

Evolution and Ethics: No Streetian Debunking of Moral Realism

FRANK HOFMANN
University of Luxembourg, Luxembourg

This paper is concerned with the reconstruction of a core argument that can be extracted from Street's 'Darwinian Dilemma' and that is intended to 'debunk' moral realism by appeal to evolution. The argument, which is best taken to have the form of an undermining defeater argument, fails, I argue. A simple, first formulation is rejected as a non sequitur, due to not distinguishing between the evolutionary process that influences moral attitudes and the cognitive system generating moral attitudes. Reformulations that respect the distinction and that could make the argument valid, however, bring in an implausible premise about an implication from evolutionary influence to unreliability. Crucially, perception provides a counterexample, and the fitness contribution of reliably accurate representation has to be taken into account. Then the moral realist can explain why and how evolution indirectly cares for the truth of moral attitudes. The one and only condition that has to be satisfied in order for this explanation to work is the sufficient epistemic accessibility of moral facts. As long as the moral facts are sufficiently reliably representable, one can see how evolution could favor getting it right about the moral facts. Interestingly, apart from this epistemic constraint no further constraint and, in particular, no objectivity constraint on what the moral facts have to be like can be derived. Thus, the only problem for the moral realist is to make good on epistemic access to moral facts—an old problem, not a new one.

Keywords: Evolution, ethics, debunking, moral realism, reactive attitudes.

1. *Introduction*

Discrediting a view or set of belief-like attitudes that aspire to truth by so-called 'debunking' has become quite popular.¹ An especially interest-

¹ See, for example, Kahane (2011) for a list of examples and some general discussion of debunking arguments.

ing case is the case of moral attitudes and evolution. The debunking argument attacks our moral attitudes (or an important part of them)² by pointing out that the force that has shaped these moral attitudes—evolution—is blind to issues of truth and only cares about fitness and survival.³ Take ‘altruistic attitudes’ as the core of our moral attitudes. An evolutionary explanation of these attitudes can be given by reference to the fitness-conduciveness of certain kinds of cooperative behavior towards one’s kin or group. If these altruistic attitudes are understood realistically, we run into a problem: they might be true, but only accidentally so, since the relevant evolutionary forces did not care about truth. Or so the debunking line of thought runs.

Not too long ago, Sharon Street has presented an evolutionary debunking argument against moral realism, called the ‘Darwinian Dilemma’.⁴ So she is turning the argument around, targeting the realistic assumption that the relevant attitudes aspire to truth and are, at least sometimes, true—what is often called ‘*moral realism*’. (Whether realism is or should be construed as implying some objectivity condition can initially be left open and will be investigated later.) The argument is clearly of high significance, since it attempts to bring to fall an entire approach in (meta-)ethics, i.e., moral realism. Street’s formulation of the argument, however, is burdened with considerations that rather distract from those valuable and interesting points that may constitute a sound argument against moral realism. One could try to disentangle the various claims and streams of thought in her presentation(s) of the argument, which would require an immense amount of careful interpretational work. Here I would like to proceed in a different way, namely, by providing a clear and systematic reconstruction of a core argument that can be extracted from Street’s considerations, without any side roads or unnecessary accompaniments. If it turned out that the argument is not really Street’s argument, this should not be too worrisome, since it is at least similar and in any case important enough.⁵

² I will take the relevant moral attitudes to be beliefs or sufficiently belief-like, since (only) they are aspiring to truth. Any difference should not matter to the argument. Street includes desires (and attitudes of approval and disapproval) among the evaluative attitudes (cf. Street 2008: fn. 3). However, these are not really at stake since they are not truth evaluable and, therefore, the issue of reliability does not arise for them, at least on a standard conception of desires.

³ A particularly succinct statement of this claim—or dogma, indeed—can be found in Burge (2010), in particular, ch. 8. Burge discusses this idea in the context of a naturalistic teleosemantics which conceives of representational functions as biological functions, and tries to argue that it has to fail for exactly that reason. For a convincing criticism of Burge’s argument see Graham (2014). For another statement of the claim see, for example, Stich (1990: 62).

⁴ The original statement is to be found in Street (2006). Street explains the argument further and, in particular, defends it against the criticism by Copp (2008) in Street (2008). Street takes value realism in general as her target. In order to keep the discussion in reasonable bounds here, I will restrict myself to the moral part of value realism, i.e., to moral realism.

⁵ There are several problematic aspects to the exposition of the ‘Darwinian Dilemma’ as to be found in Street’s writings. I only mention a few of them here. It is

In the following I will try to carve out and (re-)formulate what I take to be the heart of the considerations that can be found in Street's writings and that might provide a genuine evolutionary argument against moral realism.⁶ Exegetical issues will not be my primary concern, but systematic reconstruction.⁷ I will try to show that in the end, the argument fails. The moral realist can tell a plausible story about how evolution cares for the reliability of moral attitudes—for in a nutshell, evolution *indirectly* cares (or, at least, can care) for truth, and there is a plausible story about how this could work. The new part of the story that will be told here appeals to *reactive attitudes* (in Peter Strawson's sense) *as a social mechanism* that explains how the relevant cognitive process could be reliable after all. It provides a new picture of how evolutionary forces could have *indirectly* resulted in reliability. This adds significant support to the idea of indirect truth tracking as a response to the debunking challenge. Interestingly, this makes appeal to other strategies for defending moral realism superfluous (like appeal to the additional epistemological potential of rational reflection, (quasi-)conceptual truth, or a threat of self-defeat within the debunking argument).⁸ Fortunately, as will be argued at the end, there will not

questionable whether putting the argument in the form of a dilemma is fortunate. Construing it as a *reductio* of the moral realist assumption, by way of a defeater argument, seems more appropriate, as I will try to show. – Graber has criticized Street for setting things up in the form of a dilemma, too (cf. Graber 2012: 594). Graber suggests that we should take the argument to be an abductive one, i.e., an argument about best explanation. I disagree. Furthermore, Street's formulation also depends on a demand for reasons (for believing in the reliability of the source of our moral attitudes), or even on a demand for *independent* reasons—thus leading into issues of epistemic circularity that are at most implicitly hinted at by Street but by no means discussed explicitly and to a sufficient extent (cf. Street 2008: sc. 6). Externalists will reject such a demand for reasons. A further point of unclarity is the definition of 'moral realism' as it is the target of the argument. Street burdens 'moral realism' (and equally realism about value) with an element of objectivity which is not really required for any kind of moral realism, but only for an *objective* moral realism (cf. Street 2006: 110; Street 2008: 208 and, in particular, fn. 3). If reliability is the primary issue, why is not any kind of moral realism affected?, one can wonder. (Skarsaune presents some good critical observations about Street's characterization of 'moral realism' in Skarsaune 2011: sc. 5.)

⁶ Thus, to repeat, I will restrict the discussion to moral realism (and ignore non-moral value realism). Therein I follow Copp (2008). Street notes that she sees some difficulties with this restriction since it "introduces crucial complexities having to do with morality/reasons internalism" (Street 2008: 209). I have to confess that I fail to see any significant problem with the restriction.

⁷ Some useful exegetical work, including important interpretational questions, has been provided by Copp (2008), Enoch (2010), Skarsaune (2011), and Garber (2012).

⁸ Brosnan (2011) and FitzPatrick (2015) try to argue that *rational reflection* can lead to reliable belief formation (even if starting with initially false input beliefs). Conceptual truths, or something close enough, have been offered in response to debunking arguments by Cuneo and Shafer-Landau (2014). The charge of self-defeat is discussed in Kyriacou (2016).

be any new significant epistemological costs to the proposed response in defense of moral realism.⁹

2. *Reconstructing a Streetian evolutionary argument (crude version)*

Here is my proposal about how to formulate the initial, crude argument, argument A1:

A1:

- (1) Our moral attitudes (MAs) are significantly influenced by an evolutionary process E.
- (2) E is not truth-tracking.

-
- (3) Our moral attitudes (MAs) are not (epistemically) justified.¹⁰

Some comments on the premises are in order. Premise (1) is a rough statement that leaves out which moral attitudes exactly are at stake. I have already indicated that we are talking about the ‘altruistic core’ of our moral attitudes, concerning positive evaluation of cooperation and family support etc. (even at some cost to one’s own self-interest). For brevity’s sake I will call them the ‘MAs’. (As already mentioned, I will treat the MAs as beliefs or sufficiently belief-like to be epistemically evaluable in the relevant way. Moral attitudes that are not truth-evaluable cannot be the target of the argument.) Of course, there is no precise definition of what counts as ‘significant influence’. But it seems to be agreed on by all parties in the discussion that the influence is significant, at least for the sake of the argument. After all, it could turn out to be the case (if it has not yet). So let us take premise (1) for granted for the moment and see where it leads. (We will see soon that some reformulation is necessary.)

⁹ Similar arguments against various forms of realism (evaluative realism, religious realism, etc.) can be found in the literature. See, for example, Joyce (2008) and Ruse, Wilson (1995). Interesting comparisons of Street’s arguments and Plantinga’s evolutionary argument against atheism can be found in Crow (2015) and Moon (2016). A very good overview on evolutionary debunking arguments is Vavova (2015). Vavova also presents her own reconstruction of the most promising evolutionary argument which is different from the one I am reconstructing here. She does not recognize the distinction between evolutionary and cognitive processes that I will argue for in the next section, and she does not discuss the criticism of, and reply to, the argument that I will present in section 4. The same is true of Vavova’s earlier discussion in Vavova (2014).

¹⁰ This is essentially an instance of the argument schema that Kahane proposes as a general schema for debunking arguments, see Kahane (2011: 111). Kahane discusses Street’s argument as fitting into this schema (cf. Kahane 2011: sc. 3). However, Kahane does not present the criticism that I am going to lay out here. Rather, he gets into issues of objectivity and whether the argument over-generalizes—which are not the problems I am discussing here. Skarsaune (2011) also gets into the issue of what ‘independence of our attitudes’ means and whether philosophers like Nagel or Parfit hold that moral truths are ‘independent’ in this or that sense.

Premise (2) is the statement of evolution's blindness to truth. Evolution does not care for truth but only about survival or practical value. This slogan is taken to amount to a lack of truth-tracking. The evolutionary process E is epistemically evaluated as a bad one: it does not push toward truth, it is not truth-conducive.

Why is (3) supposed to follow?—I propose that we should take the argument to have the form of a *defeater argument* or, more precisely, an *undermining-defeater argument*.¹¹ Its premises provide an undermining defeater against the epistemic justification of the relevant moral attitudes.¹²

Famously, Pollock distinguished between two kinds of defeaters, rebutting and undermining (or undercutting) defeaters.¹³ The rebutting defeaters (against the belief that p) are simply reasons for the opposite belief, i.e., belief that non-p. The undermining defeaters consist in reasons against the reliability of the source of one's belief that p. They undermine the source as not being reliable and, thus, as not issuing (*ultima facie*) justified belief (as long as they are not themselves undermined, of course).¹⁴ I propose to take the evolutionary debunking argument as a defeater argument of the undermining sort. The source is taken to be the evolutionary process E and its lack of truth-tracking is taken to undermine E's reliability.

If taken in this way, we can see how one could think that the conclusion (3) follows from the premises. Premise (1) establishes a source of our MAs, and premise (2) discredits it, so the MAs lose the epistemic status of being justified. It is just like when a belief loses epistemic justification if one acknowledges that it has been generated by unreliable wishful thinking, for example.¹⁵

Prima facie, A1 looks like a sound defeater argument. The history of our moral attitudes seems to be discredited such that their (epistemic) justification is undermined. For what comes from a source that does not track truth could at best be accidentally true—and that is not good enough for (epistemic) justification.¹⁶

¹¹ In general, it seems that 'debunking arguments' can be best understood as undermining-defeater arguments. (Perhaps there are some exceptions, but this should be the rule.) Kahane makes the same proposal (cf. Kahane 2011: 105–6). Schafer also takes the heart of Street's considerations to be aiming at an (*a posteriori*) defeater (of an *a priori* entitlement). And he criticizes the attempt for not relying on the idea that the *empirical* facts of evolutionary theory could provide a defeater for a *normative* claim (cf. Schafer 2010).

¹² It is quite clear that our primary concern here is with *doxastic* justification, not with propositional justification or personal justification.

¹³ Cf. Pollock (1986) and, for a general overview on evidence, Kelly (2014).

¹⁴ To be more explicit, a belief might be *prima facie* justified, but any undermining defeater cancels its *ultima facie* justification (as long as it is not undermined itself). Cf., for example, Senor (1996).

¹⁵ Of course, the undermining defeater is supposed to be not defeated itself.

¹⁶ In general epistemology there is quite a controversy about how to understand how (undermining) defeaters exactly work, and some epistemological views seem to have problems here. (For some recent discussions of defeaters see, for example,

However, the appearance is misleading. There are at least two big problems with that argument. First, and most importantly, as it stands, argument A1 is a *non sequitur*. And second, and related to that, it is unclear how premise (2) is really to be understood (in the best, charitable way). Let me explain.

Firstly, the argument is a *non sequitur*. This is so since *E* is not a cognitive process. Therefore, the non-truth-tracking of *E*, stated in premise (2), is simply *not directly relevant* to the epistemic status of justification of our MAs. (It might be indirectly relevant—we will come back to that in due course.) What matters is whether the relevant cognitive process, i.e., the cognitive process of which our MAs are the outcomes, is reliable or not. But since *E* is not a cognitive process, *E*'s blindness to truth is not directly relevant. (Its indirect relevance would have to be made explicit, and I will try to do so soon.) As the argument A1 stands, it is a *non sequitur*. Premises (1) and (2) do not provide an undermining defeater against the assumption that the relevant cognitive process of which our MAs are the outcome is reliable. An undermining defeater argument has to provide reason against the reliability of *the relevant cognitive process*, since only the relevant cognitive process's reliability is what matters for the epistemic justification of its outcomes. No such reason has been provided so far.¹⁷ This is the decisive shortcoming of argument A1.¹⁸

Bergmann (2006) and Hofmann (2013).) But the phenomenon is widely acknowledged, and my discussion does not depend on any controversial account of, or assumptions about, undermining defeaters. In particular, it does not depend on whether one goes for a psychologistic or an anti-psychologistic conception of defeaters (i.e., whether one conceives of defeaters as beliefs or the propositions believed).

¹⁷ That cognitive processes are the relevant items when it comes to epistemic justification and defeaters is a common assumption in general epistemology, at least for those epistemologists which admit the relevance of (certain aspects of) the history of a belief (cf., for example, Goldman 1979). If one opts for an a-historical, 'current time-slice' epistemology, such as, for example, Pryor's dogmatism, then no historical processes are directly relevant to the status of epistemic justification. Then, however, the prospects for an evolutionary undermining defeater argument are even dimmer. So the assumption that the history of an attitude matters is granted for the sake of the argument, and not a substantial assumption that I am making. But if the history is relevant, we have to be clear about which parts or aspects of the history of a belief are relevant. There is of course some significant controversy about the individuation of cognitive processes, whether they are more narrowly or more broadly construed and whether they are to be construed individually or socially. The generality problem is a major topic here. (Cf., for example, Goldman 2008. For a social individuation see Goldberg 2010.) But some limits are highly plausible and commonly accepted. In particular, it is quite clear that a cognitive process does not extend temporally back beyond the individual's existence. The influence of evolution lies of course way back in the past, much beyond the individual's life. Thus the evolutionary process *E* cannot count as a cognitive process.

¹⁸ The very idea of this criticism can be found in Wielenberg (2014). Wielenberg, however, does not try to re-formulate the argument and, thus, does not arrive at the—very significant—result that will be forthcoming from the following reformulations, namely, that the indirect 'truth tracking' account can be upheld and spelled out by means of a social reactive attitudes story (see section 4, below).

Secondly, premise (2) is unclear. What does it mean to say that ‘E is not truth-tracking’? Once we have distinguished between the cognitive process that is responsible for our MAs and the evolutionary process that has significantly influenced our cognitive systems, we have to spell out the ‘non-truth-tracking’ of the evolutionary process. It is not a cognitive process, and so talk of reliability is not appropriate or must be understood in some other way (different from how it is understood when applied to cognitive systems). How then is the lack of ‘truth-tracking’ of E to be understood?¹⁹

Can we save the argument by reformulation?, you may ask. If we want to formulate a *sequitur* argument, we have to reformulate everything in terms of the cognitive process—call it ‘C’—which produces the MAs in us. In effect, the recipe for reformulation is quite clear. First, we have to split premise (1) into two, so to speak, one pertaining to the cognitive process C and another one pertaining to the evolutionary process E. Second, we have to find a suitable explication or replacement for premise (2).

3. *Refining the argument*

Following the recipe leads quite naturally to the following reformulation of the argument—call it argument A2:

A2:

- (1a) Our MAs are produced by a cognitive process C.
- (1b) C has been significantly influenced by E.
- (2) E influences cognitive processes significantly because of their fitness contribution.²⁰

(3) Our MAs are not justified.

The problem with this reformulation, however, is quite clear: it is still a *non sequitur*. The premises are silent about the reliability of the relevant cognitive process C. By merely stating that the evolutionary process E goes by fitness contribution, we have not yet been told whether the result of E’s influence is reliable or not. In other words, we have not yet spelled out the first part of the slogan, ‘evolution does not care for truth ...’, but only the second part, ‘... but (only) about survival’. Even if we leave in the ‘only’ it is not clear what follows for the issue of C’s reliability. We could change premise (2) accordingly:

¹⁹ A note on the exegetical debate may be in order here. Copp (2008), Enoch (2010), and Skarsaune (2011) struggle with understanding Street’s ‘tracking account’. But it seems that the issue is undecidable since unclear. Only cognitive processes, or systems, that produce truth-evaluable states are in the business of ‘tracking truth’, i.e., are supposed to be reliable. The evolutionary process is no such process and, thus, falls outside of any reliable/unreliable classification.

²⁰ I choose the ‘because of’ formulation in order to avoid any controversial commitment to teleological notions, such as selection-for. One could equally well speak of ‘according to’. Nothing hangs on this.

- (2') E influences cognitive processes significantly because of, and only because of, their fitness contribution.

Still it would be unclear what consequences for reliability this would have. Any such consequences should be spelled out explicitly, since they are not obvious at all and are the ones that do the crucial work in the argument. All we are entitled to assume from contemporary evolutionary theory is (1b) and (2) or (2'), but evolutionary theory neither contains nor implies any claim about the unreliability of C.

So how are we to fill in the required additional premise in the best possible way?²¹ I submit that the best supplemented argument, argument A3, reads like this:

A3:

- (1a) Our MAs are produced by a cognitive process C.
- (1b) C has been significantly influenced by E.
- (2a) E influences cognitive processes significantly because of their fitness contribution.²²
- (2b) If C has been significantly influenced by E because of a fitness contribution, then C is not reliable.

(3) Our MAs are not justified.

Now the problem lies with *premise (2b)*. This further premise (2b) is explicitly about reliability or truth-tracking, and so the gap is filled. But (2b) is *not plausible* (as will be shown in a minute). We have turned the argument into a *sequitur*, but only at the price of introducing an implausible assumption. Therefore, the undermining-defeater argument fails. And so the overall conclusion is that no reason against moral realism has been presented so far. Moreover, this result remains even if we switch from (2a) and (2b) to (2a') and (2b'):

- (2a') E influences cognitive processes significantly because of, and only because of, their fitness contribution.
- (2b') If C has been significantly influenced by E because of, and only because of, its fitness contribution, then C is not reliable.

It remains to be shown that premise (2b) is not plausible. There are at least two reasons. The first reason is that there is a *clear counterexample*, namely, *perception*. Perceptual systems count as cognitive systems in the relevant sense, since they provide perceptual representations which can be correct or incorrect (accurate or inaccurate) and, thus, are evaluable as reliable and unreliable.²³ Quite plausibly, perceptual systems (like the visual system) have been under the significant influence

²¹ There might be some unclarity about what it means to say that E exerts influence 'because of fitness contribution', or that E 'goes by fitness'. But whatever unclarity there is, it is a further problem, not the problem I am belaboring here.

²² (2a) = (2). I have chosen the renaming simply because of the nicer partitioning that results.

²³ I will switch back and forth between processes and systems. Nothing should hang on that.

of evolution, and they have been so because of their fitness contribution. But it is not plausible to think that perceptual systems are *therefore* unreliable. (2b) does not state a true connection between evolution's influence and the reliability of its outcome, the evolved cognitive processes.²⁴ The second reason for rejecting (2b) is a *theoretical consideration* (that somehow generalizes the first point). In general, a fitness contribution may *consist exactly* in reliably correct representation generated by a cognitive system (of the relevant part of reality, under favorable conditions, and in cooperation with a suitable action-control system, of course).²⁵ Such reliability might be useful, or even very useful. This applies to our MAs just as well as to any other cognitive systems and processes—as long as we are dealing with truth-evaluable attitudes or states.²⁶ For example, to have moral attitudes that favor cooperation with cooperative partners is of course very useful for getting along well with other individuals, on the whole and at large. (We will take up and describe this kind of social advantage further in the next section.) Being reliably correct about the moral facts can thus be very beneficial. If there is any truth to this theoretical consideration, then premise (2b) cannot be true. So argument A3 is not sound.²⁷

24 The case of perception as a counterexample has been noticed by others, for example, by Brosnan (2011).

25 In this connection Peter Graham has put forward very interesting points about the more precise way in which the truth (veridicality, accuracy) of representations can be understood as contributing to the fitness of their possessors, even without those representations having (directly) the (teleo-)function of increasing fitness. See Graham (2014). One of the most important points is that such contributions take the form of a whole package with a functional analysis such that each element does what the other elements need, if everything goes the normal way, in order to produce the fitness increasing behavior. Producing true (veridical, accurate) representations is what the cognitive process C does (in normal conditions), and other elements have the job of producing behavior which is appropriate to the corresponding truths, based on these true representations. This seems to me to be exactly the way in which it could be said that 'evolution *indirectly* cares for truth'. In this way, any suggestion of an incompatibility of adaptation on the one hand and reliably correct representation on the other hand can be rejected. Our cognitive process C can be seen to be both an adaptation (with the function of reliably producing true representations of moral facts) and reliable or 'truth tracking' (in favorable or normal conditions for which it is made). This is exactly *how* C can contribute to fitness.

26 Teleosemantics even builds an account of representation on this idea of usefulness. There are many versions of teleosemantics that differ in details. But a common core is that the usefulness of (the use of a significant amount of earlier) correct information is partly constitutive of representation (now). Cf., for example, Millikan (1984) and Dretske (1995).

27 Huemer has tried to present a kind of argument, albeit in a sketchy form (as Huemer himself admits), for the reliability of our moral attitudes or "for why people should have correct ethical beliefs", as he puts it (Huemer 2005: 2). He argues that evolution favors having correct ethical beliefs given that these have a certain, 'altruistic' content (Huemer 2005: 218–9). As I understand his sketchy argument, it in effect amounts to arguing against (2b).

4. *Constraints on moral facts*

Now the situation is the following one. In order to block the argument, the moral realist appeals to the idea that evolution can *indirectly* care for truth (and, plausibly, does indirectly care for truth), namely, in the sense of truth providing a contribution to fitness.²⁸ We can now ask the question: what do the moral facts have to look like in order for this idea to work? What is required for such a contribution? Does the idea put any interesting constraint on the moral facts? For example, do they have to be ‘objective’ in any sense? Do they have to be causally efficacious? One could think that there are some interesting necessary and/or sufficient condition that could be derived from the moral realist’s idea of reliable representation of moral facts playing the role of a contribution to fitness. Which ones could that be?—Call this ‘the constraint question’.

The first, immediate answer to the constraint question that I will argue for is that the moral facts have to be capable of playing a certain role—call it role ‘R’. Role R can be described as follows:

- (R) The moral facts must be sufficiently reliably representable by the members of the group S and it must be possible that sufficiently many members of S (actively and passively) respect the moral facts.

The first part of (R) is the more important one, for our purposes. It is sufficiently reliable representability of the moral facts. In other words, the moral facts must not be too hard to detect. The members of the group have to have a quite easily available means of reliably representing the moral facts. Of course, they need not be infallible and may make mistakes about the moral facts sometimes. But a sufficient amount of accurate representation must be guaranteed. This the *epistemic condition* that the first part of (R) expresses.

The second, ‘practical’ part is rather obvious, but it is mentioned since it is required to understand the full story. It concerns the possibility of doing what the moral facts require and of showing the reactive attitudes of praise and blame, and all the rest of the reactive attitudes that have been discussed since Strawson, towards someone who is, or is not, acting in accordance with the moral requirements.²⁹ If it is a moral fact, for example, that Kim ought to help Jones (because Jones has been injured in some accident), then typically, Kim must be capable of doing what she is required. Call this the ‘active respecting of moral facts’. Equally, the members of the group must be capable of showing the positive and negative reactive attitudes towards the agents who act either in accordance or in violation of the moral facts, at least sufficiently often. Call this the ‘passive respecting of the moral facts’. Active

²⁸ For the present purposes, we can count “the four Fs” as what evolution *directly* cares for, as Graham puts it nicely: “feeding, fleeing, fighting, and reproducing” (Graham 2014: 19).

²⁹ Cf. Strawson (1962).

and passive respecting of moral facts have to be sufficiently possible, says the second, ‘practical’ part of (R).

Now the new story about how evolution could indirectly favor the reliability of C can be told. It is a *social reactive attitudes story* that fills the gap in the defense of moral realism against the evolutionary debunking attempt. As long as sufficient reliable representation in a group S is secured, and the members of S are sufficiently capable of actively and passively respecting the moral facts, it will in general be of advantage to any individual of S to correctly represent the moral facts and to act in accordance with them. In our example, crudely put, if Kim helps Jones, she will probably be recognized as a kind, supportive person and will receive praise and other positive reactions from the others. The moral attitudes (MAs) tend to favor such helping behavior, and so they indirectly contribute to the positive reactions received from other. (That the members of S are capable of reliably representing the *actions* performed by members of the group must also be secured, of course. They must be able to see whether the action fits the moral facts or not, at least sufficiently often. I take this to be granted since uncontroversial.) It is thus entirely unmythical how a reliably working cognitive process C governing MAs can contribute to fitness and survival via social interactions and the reactive attitudes therein. Linking positive reactions to conformity with the moral facts can do the job. MAs can benefit their subjects even if they do not benefit because they are true—they can benefit indirectly.³⁰

The advantages provided by the social mechanism of reactive attitudes in a group can be quite high, indeed, they can be extremely high. It all depends on how strong the reactive attitudes are. (As we all know, in fact they are quite high nowadays and include all kinds of social exclusion or punishments etc.)³¹

At this point, we can connect the discussion directly to Street’s writings. The just-mentioned social reactive attitudes account directly rebuts Street’s ‘implausible coincidence objection’. According to the moral realist, it is no mere coincidence, unexplicable or mysterious, that our

³⁰ Note that the proposal is not committed to any substantive normative claim that other third-factor accounts are committed to. For example, Enoch proposes the normative claim that survival and reproductive success are somewhat good (cf. Enoch 2010: 430). Wielenberg assumes the substantive claim that human beings have rights (cf. Wielenberg 2010). Normative claims like these are used by their proponents to argue for the reliability of our MAs. It is not entirely clear, however, how this argument is supposed to run. And, more importantly, it is very controversial whether one can rely on some such morality claim or not, since a quite serious suspicion of circularity or question-begging arises here. (Cf. Vavova 2014, Moon 2016, and Klenk 2017 for discussions on this point).

³¹ To connect the discussion directly to Street’s writings at this point: The consideration just given directly rebuts Street’s ‘implausible coincidence objection’. According to the moral realist, it is no mere coincidence, unexplicable or mysterious, that our MAs reliably represent the moral facts, and it is no mere coincidence that their reliably representing these moral facts is beneficial. Cf. Street (2006: 125).

MAAs reliably represent the moral facts, and it is no mere coincidence that their reliably representing these moral facts is beneficial.³²

The first or immediate answer to the question just given is not the end of it. We can go on and ask what further conditions have to be in place in order for truth making the envisioned contribution to fitness. For we can ask what the moral facts have to be like in order to be sufficiently reliably representable by the members of the group. For example, do they have to be ‘objective’?

The second answer to the constraint question then is that there is *no further interesting constraint on moral facts* that could be derived from sufficiently reliable representability. As long as the moral facts are sufficiently reliable representable, the social mechanism just described can work. Sufficient reliable representability may have its preconditions. And whatever is required for sufficient reliable representability has to be the case in order for the story of indirect contribution to fitness to work. If we leave to one side the second, practical part of role (R)—which is appropriate in the present context since it is not questioned by any party in the debate—what remains is simply the *epistemic condition of being sufficiently reliably representable*. Succinctly put, as long as the moral facts are *sufficiently epistemically accessible*, the moral realist can make good on the idea that evolution indirectly cares for the truth of moral attitudes.³³

In order to make this second answer plausible, let us run through a number of candidate conditions that might easily come to mind. Let us begin with *full objectivity*, i.e., total mind-independence. This is not a requirement for the story to work since if the moral facts are like some consequentialist think they are—i.e., an action’s maximizing *pleasure* or maximizing *desire* satisfaction, which is belief-independent but not entirely mind-independent—the story is in no way excluded. But fully objective, entirely mind-independent moral facts could fit the same bill. If the moral facts ultimately consisted in some primitive moral reasons relations, holding between some descriptive facts and certain responses (actions and/or attitudes), they could play the very same role—as long as the moral reasons facts are sufficiently epistemically accessible. Next, the moral facts could even be *subjective* in the sense of being *relative* to persons. Suppose that what is morally good or what one morally ought to do varies from person to person. Even this would not undermine the story. As long as the moral facts are sufficiently epistemically accessible (and practicable ...) acting in accordance with the moral facts would be likely to contribute to fitness. Finally, it does also not matter whether the moral facts are *reducible* to descriptive facts or *causally efficacious*. Perhaps, some condition like *supervenience* is

³² Cf. Street (2006: 125).

³³ It seems appropriate to call this condition ‘epistemic’ since reliable representation is sufficiently similar to knowledge, and knowability would of course be fine, too.

necessary for sufficiently reliable representation. But causal efficacy does not seem to be required (as the case of knowledge of mathematics shows).³⁴ And supervenience of the moral on the non-moral facts is a thesis that is widely accepted among moral realists, both naturalistic and non-naturalistic ones. To be sure, reliable representation (or even knowledge) is not for free. Arguably, it requires *some mechanism* which supports the reliable tokening of representations. But the important point is that it is not as though *the social role of moral facts* described in the story required some robust anchoring of moral facts in non-moral, descriptive facts. It is only the *epistemic accessibility* that might require this. And supervenience might be all that is needed for that.

The conclusion of these considerations is thus easily stated. *The one and only constraint* that the role (R) puts on the moral facts is their *epistemic accessibility*. As long as the moral facts are sufficiently reliably representable the social story of reactive attitudes can work and thus ‘implement’ a way of truth being highly significant to fitness—truth about the moral facts. Therefore, the moral realist has a plausible story about how to explain why premise (2b) in the argument above fails. And the story is fully in line with, and spells out, the idea of evolution caring *indirectly* for the reliability of our MAs. The evolutionary influence on our cognitive process C is fully compatible with its favoring the reliability of C. The MA’s *raison d’être* is tied to the reliability of their generating process C. Whether the moral facts have to be objective in any interesting sense can be left open—unless it is entailed by the epistemic accessibility of the moral facts.

If this conclusion is correct, it follows that *the evolutionary debunking argument does not pose any new problem for moral realism*, since the epistemic accessibility of moral facts is an old problem that has long been recognized and discussed. So interestingly, we have been lead back to the old epistemic problem, and no new problem arising from evolution has been discovered.

In sum, the evolutionary debunking argument that has been extracted from the Streetian considerations does not yield a sound argument against moral realism—not even if some objectivity requirement on moral realism is imposed. Once a proper statement of the argument

³⁴ Here is the right place to critically comment on Mogensen’s distinction between proximate and ultimate causes of MAs and corresponding kinds of biological explanation (cf. Mogensen 2015). The application of this distinction matches the distinction between the evolutionary process E and the cognitive process C insofar as the former concerns phylogeny and the latter concerns the individual. But there is no *causal* implication: the cognitive process need not have moral facts as proximate causes. So Mogensen’s focus is too much on causation, whereas the appropriate focus should be on reliable representation. In addition, the explanatory story on offer here gives moral facts a role in the *phylogenetic, evolutionary genesis* of our MAs—so an ‘ultimate’ role, if you like—and not (only) in the *individual’s* MAs—the ‘proximate’ role—, as Mogensen wants to have it (cf. Mogensen 2015: 197). Therefore, Severini’s criticism of Mogensen’s use of the proximate/ultimate distinction does not apply to the story on offer here (cf. Severini 2016).

has been formulated—argument A3—the crucial weakness becomes apparent: evolution’s pressure towards fitness does by no means exclude that reliable cognitive processes will be favored. Quite the contrary, it seems that evolution *indirectly* cares for truth, or at least can indirectly care for truth. The moral realist can tell a plausible explanatory story about how this influence could have developed, the story of reactive attitudes. Thus, the response to the debunking argument is no longer just an in-principle possibility of ‘indirect truth caring’ but a concrete, though sketchy, explanatory account of how this in-principle possibility can be realized. In addition, we can see what the moral facts have to be like in order to play their role in this story. They simply have to be sufficiently epistemically accessible, and no more. The evolutionary considerations, therefore, bring out the importance of the epistemology of moral facts. Some solution to this epistemological challenge has to be found. But this is not a new problem, and so the evolutionary argument does not yield any new constraint on moral realism.³⁵

References

- Alston, W. 1993. *The Reliability of Sense Perception*. Ithaca: Cornell University Press.
- Alston, W. 1986. “Epistemic circularity.” *Philosophy and Phenomenological Research* 47 (1) 1–30.
- Bergmann, M. 2006. *Justification Without Awareness*. Oxford: Oxford University Press.
- Brosnan, K. 2011. “Do the evolutionary origins of our moral beliefs undermine moral knowledge?” *Biology and Philosophy* 26 (1): 51–64.
- Burge, T. 2010. *Origins of Objectivity*. Oxford: Oxford University Press.
- Copp, D. 2008. “Darwinian skepticism about moral realism.” *Philosophical Issues* 18: 186–206.
- Crow, D. 2015. “A Plantingian pickle for a Darwinian dilemma: evolutionary arguments against atheism and normative realism.” *Ratio* 24 (2): 130–48.
- Cuneo, T and Shafer-Landau, R. 2014. “The moral fixed points: New directions for moral nonnaturalism.” *Philosophical Studies* 171: 399–443.
- FitzPatrick, W. J. 2015. “Debunking evolutionary debunking of ethical realism.” *Philosophical Studies* 172: 883–904.
- Dretske, F. 1995. *Naturalizing the Mind*. Cambridge: MIT 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.
- Goldberg, S. 2010. *Relying on Others*. Oxford: Oxford University Press.

³⁵ I am very grateful for helpful comments and discussions to Christoph Fehige, Thomas Grundmann, Susanne Mantel, and Peter Schulte. An earlier version of the paper has been presented in the Colloquium in Practical Philosophy at the Saarland University, Saarbruecken, Germany, organized by Christoph Fehige and Ulla Wessels, in January 2015. Many thanks to the organizers and participants for stimulating comments and valuable points of criticism.

- Goldman, A. 1979. "What is justified belief?" G. S. Pappas (ed.). *Justification and Knowledge*. Dordrecht: Reidel: 1–23.
- Goldman, A. 2008. "Reliabilism." *Stanford Encyclopedia of Philosophy*, ed. Ed Zalta.
- Graber, A. 2012. "Medusa's gaze reflected: a Darwinian Dilemma for anti-realist theories of value." *Ethical Theory and Moral Practice* 15: 589–601.
- Graham, P. 2014. "The function of perception." In A. Fairweather (ed.). *Virtue Epistemology Naturalized: Bridges Between Virtue Epistemology and Philosophy of Science*. Synthese Library 366: 13–31.
- Huemer, M. 2005. *Ethical Intuitionism*. London: Palgrave Macmillan.
- Hofmann, F. 2013. "Three kinds of reliabilism." *Philosophical Explorations* 16 (1): 59–80.
- Joyce, R. 2007. *The Evolution of Morality*. Cambridge: Cambridge University Press.
- Kahane, G. 2011. "Evolutionary debunking arguments." *Nous* 45 (1): 103–25.
- Kelley, T. 2014. "Evidence." *Stanford Encyclopedia of Philosophy*.
- Klenk, M. 2017. "Can moral realists deflect defeat due to evolutionary explanations of morality?" *Pacific Philosophical Quarterly* 98: 227–48.
- Kyriacou, C. 2016. "Are evolutionary debunking arguments self-debunking?" *Philosophia* 44: 1351–1366.
- Pollock, J. 1986. *Contemporary Theories of Knowledge*. Lanham: Rowman and Littlefield.
- Millikan, R. G. 1984. *Language, Thought, and Other Biological Categories*. Cambridge: MIT Press.
- Mogensen, A. 2015. "Evolutionary debunking arguments and the proximate/ultimate distinction." *Analysis* 75 (2): 196–203.
- Moon, A. 2016. "Debunking morality: Lessons from the EAAN Literature." *Pacific Philosophical Quarterly* 98: 208–26.
- Ruse, M. and Wilson, E. O. 1995. "Moral philosophy as applied science." In E. Sober (ed.). *Conceptual Issues in Evolutionary Biology*. Cambridge: MIT Press: 421–38.
- Senor, T. S. 1996. "The prima/ultima facie justification distinction in epistemology." *Philosophy and Phenomenological Research* 56: 551–66.
- Severini, E. 2016. "Evolutionary debunking arguments and the moral niche." *Philosophia* 44 (3): 865–75.
- Skarsaune, K. O. 2011. "Darwin and moral realism: survival of the iffiest." *Philosophical Studies* 152: 229–43.
- Strawson, P. 1962. "Freedom and resentment." *Proceedings of the British Academy* 48: 1–25.
- Street, S. 2008. "Reply to Copp: naturalism, normativity, and the varieties of realism worth worrying about." *Philosophical Issues* 18: 207–28.
- Street, S. 2006. "A Darwinian Dilemma for Realist Theories of Value." *Philosophical Studies* 127: 109–166.
- Smith, M. 1994. *The Moral Problem*. Oxford: Oxford University Press.
- Stich, S. 1990. *The Fragmentation of Reason*. Cambridge: MIT Press.
- Stroud, B. 1984. "The problem of the external world." In B. Stroud. *The Significance of Philosophical Skepticism*. Oxford: Oxford University Press: 1–38.

- Vavova, E. 2015. "Evolutionary debunking of moral realism." *Philosophy Compass* 10 (2): 104–16.
- Vavova, E. 2014. "Debunking evolutionary debunking." *Oxford Studies in Metaethics* 9: 76–101.
- Wielenberg, E. J. 2010. "On the evolutionary debunking of morality." *Ethics* 12 (3): 441–64.
- Wielenberg, E. J. 2014. *Robust Ethics*. Oxford: Oxford University Press.