough for justification or epistemic permissibility?

The answer must presumably be that there is something special about the latter type of support that makes the absence of it especially significant. One option is to stress that it is external to, and/or independent of, the (the truth of the) target beliefs. If the best explanation of our having the belief that P implies or strongly indicates the truth of P, then the fact that we have it is evidence for P, regardless of the content of P. The fact that we have the belief is a psychological and natural fact, and can be established without settling any moral or mathematical issues. This is why—on one idea—the availability of an explanation of the reliability of our moral or mathematical beliefs provides crucial support for them and why the principled absence of such an explanation may have skeptical implications.

The reasoning just hinted at needs to be extensively elaborated to have any prospects of success, for example regarding the concept of “independence” that might plausibly be invoked. Let me again stress, however, that my aim in this paper is not to defend any particular construal of the BF-argument. The point I want to make is rather that it makes the request to show how the reliability of moral or mathematical beliefs could possibly be explained on non-naturalist or Platonist grounds collapse into the more familiar demand to show how the permissibility of our beliefs in the target areas can be reconciled with a naturalistic epistemology.

Note that, if it is the naturalness of the fact that we have a belief that is the source of the significance the BF-argument assigns to explanations of reliability, then support from *any* kind of natural facts would do just as well. For example, a Platonist could try to respond to the argument by stressing, à la Quine, that (at least some subset of) our mathematical beliefs are justified in virtue of playing an indispensable role in physics. And a moral realist may hope to do the same with regards to morality by appealing to the role moral assumptions possibly play

in the explanations of phenomena such as, to use Nicholas Sturgeon's example (Sturgeon 1988), social revolutions. Of course, both these defensive projects might fail. But the thing to note is that if an argument to the effect that the best explanation of some natural phenomenon assumes moral facts counts as a response to the moral version of the BF-argument, then it coincides with Harman's argument to the effect that we have no reason to believe in the existence of objective moral facts because they are never posited by the best explanation of anything observable.

Hartry Field is of course associated with the view that the Quinian argument *does* fail, through his attempts to "nominalize" physics and to show that the role mathematics play isn't indispensable after all.⁹ But he has also suggested that the outcome of that discussion is really irrelevant to the Platonist's ability to handle the BF-challenge and that a successful indispensability argument would not really help the Platonist against it. One may accordingly conclude that the construal of the argument developed in this section fails to capture Field's intentions. However, Clarke-Doane's argumentation illustrates that the room for an alternative construal is slight. This provides further support for the thesis that, in so far as the BF-argument has any weight at all, it does not constitute a distinctive, separate objection.¹⁰

6. *Concluding remarks*

We started out by wondering if the Benacerraf challenge allows the critics of non-naturalism to expand their dialectical arsenal, by offering a way of reaching their conclusions without having to get their hands dirty through empirical speculations. The upshot of my discussion is that it doesn't really. When we seek to reconstruct the challenge in a way that avoids crucial

⁹ See Field (1980). For a defense of mathematical indispensability arguments, see Colyvan (2001).

¹⁰ For further discussion of Clarke-Doane's arguments and their implications to skeptical challenges against moral realism, see Tersman (2016).

objections it is transformed into a version of Harman's well known debunking-style argument and must be supplemented for example by considerations invoked by the argument from disagreement.

I would like to end, however, with some partially self-critical remarks. The aim to identify distinct arguments for a philosophical position is understandable because it may enable us to proceed in a piecemeal fashion when we want to assess it. We could start by looking at one argument, and when we're done we turn to the next, and so on, until an overall evaluation is reached. In reality, however, many arguments in fact interact in different ways, and the piecemeal strategy sometimes leads us to ignore the existing dependencies, which in turn makes us less well placed to assess the argument. This observation calls, I think, for a more holistic approach. As John Rawls has stressed, the justification of a philosophical position is "a matter of the mutual support of many considerations, of everything fitting together into one coherent view." (Rawls 1971, 21) It may seem that I have violated this dictum by focusing on the question of whether the moral BF-argument is a distinct objection. However, the conclusion I have defended can be formulated in a way that acknowledges it. The conclusion is simply that the BF-argument doesn't contribute much to the relevant whole. It relies on basically the same epistemic ideas and machinery as the more familiar challenges but adds very little, if anything, to the bigger picture.

References

- Benacerraf, P. (1973). "Mathematical Truth", *Journal of Philosophy* 70, 661-79.
- Brosnan, K. (2011). "Do the evolutionary origins of our moral beliefs undermine moral knowledge?", *Biology and Philosophy* 26, 51-64.
- Clarke-Doane, J. (2015). "What is the Benacerraf Problem?", in F. Pataut (ed.), *New Perspectives on the Philosophy of Paul Benacerraf: Truth, Objects, Infinity*, forthcoming.

- Colyvan, M. (2001). *The Indispensability of Mathematics*, Oxford: Oxford University Press.
- Enoch, D. (2010). “The epistemological challenge to metanormative realism: how best to understand it, and how to cope with it”, *Philosophical Studies* 148, 413–38.
- Enoch, D. (2011). *Taking Morality Seriously: A Defense of Robust Realism*, Oxford: Oxford University Press.
- Field, H. (1980). *Science Without Numbers*, Princeton: Princeton University Press.
- Field, H. (1989). *Realism, Mathematics, and Modality*, Oxford: Blackwell.
- Field, H. (1996). “The A Prioricity of Logic”, *Proceedings of the Aristotelian Society* 96, 359-79.
- Field, H. (2005). “Recent Debates About the A Priori”, in Gendler and Hawthorne (eds.), *Oxford Studies in Epistemology*. Oxford: Oxford University Press, 69–88.
- Field, H. (2009), “Epistemology without metaphysics”, *Philosophical Studies* 143, 249-90.
- Harman, G. (1977). *The Nature of Morality*, New York: Oxford University Press.
- Joyce, R. (2013). ”The Evolutionary Debunking of Morality’, in J. Feinberg & R. Shafer-Landau (eds.), *Reason and Responsibility*, 15th edition, Cengage.
- Rawls, J. (1971). *A Theory of Justice*, Cambridge, Mass.: Harvard University Press.
- Schafer, K. (2010). “Evolution and Normative Scepticism”, *Australasian Journal of Philosophy* 88, 471-488.
- Shafer-Landau, R. (2003). *Moral realism: A defence*, Oxford: Oxford University Press.
- Skarsaune, K. O. (2011). “Darwin and Moral Realism: Survival of the Iffiest”, *Philosophical Studies* 152, 2011, 229-243.
- Smith, M. (1994). *The Moral Problem*, Oxford: Blackwell.
- Street, S. (2006). “A Darwinian Dilemma for Realist Theories of Value”, *Philosophical Studies* 127, 109-66.

Sturgeon, N. (1988). "Moral Explanations", in G. Sayre-McCord (Ed.), *Essays on moral realism*, Ithaca: Cornell University Press, 229-255.

Tersman, F. (2006). *Moral Disagreement*, New York: Cambridge University Press.

Tersman, F. (2016). "Explaining the reliability of moral beliefs". In *Ethics and Explanation* (eds. Leibowitz, U., and Sinclair, N.), Oxford: Oxford University Press, 37-57.