



UNIVERSITY OF LEEDS

This is a repository copy of *Evolutionary Hypotheses and Moral Skepticism*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/135455/>

Version: Accepted Version

Article:

Isserow, J orcid.org/0000-0001-5900-8363 (2019) Evolutionary Hypotheses and Moral Skepticism. *Erkenntnis*, 84 (5). pp. 1025-1045. ISSN 0165-0106

<https://doi.org/10.1007/s10670-018-9993-8>

© Springer Science+Business Media B.V., part of Springer Nature 2018. This is a post-peer-review, pre-copyedit version of an article published in *Erkenntnis*. The final authenticated version is available online at: <https://doi.org/10.1007/s10670-018-9993-8>.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Evolutionary Hypotheses and Moral Skepticism

Jessica Isserow

School of Philosophy, Australian National University

This is a post-peer-review, pre-copyedit version of an article published in *Erkenntnis*. The final authenticated version is available online at https://link.springer.com/article/10.1007/s10670-018-9993-8?wt_mc=Internal.Event.1.SEM.ArticleAuthorOnlineFirst

Abstract

Proponents of evolutionary debunking arguments aim to show that certain genealogical explanations of our moral faculties, if true, undermine our claim to moral knowledge. Criticisms of these arguments generally take the debunker's genealogical explanation for granted. The task of the anti-debunker is thought to be that of reconciling the (supposed) truth of this hypothesis with moral knowledge. In this paper, I shift the critical focus instead to the debunker's empirical hypothesis and argue that the skeptical strength of an evolutionary debunking argument is dependent upon the evidence for that hypothesis—evidence which, upon further inspection, proves far from compelling. Following that, however, I suggest that the same considerations which spell trouble for the empirical hypotheses of traditional debunking arguments can also be taken to give rise to an alternative—and better supported—style of debunking argument.

1. Introduction

According to proponents of evolutionary debunking arguments (EDAs), explanations as to how our moral faculties might possibly have arisen can be recruited as a tool for undermining moral knowledge (Ruse & Wilson 1986, Joyce 2006, Street 2006).¹ EDAs traditionally rest upon a 'how-possibly' story that explains the emergence of our moral faculties by appeal to the fitness benefits that they conferred at some stage in our evolutionary past. Debunkers argue that supposing the rough accuracy of this story, it would be a lucky coincidence if the moral beliefs shaped by these faculties had stumbled upon the moral truths. Tracking facts about reproductive fitness is one thing, tracking moral truth is quite another. Insofar as our moral beliefs require luck in order to secure truth, they plausibly fall short of knowledge.

It has become customary for critics of EDAs to take the debunker's genealogy for granted. The anti-debunker's burden is usually thought to be that of reconciling the supposed truth of this genealogy with moral knowledge (see Brosnan 2011, Vavova 2014). My strategy is different. For the sake of argument, I am willing to grant to debunkers that their genealogies *do* cast the epistemic credentials of our moral beliefs into doubt.

¹ A note on terminology: I use 'moral faculties' to refer to the psychological mechanisms and processes that explain why we make the moral judgments that we do. EDAs are not aimed at moral beliefs directly, but at the psychological faculties that shape them; the influence of evolutionary forces upon moral beliefs themselves is usually thought to be indirect. (See Joyce 2006, pp.180-181; Street 2006, p.119.)

My own strategy applies a critical eye to debunkers' evolutionary hypotheses. To my mind, debunkers have not done enough to motivate the empirical plausibility of these accounts of moral evolution. In what follows, I shall argue that this spells trouble for their arguments, which are in danger of losing their skeptical teeth if they show, at most, that highly questionable scientific hypotheses are inconsistent with moral knowledge. Yet all is not lost for evolutionary debunking; for the very same considerations which serve to undermine the empirical hypotheses of traditional debunking arguments can also be taken to provide a new route to moral skepticism. Ultimately, I will suggest that these considerations give rise to an alternative—and better supported—style of debunking argument.

My first order of business will be to articulate the basic structure of EDAs (Sect. 2). I shall then direct my focus to the accounts of moral evolution that debunkers have proposed (Sect. 3). Such accounts do not survive critical scrutiny. Debunkers' genealogies are vulnerable to substantial empirical challenges, and these challenges should shake our confidence in them. In Sect. 4, I consider a genealogy which lacks those features that debunkers take to undermine moral knowledge. I then provide reason to suppose that we currently lack the empirical resources needed to adjudicate between debunking and non-debunking accounts of moral evolution (Sect. 5). These considerations, I will suggest, give rise to a novel epistemic situation. Drawing upon recent work in epistemology, I explain how this situation may itself give rise to moral skepticism (Sect. 6).

2. Evolutionary Debunking Arguments

The debunker's argument rests upon three key steps, which I will here outline in brief. First of all, the debunker adopts a particular *metaethical* position, generally a variety of moral realism according to which moral truths hold independently of moral evaluators. Second, the debunker posits some evolutionary *how-possibly story*, an empirical hypothesis as to how our moral beliefs might have come to be as they are through purely naturalistic processes. In general, the explanation will advert primarily to natural selection, but I defer discussion of the finer details here to Sect. 3. What's centrally important is that the evolutionary hypothesis makes (or purports to) appeal only to factors which are logically and causally independent of the supposed moral truths. Given the truth of that hypothesis, then, and the presupposition of realism, we are led to the third, epistemic stage of the argument; that any moral judgements we make are, if true at all, only true through some unlikely and epistemically problematic kind of coincidence. The debunker concludes that our moral beliefs (judgements) are either false or, at best, coincidentally true and certainly not worthy of being called 'knowledge'.

My foremost concern in what follows will be the empirical background to EDAs, but let me say a few more words about the metaethical and epistemic steps of the argument just sketched. It is common to take debunkers to operate under the assumption of moral realism (see for example, Kahane 2011, Brosnan 2011, Clarke-Doane 2012).² On this view, moral

² I should note that Street's (2006) true target is evaluative realism; she extends the skeptical challenge to all evaluative beliefs (when their contents are construed realistically), of which moral beliefs form a proper subset. Having noted this, I will, for ease of exposition, mostly formulate EDAs in terms of an epistemic challenge to moral beliefs. 'Evaluative' can be substituted for 'moral' so as to make the formulation fit Street's own articulation.

truths hold independently of our attitudes; what is morally right is right regardless of whether we happen to think that it is, or how we happen to feel about it. The extent to which EDAs rest upon realist assumptions is admittedly a matter of contention (see Joyce 2016). But realists are at least thought to be especially vulnerable to the challenge that debunkers pose. Russ Shafer-Landau (a realist himself) offers a helpful summary of the core concern: “It would be a miracle were there anything like a close correspondence between the deliverances of a faculty shaped by evolutionary pressures, and a set of moral truths whose contents (if realists are correct) are fixed independently of the outputs of this faculty” (2012, p.3). Since the moral realist lies squarely in the debunker’s line of attack, I will take her to be the primary target of these skeptical arguments.

On the matter of the epistemological step, it’s far from straightforward how we ought to translate the debunker’s talk of a “coincidence” into more familiar terminology. Some interpret debunkers as aiming to show that our moral beliefs are *insensitive* to the moral truths; that had the moral truths had been different, our moral beliefs would have been the same (Clarke-Doane 2012). Yet others have cautioned against formulating EDAs in these terms (Korman 2014), and alternative epistemic principles have been explored (Bogardus 2016).

It will not be necessary to take a strong stand on this issue for my own purposes. What is centrally important is that all EDAs trade upon empirical claims about moral evolution. The intended lesson of these evolutionary hypotheses may be that our moral faculties are insensitive to the moral truths. Or, it may be that these faculties are unreliable, or that they are unsafe. But no such lessons can reasonably be drawn if the evolutionary hypotheses themselves don’t pass muster. Insofar as these empirical claims are not plausible, they are not plausibly cause for epistemological alarm.

Though taking a strong stand will not be necessary, it will still be useful to have a reasonable reconstruction of the epistemological step to work with for the purposes of the ensuing discussion. To this end, I propose that we take EDAs to operate upon the following two assumptions:³

Epistemic Assumption #1

In light of the debunker’s evolutionary hypothesis, we would require an epistemically problematic sort of luck in order for our moral beliefs to be true.

Epistemic Assumption #2

Knowledge requires the absence of such luck.

The first assumption is suggested by the debunker’s claim that natural selection steered our moral faculties in directions having nothing at all to do with the moral truths (Joyce 2006, p.222; Street 2006, p.122).⁴ The accompanying concern is that there is no appropriate explanatory connection between our evolved moral faculties and the moral truths if the debunking genealogy is correct.⁵ Even if our moral beliefs had come to line up in the right way with the moral truths, this would, given their evolutionary origins, be “... a matter of

³ For similar diagnoses, see Shafer-Landau (2012) and Dunaway (2017).

⁴ All unattributed citations to Joyce and Street henceforth will be to their respective 2006 publications.

⁵ I borrow this way of framing things from Korman (2014).

sheer luck” (Street 2008, p.208). The second assumption is something that I don’t take to require much in the way of elaboration; it is widely (even if not unanimously) accepted that epistemic luck is incompatible with knowledge.

This admittedly broad characterisation of the debunker’s epistemological step will suffice for my purposes. My foremost concern is to put pressure on the empirical merits of the debunker’s genealogy—not to explore the finer contours of her epistemic reasoning. In what follows, then, I shall be granting to the debunker *arguendo* that particular genealogies which strongly suggest an absence of an appropriate explanatory connection between our moral faculties and the moral truths undermine moral knowledge.

3. Debunking Genealogies

The debunker’s empirical hypothesis plays an important role in her debunking argument. If that hypothesis is close to correct, then moral skepticism looms (given the assumption of realism). If, on the other hand, that hypothesis is lacking in empirical plausibility, then at best she will have shown us that a questionable genealogy is inconsistent with moral knowledge. And that’s hardly likely to leave the realist quivering in her boots.

Though Richard Joyce and Sharon Street are not the only philosophers to have recruited an evolutionary hypothesis in service of moral skepticism, their genealogies will form our focus in the ensuing discussion. Though my aim will be to put pressure upon the empirical plausibility of these how-possibly stories, it should be noted from the outset that my intention is not to show them to be *false*. The primary purpose of gesturing towards these issues is to arouse a sense of epistemic unease.

It should also be noted that Joyce and Street are not being singled out for target practice. Ultimately, I will argue that other accounts of moral evolution likewise encounter significant problems, and will suggest that these problems may very well be a symptom of a deeper issue—it may be that we simply aren’t in a position to place a great deal of our confidence in *any* account of our moral past. The ultimate aim of Sections 3-5 then, will be to show that there is good reason to regard our position with respect to our moral history as one fraught with epistemic limitations (the implications of which will then be discussed in Sect. 6).

3.1 Joyce’s genealogy

Joyce begins his how-possibly story by noting that evolution furnished us with various kinds of prosocial propensities that inclined us to cooperate with one another (pp.47-51). Though helpful, these inclinations were unreliable; our desires to cooperate can and often do falter. Given that the cooperative sphere is of far too great evolutionary significance to be at the mercy of such fickle inclinations, more effective motivational mechanisms were required. Joyce hypothesizes that these mechanisms were our moral faculties (pp.111-3).

It is important to note from the outset that Joyce understands this capacity for moralized thought in a cognitively rich way. Moral faculties don’t come cheap; they require “cognitive and conceptual sophistication” (p.94). The mere capacity to classify actions as right or wrong, for instance, does not suffice. One must also be able to grasp certain formal properties of

and relations between moral concepts (pp.67-9). Moral faculties require something above and beyond a suite of emotions, or emotional dispositions.

On Joyce's hypothesis, then, our moral faculties were selected for to boost our cooperative dispositions. Cooperation is more likely to transpire when one believes it to be a moral requirement. When early humans developed these moral faculties, they came to conceive of helping behavior as an inescapable obligation, the force of which could not be overridden by competing interests which favored non-compliance (pp.60-3). Moral faculties therefore earned their evolutionary keep by enabling individuals to block competing interests that would interfere with prosocial motivation from the deliberative sphere—they functioned as devices of “personal commitment” (pp.111-122). Given that prosocial behavior is often costly, moral faculties also came to serve as signs of interpersonal commitment; they offered early humans a means by which to convincingly signal their prosocial dispositions to others.

Though it is admittedly difficult to do complete justice to the impressively detailed hypothesis that Joyce develops, the above seems to me to be a faithful summary of its key claims. However, it strikes me that this hypothesis is vulnerable to a number of substantial challenges. My first concern pertains to Joyce's claim that our moral faculties were selected for because they furnished early humans with a capacity to commit themselves to prosocial behavior. There are alternative explanations for our capacity to commit to prosocial behavior—explanations which strongly suggest that there is no need to posit any moral faculties to do the work.

Robert Frank's (1988) research on the moral passions suggests that a capacity to commit to prosocial action merely requires a disposition to experience particular, motivationally powerful emotions. Frank proposes that “moral passions” evolved to control our temptation to favor immediate over long-term rewards (1988, p.82). An agent who is inclined to experience the aversive emotion of guilt, for instance, is better equipped with the psychological resources to resist the temptation to betray a hunting partner. This, in turn, earns her a reputation as a trustworthy individual, and thereby makes her more likely to be chosen as a partner in beneficial, cooperative ventures in the future (1988, p.17). According to Frank, these emotions, when activated, are accompanied by involuntary and observable symptoms—perspiration and blushing, for example. Over time, these symptoms came to be associated with the presence of the emotions with which they were reliably correlated. This enabled the moral emotions to function as hard-to-fake, honest signals of commitment and trustworthiness.⁶

In fairness to Joyce, he takes himself to be supplementing Frank's account rather than objecting to it (p.122). But their accounts diverge in a crucial respect: Frank's moral passions are not *moral faculties*. For Joyce, moral faculties are underwritten not only by emotions, but also by much more phylogenetically recent cognitive machinery—rich conceptual capacities (p.121) and language (p.84), among them. Frank's account therefore constitutes a genuine

⁶ Detailing the finer contours of hard to fake, honest signals is well beyond the scope of this paper. See Frank (1988) for an informative discussion.

alternative to Joyce's hypothesis; it does not posit moral faculties to explain our capacity to commit to prosocial behavior. Indeed, Frank's account suggests that there is no need to invoke moral faculties to do the explanatory work. *Pace* Joyce, we can explain our capacity for prosocial commitment by appealing to phylogenetically ancient emotional machinery—most centrally, a capacity to experience emotions such as aggression and guilt.

Of course, Joyce doesn't claim that making moral judgments is the only or even the best way to motivate prosocial behavior. He only claims that it helps enough to be selectively advantageous. But evolution is a satisficer—not a maximizer. Wholly new motivational mechanisms are unlikely to be selected for when existing structures (in this case, our emotional centres) can be manipulated to fulfil the same evolutionary end. Given that motivational mechanisms resembling Frank's moral passions were already available to enhance our cooperative dispositions, it becomes much less plausible that there were strong selective pressures favoring the evolution of moral faculties of the kind that Joyce envisages. This is not to say that wholly new or specialized motivational machinery is never selected for. (Joyce (pp.114-5) cites the example of the orgasm.) It is only to suggest that a hypothesis which takes such mechanisms to have been selected for is at a theoretical disadvantage compared to one which does the same explanatory work without them.

There are internal tensions within Joyce's hypothesis as well. Joyce claims both that (i) the need for cooperative motivation explains the emergence of our moral faculties, and that (ii) language is a precondition for morality (pp.84-5). This package of claims becomes problematic once we reflect upon the relationship between language and cooperation.⁷ Language is itself an expression of cooperation; it a subtle form of informational cooperation, and one that arguably requires a rich level of cooperation as a precondition for its emergence (Sterelny 2012b, p.105; see also Tomasello 2008; Hurford 2007). Moreover, alongside the evolution of language, we could expect other forms of cooperation to have developed; for through enriching communication, language also *enhances* cooperation, paving the way for reputation effects and the resolution of coordination issues (see Smith 2010).

Joyce's hypothesis that our moral faculties were an adaptation that enhanced cooperative motivation is therefore even more dubious once we factor in his claim that language is a necessary precondition for morality. A social world in which language is present is a social world "that has long-been cooperative" (Sterelny 2012b, p.105). Plausibly, a substantial motivation to cooperate was *already* present prior to the arrival of our moral faculties. This puts quite a bit of pressure on the hypothesis that those faculties were selected for to provide a much-needed boost to our cooperative dispositions.⁸

⁷ My criticism here is structurally identical to one that Kim Sterelny (2012b) has raised against Philip Kitcher (2011), whose genealogy shares similar assumptions to Joyce's.

⁸ A possible recourse for Joyce would be to argue that *certain forms* of cooperation are a prerequisite to language use, whereas other (more sophisticated) forms of cooperation require language—though this would require spelling out both why (i) more sophisticated forms of co-operation were plausibly needed, and (ii) the degree to which moral faculties can plausibly be taken to have been important for their emergence and persistence.

3.2 Street's genealogy

Street's genealogy takes off from the intuitive idea that the kinds of evaluative judgments that an organism is disposed to make have important consequences for its reproductive success (pp.114-6). It is, she suggests, "only reasonable to expect there to have been, over the course of our evolutionary history, relentless selective pressure on the content of our evaluative judgements, or... "proto" versions thereof" (p.114). Street likens these basic, proto-evaluative responses to the hard-wired response of a Venus fly-trap when it detects prey: "each may be seen as having the same practical point: to get the organism to respond to its circumstances in a way that is adaptive" (p.128).

Street's story posits continuity between the basic evaluative tendencies of our close ape relatives and our own. She understands chimpanzees, for instance, to be capable of some form of proto-evaluation, whereby they experience certain things in their world as counting in favor of certain responses on their behalf (pp.117-9). The uniquely human capacity for reflective and linguistic evaluative judgments is understood as "a relatively late evolutionary add-on, superimposed on top of much more basic behavioral and motivational tendencies". Thus, evolution is understood to have influenced the content of human evaluative judgments indirectly, by having first directly influenced the more basic evaluative tendencies upon which they were constructed.

This story is implausible in certain respects. Street claims that basic evaluative judgments are "analogous to" certain kinds of hardwired mechanisms that we see in other organisms. In each case, a particular adaptive response comes to be paired with particular cues in the environment. Whereas a "... reflex mechanism does this through a particular hard-wiring of the nervous system", an evaluative judgment "does this by having the organism experience a particular response as called for, or as demanded by, the circumstance in question" (pp.127-28). The analogy suggests that the content of evaluative judgments is fairly inflexible and cue-bound. Yet this is questionable on empirical grounds. Moral judgments in particular exhibit remarkable phenotypic plasticity; their contents vary with different cultural contexts (along with other environmental factors). Individual rights are a primary source of moral concern in some cultures. In others, communal duties are emphasized (Vauclair & Fisher 2011). Cultures also differ in the extent to which they endorse moral values pertaining to harm, fairness, loyalty, authority, and purity (Graham et. al 2011). Of course, basic evaluative tendencies may plausibly play a role in shaping the content of these moral judgments. But their role is not nearly as central as Street supposes. The impact of ontogenetic factors strongly suggests that the content of evaluative judgments is deeply sensitive to other influences.⁹

Another issue with Street's account is the omission of empirical detail. How-possibly stories typically constitute cumulative narratives; they specify the succession of environmental pressures and steps in the evolutionary trajectory that may have led to the

⁹ See Deem (2016), who also takes issue with this feature of Street's account, and considers a host of other important ontogenetic factors. See also Prinz (2009), who doubts that there is even canalization for very basic and general moral values.

development of a known phenomenon.¹⁰ Presumably, this applies to our evaluative capacities—no one thinks that we suddenly woke up one day making full-blown moral judgments. Although I don't mean to suggest that Street is committed to such a claim, there are nonetheless important gaps in her genealogy and among the most important of these are the steps that led from proto-evaluative tendencies to full-blown evaluative capacities.

Street's evolutionary story moves quickly from basic evaluative tendencies to the capacity for reflective evaluative judgment. Yet there are surely other psychological precursors to our capacity for reflective evaluative judgment that our proto-evaluative predecessors are unlikely to have had—a shared, verbal language (Joyce 2006), and improved impulse control (Sterelny 2012b), among them. Without these details filled in, a central claim of Street's genealogy becomes less plausible; it remains an open possibility that our reflective evaluative capacities are *not* a mere extension of more basic evaluative tendencies. Other psychological faculties may very well have constituted an important part of the foundation upon which our capacity for reasoned, evaluative judgment was constructed.

This is at once an empirical and a philosophical problem. Leaving these details unspecified impacts upon the account's empirical plausibility. But the omission of detail also suggests that the genealogy cannot so easily be recruited for skeptical purposes. Given that Street's how-possibly story is considerably minimal in its exposition of the *how*, it underdetermines the evolution of our evaluative capacities. And this, in turn, threatens its skeptical potential; for the realist may very well fill in these gaps in a manner that is hostile to the debunking project. (I will explore one promising way in which she might go about doing so in Sect. 4.2.)

3.3 Take away

I have argued that debunkers' genealogies are vulnerable to substantial empirical challenges. While such challenges are by no means decisive against their hypotheses, they should shake our confidence in them. However, one might worry that all of this is of cold comfort to the moral realist. It may very well be true that debunkers construct their skeptical arguments upon dubious empirical foundations. But perhaps no empirical hypothesis is friendly to the claim that we have moral knowledge. I will now suggest that this is not the case.

4. A non-debunking genealogy

Neither Joyce nor Street thinks that just *any* account of moral evolution is capable of undermining moral knowledge; they do not rest their skeptical hopes upon the claim that our moral faculties can be afforded some evolutionary explanation or other. Their arguments depend quite heavily upon the assumption that *a particular kind* of genealogy explains why we have the moral faculties that we have—one that has certain features which suggest a lack of an appropriate connection between our moral faculties and the moral truths.

¹⁰ Or at least, *plausible* how-possibly stories arguably do so. See Sterelny (2012a) for a compelling criticism of “key-innovation” models of human evolution.

On the most common sort of debunking hypothesis, the feature in question is the doxastically troubling influence of natural selection.¹¹ Since these selective forces seem to be a skepticism-producing feature par excellence for debunkers, a natural move is to explore the promise of a by-product account of moral evolution, according to which our moral faculties were a side-effect of other capacities that may or may not themselves have been selected for. It is to such an account that I now turn.¹²

4.1 Ayala's genealogy

Francisco Ayala (1987, pp.327-9) proposes that our moral faculties were a by-product of the following intellectual capacities:

- (A1) The capacity to anticipate the consequences of one's actions.
- (A2) The capacity to judge certain things as more desirable than others.
- (A3) The capacity to choose between alternative courses of actions.

Although Ayala believes that these capacities were likely the result of selection, he denies that moral faculties themselves were adaptive (1987, pp.236-9). Our moral faculties were simply a by-product of such capacities; they came along for the ride with our sophisticated intellectual repertoire. Ayala takes these intellectual capacities to be unique in their most developed form to human beings. Indeed, he denies any meaningful relation between our moral faculties and the capacities that underwrite the social behavior of other animals.

Insofar as moral norms are concerned, Ayala concedes that our moral codes often appear to be consistent with the biological dispositions of our own and other species (1987, pp.237-242). He emphasizes, however, that this congruency is neither necessary nor universal. Although moral codes cannot deviate too far from our reproductive interests (lest they promote their own demise), they are not influenced by our biological nature to any significant degree.

Like debunker's accounts, Ayala's genealogy is not without its problems. One issue is his denying any meaningful relation between the capacities that underwrite social behavior in other animals and our moral competence. Even if the three intellectual capacities that Ayala lists are jointly necessary for moral agency, it seems exceedingly unlikely that they are sufficient. Other cognitive capacities and emotional dispositions are needed to underwrite moral competence. Importantly, a number of these are found in other apes, including a capacity for empathy (Bekoff & Pierce 2009, Flack & de Waal 2000); and other emotions

¹¹ I do not mean to suggest that *any* adaptive explanation of a cognitive faculty precludes us from forging an appropriate explanatory connection between that faculty and a particular domain of truths. Plausibly, natural selection has not been a distorting influence upon various commonsense beliefs that we hold (e.g., beliefs in the existence of ourselves and other bodies). As Wilkins and Griffiths (2013) point out, fitness-tracking and truth-tracking don't seem to come apart here; true commonsense beliefs are plausibly linked to evolutionary success.

¹² I am here granting the assumption that human morality is a unified phenomenon that may properly be described as an adaptation, or a by-product. Sterelny and Fraser voice an appropriate suspicion of this assumption, noting that human morality is something of a "complex mosaic", the many elements of which plausibly "have different origins, respond to different selective forces [and] depend on different cognitive capacities" (2016, p.983). Though this does suggest that debunking genealogies tend to oversimplify things, it does not undercut the basic argument to be developed in what follows. If anything, it supports my contention that extant accounts of moral evolution are empirically questionable.

(Fessler & Gervais 2010); along with—arguably but by no means uncontroversially—a sense of fairness (Brosnan & de Waal 2003); and some capacity for social learning (Tennie et. al 2009). Considerations of evolutionary parsimony suggest that we should take many of these psychological capacities of our close ape relatives to be homologous to those which explain our moral competence. (See de Waal 2006; Boehm 2012.)

A further issue is that our biological nature shapes the norms that we endorse to a greater degree than Ayala supposes. Norms that arouse our emotional centres often have more cultural fitness (Nichols 2004). Learning biases can likewise have important effects upon norm transmission (Richerson & Boyd 2005).

4.2 Ayala's account: a non-debunking genealogy

Debunkers would be hard-pressed to use Ayala's hypothesis in service of their skeptical ends. The traditional method of debunking won't work here; the claim that there is no appropriate explanatory connection between our moral faculties qua products of natural selection and the moral truths no longer has any bite if our moral faculties were not the products of selective forces. Ayala's account seems to shelter our moral beliefs from these sorts of distorting influences.

There is another, more important respect in which Ayala's moral genealogy is unfriendly to the debunking project: it is considerably friendly to *the realist*. Many realists believe that the cognitive skills that enable us to discover moral truths are “more general powers of reasoning” (Cuneo & Shafer-Landau 2014, p.427), or abilities to subject ourselves to certain “norms of thought” (Nagel 2012, p.21). Ayala hypothesizes that our capacity to make moral judgments is a by-product of intellectual capacities such as these. This, in turn, lays a suitable foundation upon which to construct an appropriate explanatory connection between our moral faculties and the moral truths. The realist can argue that we have the ability to track moral truths because the capacities necessary to do so—general powers of reasoning and the like—were adaptive.

However, we shouldn't declare victory on behalf of the realist just yet. We must consider a pre-emptive criticism from Street, who is suspicious that by-product accounts evade her skeptical challenge (pp.142-44). Whatever capacity C the realist thinks our moral faculties are a by-product of, Street claims that she must either affirm or deny a relation between *that* capacity and the moral truths. Suppose that the realist chooses the former. The challenge now is to explain what this relation is. And at this stage, Street doubts that the realist can say anything other than this: C involves a basic ability to grasp moral truths, of which our current ability to grasp moral truths is a sophisticated extension. Yet this leaves her vulnerable to a dilemma. Either (1) the more basic capacity C arose by fluke (a remarkable coincidence), or (2) it was selected for in order to track moral truths (a scientifically implausible claim).

On closer inspection, however, Street's dilemma turns out to be a false dilemma. Let's take the second horn (2) first. None of the abilities which constitute Ayala's capacity C seem to have been selected for to track moral truths. Indeed, Ayala explains the selection for an ability to anticipate the consequences of one's actions (A1) by appeal to the cognitive abilities needed for tool construction. The construction of tools requires anticipating the uses to

which they will be put, conceiving of them “... as means that serve certain ends or purposes” (1987, p.238).

It’s also worth noting that there are a number of candidates for capacity C aside from those Ayala considers. On a recent by-product account developed by Michael J. Deem, for example, capacity C is thought to involve “the ability to discriminate, abstract from, and classify persons, events, and actions, the ability to anticipate and predict consequences of one’s own behavior and that of one’s conspecifics, and the capacity for belief formation” (2016, p.738). Once again, there seems to be no good grounds for supposing that these abilities were selected for to track moral truths. As Deem notes, these features of our cognitive architecture were arguably important for enabling early humans to navigate their way around ecologically varied or informationally opaque environments (see Godfrey-Smith 1996, Sterelny 2003).

It’s true that if the realist were to affirm a relation between C and the moral truths, then there would need to be an appropriate explanatory connection between the cognitive abilities to which C refers and a capacity for moral truth-tracking. But it’s unclear why Street thinks that this connection could only be explained by saying that C itself involves a basic ability to grasp moral truths. C can instead be regarded as a foundation upon which the skills that are relevant to moral competence are able to develop—skills such as abilities to form beliefs, reason abstractly, and anticipate the likely consequences of one’s actions.

As far as I can tell, none of the above seems to impale the realist upon Street’s first horn; nothing commits her to the claim that capacity C arose *by fluke*. Consider the suggestion that C might involve discriminatory abilities, together with a capacity for belief-formation. As I have noted, it would be implausible to suggest that *these* capacities arose by fluke; it is reasonable to suppose that there was a need for more sophisticated cognitive abilities in informationally opaque environments.

The by-product views canvassed above would therefore seem to escape the dilemma unscathed. In what follows, then, I will assume that such accounts constitute non-debunking genealogies; they allow us to forge an appropriate explanatory connection between our moral faculties and the moral truths.

5. Our epistemic situation

So far, my arguments seem to work to the advantage of the moral realist; she can opt for a non-debunking genealogy, and in doing so, avoid the skeptical consequences of a debunking one. Yet one might think that such arguments are at least somewhat to the advantage of the debunker as well. After all, it is still open *to her* to embrace a debunking genealogy.

It would be a happy result if we could simply let a thousand flowers bloom. But of course, we cannot simply permit the realist or the debunker to affirm whichever genealogy is to their philosophical tastes. In what follows, I shall argue that insofar as the mystery of our moral past remains a mystery, the most reasonable course of epistemic action may very well be for each party to suspend judgment upon which genealogy is correct.

5.1 The mystery of the moral past

Debunking and non-debunking genealogies alike are vulnerable to substantial empirical challenges. The proposal that I now want to explore is that these challenges may actually be a symptom of a deeper issue: that all of the evidence is not in (and perhaps is not likely ever to be in) when it comes to our moral past.

This proposal has been developed in detail by the evolutionary biologist Richard Lewontin (1998), who has long argued that we are drastically lacking in the sort of data that would be needed to understand the evolution of human cognition more generally. But the skeptical sentiment is by no means unique to him. (Richardson (1996) pushes a similar line.) Even those in the business of reconstructing the trajectory that our moral evolution might have taken are careful to acknowledge the epistemic limitations that we face. Joyce observes that arguments for and against human psychological nativism “seem to generate a peculiar amount of entrenched intellectual acerbity...it is a useful palliative to frequently remind ourselves of the common ground shared by all reasonable advocates: that we really don’t know yet” (2008, p.195). Phillip Kitcher similarly concedes that reconstructing our actual moral past “... is plainly beyond the evidence available”, noting that the “data are too sparse to screen out rival hypotheses about the sequence of events” (2011, p.11).

These remarks suggest a worrying under-determination of theory by available evidence. Insofar as we cannot gain empirical traction on a range of relevant issues (e.g., facts about genetic differences and ancestral environments), we cannot be sufficiently confident that certain features of human cognition (e.g., moral cognition) evolved one way rather than another (Lewontin 1998). The available evidence does not allow us to reliably adjudicate the competing hypotheses we have so far developed.

Indeed, we arguably do not even have access to the full space of reasonable empirical hypotheses about the evolution of moral cognition. Certainly, the possibilities have not yet been mapped out by anyone, even from a very coarse-grained point of view. What we have now is likely to be but a small scattering of possible hypotheses that, for all we know, may not be representative of the wider range of possibilities. One can certainly imagine that there are a great many creative perspectives on the evolution of moral cognition that have yet to be considered.

If this is the correct assessment, then it would seem that we are not in a position to place a great deal of confidence in *any* account of our moral past. Indeed, if I’m right and we don’t even have a full grasp of the space of hypotheses, then the most appropriate epistemic action we can take is not just to assign a low (but definite) probability to any one hypothesis, it’s to avoid assigning definite probabilities altogether.¹³ Insofar as our epistemic situation with respect to our moral past is one of highly ambiguous evidence and an apparent unawareness of the wider space of options, it would seem that the most epistemically appropriate course of action for us is to avoid taking a stand prior to receiving more information—to suspend judgment on the matter of how things actually went.

¹³ For discussion relating to probabilistic inferences in the absence of determinate evidence, see Joyce (2005).

It might be protested that the mystery of the moral past is not nearly as mysterious as the above discussion makes it out to be. In particular, one might suspect that we can rule out some moral genealogies owing to their manifest implausibility. My reply to this suggestion varies depending upon whether it is intended as a challenge to the letter or to the spirit of these accounts. Suppose firstly that it is a challenge to their *letter*. The claim is that—to take Joyce’s genealogy as a demonstrative example of its target—this *particular execution* of the hypothesis that our moral conscience is an adaptation is unconvincing. I very much agree with this line of response. That is to say, I agree that one could reasonably take issue with the finer details of any of these genealogies (for that claim is crucial to my argument). But this challenge still leaves open the possibility of a more plausible execution of the hypothesis that our moral conscience is an adaptation. Thus, the suggestion doesn’t undermine my claim regarding our current epistemic situation; for that situation is consistent with (and indeed, largely depends upon) the assumption that the moral genealogies available are amenable to improvement.

Suppose, however, that the suggestion is intended as an attack upon the *spirit* of these genealogies. The suggestion is that *any* account that construes human morality as an adaptation can be ruled out automatically—perhaps owing to the well-known controversy surrounding moral nativism. (See Prinz 2009.) For my part, this challenge seems to amount to unlicensed dogmatism. No moral genealogy is free from controversy; each is inevitably confronted by challenges that partially undermine its scientific plausibility. For some theorists, moral nativism constitutes a controversial empirical assumption, whereas for others, assuming the falsity of nativism begs important scientific and methodological questions. It would appear dogmatic, then, to view one’s own theoretical inclinations as providing decisive reason to completely rule out an entire class of hypotheses. And even if we were to permit such dogmatism, doing so would only narrow the domain of uncertainty. Even among the various non-nativist stories that remained, there would still remain considerable doubt concerning which *of these* is correct. And moral nativism is certainly not necessary to undermine moral knowledge; genealogies that do without this assumption can similarly carry skeptical implications. (Presumably, a hypothesis according to which our moral faculties were the result of random genetic drift would not spell victory for moral realism (see Kahane 2011, pp.111-2).)

Accordingly, it seems that some degree of epistemic caution is warranted on our part. The considerations above ought to disincline us from placing a great deal of confidence in any one account (or kind of account) of moral evolution. We must instead await a future time (if any there be) when we are better equipped in the way of evidence to properly adjudicate between these hypotheses. Given our current position, it would be premature—and indeed, epistemically irresponsible—to place our faith in (or wholly rule out) any one hypothesis without further empirical findings to back up our choice. If we are to be epistemically responsible, then it seems that we must suspend judgment upon which account of moral evolution is the right one.

5.2 The mystery of the moral past: *cui bono?*

Our epistemic situation with regard to the moral past is bad news for traditional debunking arguments. If debunkers are to transition from their empirical premises to their skeptical conclusion, then they must affirm that a particular kind of genealogy explains why we have the moral faculties that we have. But we can now see that this would be premature. Pace debunkers, we are simply not in a position to place a great deal of confidence in any account (or kind of account) of moral evolution.

Our epistemic situation is, however, also bad news for the realist. Earlier I suggested that it might be open to the realist to embrace a non-debunking genealogy. Yet we can now see that she is no more in a position to affirm the truth of non-debunking genealogy than the debunker is to affirm the truth of her own. If they are to be epistemically responsible, then both the realist and the debunker must suspend judgment as to whether or not our evolutionary history undermines moral knowledge. Yet where does this leave moral knowledge? I will now suggest that these considerations may leave it in a very bad place indeed.

6. A new debunking argument

I have granted to the debunker that if her how-possibly story turned out to be the how-actually story, then we would be led into wholesale moral skepticism. Unfortunately for the debunker, there are also non-debunking genealogies available, which, if true, put us in a position to deny that our evolutionary history undermines moral knowledge. But equally unfortunately for the realist, we cannot simply assume the truth of the latter hypotheses; we must suspend judgment upon which kind of genealogy is correct. In what follows, I suggest that this epistemic situation can be taken to support a new style of debunking argument.

6.1 Debunking: old and new

Our epistemic situation is as follows. We know that we are the products of a particular evolutionary history. But we do not know what kind of evolutionary history that is. Specifically, we do not know whether or not it is friendly to moral skepticism. I now want to propose that this epistemic situation opens the way for a new debunking argument (henceforth NDA):

NDA1. Our evolution could have occurred in a debunking or in a non-debunking way.

NDA2. If our evolution occurred in a debunking way, then we do not have moral knowledge.

NDA3. If our evolution occurred in a non-debunking way, then we do have moral knowledge.¹⁴

NDA4. We do not know whether our evolution occurred in a debunking or in a non-debunking way.

We do not know whether or not we have moral knowledge

¹⁴ I assume here that we do not have any *independent* reasons for doubting our claim to moral knowledge.

Notice that the NDA improves upon the traditional EDA in two key respects. Firstly, the new debunker does not find herself in the same precarious position as that of her predecessor, who rested her skeptical hopes upon a particular account of moral evolution. Unlike the traditional debunker, the new debunker doesn't assume the truth of any particular empirical hypothesis. This suggests a second advantage of the NDA: it better respects the open-endedness of scientific enquiry. The new debunker acknowledges affirms that further developments in the relevant areas of research are needed before we can place more confidence in any account of our moral past.

Yet one might wonder why the NDA counts as a *debunking* argument. Consider the position of the moral realist now. Having been presented with the NDA, she acknowledges that she does not know whether her moral faculties are the outcome of a debunking or a non-debunking evolutionary history. Thus, she happily acknowledges that she does not know whether she knows some moral proposition *m*. But this seems like a far cry from conceding that she does not *know* that *m*.¹⁵ When all is said and done, the NDA merely tells us that:

- (a) We do not know whether or not we have moral knowledge

It does not tell us that:

- (b) We lack moral knowledge

However, I think that the right conclusion to draw from the NDA may in fact be (b). In what follows, I make my case.

6.2 The return of moral skepticism

In the remainder of this paper, I want to explore a possible route from the NDA to moral skepticism. In particular, I want to propose that the realist's suspension of judgment upon the content of her evolutionary past provides her with a *defeater* for her moral beliefs. To motivate this suggestion, it will be helpful to first consider the following case:

Leia agrees to be a subject in an experiment. At some stage during the experiment, she will be asked to sit a test which requires her to identify the colors of shapes on a sheet of paper. Before the test, some of the subjects are given a drug. Others are given a placebo. No one is told which they received. Nor are they told how many received the drug or the placebo. The drug has the following effect: it leads subjects to form beliefs about the colors of shapes—beliefs which are generated in a way that has nothing at all to do with what color any given shape happens to be. The drug has no other symptoms, and the subjects have been informed about its effects. Leia takes a pill, and she begins the test. The first question asks her to identify the color of the square. She writes 'yellow'.

¹⁵ Fans of the KK-principle might beg to differ. There is a straightforward route from the NDA to moral skepticism for those who take knowing that *p* to entail being in a position to know that one knows that *p*. But I do not wish to rest my arguments for the NDA's skeptical potential upon the KK-principle, whose fans are in relatively short supply nowadays.

Assume that Leia's color vision is generally reliable in the absence of any drugs. Now, a question: is it plausible to say of Leia that she knows that y , 'the square is yellow'? Addressing this question will no doubt require us to attend to various features of Leia's epistemic situation. Let's focus upon her higher-order beliefs to begin with. How is Leia likely to respond were we to ask her, 'was your belief that y formed in a reliable way?' Since Leia has no clue whatsoever as to whether she received the drug, she would presumably refrain from either affirming or denying this; she would suspend judgment upon the proposition y^* , 'my belief that y was formed in a reliable way', and this would seem to be the appropriate response.¹⁶

Now, a further question: what (if any) implications might Leia's suspension of judgment upon y^* have for her belief that y ? It seems to me that taking this sort of attitude towards y^* has very bad implications indeed. In suspending her judgment, Leia suggests that she harbors significant doubts regarding the positive epistemic status of her belief that y . Indeed, she suggests that she would count herself rather lucky if that belief turned out to be true. In the face of such doubt and acknowledged dependence upon luck, I submit that it is incorrect to say of Leia that she really knows that y .

A suggestion owing to Michael Bergmann will help us to put some more flesh on the bones here. Bergmann's focus, like ours, is an agent who suspends judgment on the matter of whether a particular belief (or class of beliefs) that she holds was formed in a reliable way. Bergmann proposes to construe a suspension of judgment under such circumstances as grounds for epistemic defeat. In his view, an agent's suspension of judgment upon a higher-order proposition like y^* is a doxastic attitude that provides her with an undercutting defeater for her belief that y ; it provides her with grounds for thinking that her reasons for believing that y may not be indicative of its truth (2005, pp.426-7). Assuming (as is plausible) that knowledge requires an absence of defeaters, an agent who suspends judgment upon y^* cannot properly be said to know that y .

Bergmann's framework suggests the following plausible assessment of Leia's epistemic situation: Leia's belief that y falls short of knowledge in virtue of the undercutting defeater that she has acquired for that belief. The defeater in question is a 'mental state defeater'. It is a doxastic attitude—specifically, a suspension of judgment upon y^* —that reduces the positive epistemic status of the belief that y .

The appeal to mental state defeaters here may arouse the suspicion that our assessment is one that only epistemological internalists would find appealing. However, this is not so; the importance of mental state defeaters is something acknowledged by externalists as well. (See for example, Nozick 1981, p.196; Goldman 1986, pp. 62-63, pp.111-112.) Thus, Bergmann's suggestion is not especially partisan, and it should, I think, strike us as plausible. It is surely bad news for a belief, epistemically speaking, if one comes to believe that it may very well have been formed in an untrustworthy way.

¹⁶ For those who are unsure as to whether this would indeed be the appropriate response, we can add that Leia has never been a subject in experiment before, and that she is unaware of the typical distribution of placebos under such conditions.

None of this is to suggest that for any candidate belief to count as knowledge, one must additionally hold a higher-order belief that that belief was formed in a trustworthy way. (It is possible that an agent has *never even considered* the latter, higher-order proposition.) We are only assuming here that a suspension of judgment upon p^* , ‘my belief that p was formed in a trustworthy way’, constitutes an undercutting defeater for the belief that p . Thus, I do not claim that generally speaking, information regarding the positive epistemic credentials of one’s beliefs is necessary in order for those beliefs to qualify as knowledge. I am only suggesting that such information can have important implications for one’s knowledge-claims.

But what’s all of this to the moral realist? I now want to propose that she finds herself in a parallel epistemic predicament upon being presented with the NDA. Insofar as the realist suspends judgment upon the higher-order proposition m^* , ‘my moral beliefs were formed in a reliable way’¹⁷(in light of considerations of the kind discussed in Sect. 5.1), what is true of Leia seems equally true of her, *mutatis mutandis*. If the realist’s moral beliefs are to qualify as knowledge, then she cannot believe that those beliefs may very well have been formed in an untrustworthy way. In suspending judgment upon m^* , the realist acquires an undercutting defeater for her moral beliefs. And insofar as her moral beliefs are defeated, they plausibly fall short of knowledge.

One may object that there is a crucial disanalogy between Leia and the moral realist. Leia’s color judgments are formed on the basis of how things appear to her senses. She has no way to tell whether she has been given the drug, which operates precisely by exploiting the appearance–reality distinction. But a moral realist may very well take us to have direct intuitive access to the truth of moral propositions.¹⁸ If this is right, then perhaps we could check—“from the inside”, as it were—whether or not we are the products of a debunking genealogy. One need only consider gratuitous killing, intuit that it is obviously morally wrong, and triumphantly conclude that our faculties did not evolve so as to be unreliable.

In my view, however, the important parallel between Leia and the moral realist is not how they form their respective judgments, but their assessment of how trustworthy their mechanisms of judgment-formation are. So long as the trustworthiness of one’s moral intuitions (or faculties of moral insight or what-have-you) is something that can be brought into question, the basic challenge applies equally well to moral judgments. And it seems to me that moral intuitions can be brought into question by parallel considerations. At the very least, one ought to have less confidence in the deliverances of one’s faculty of moral intuition when participating in experiments with belief-inducing drugs, or upon being informed that there is *a priori* gas in one’s vicinity.¹⁹

It is of course open to the moral realist to adopt a Cartesian posture in response. She may insist that the intuition that p is morally wrong is such that it renders p immune to rational doubt. But to do so would be to saddle herself with (what is to my mind) an implausible

¹⁷ This is not to confuse two levels of explanation. M^* concerns moral faculties themselves, the reliability of which is cast into doubt by some genealogies but not others.

¹⁸ I thank an anonymous referee for raising this important challenge.

¹⁹ So-called *a priori* gas is said to induce a “...phenomenology of blatant obviousness” (Hawthorne 2007, p.205).

moral epistemology—and one that is radically out of step with what contemporary intuitionists believe. The vast majority regard intuitions as defeasible. It is commonly acknowledged that their justification-conferring power can be undercut by empirical considerations—facts about one’s cognitive biases, say. (Bedke 2008, p.264; see also Audi 2009, p.31).

A small caveat is in order prior to concluding. I have spoken of ‘us’ and of ‘our’ beliefs. But EDAs may very well affect different people in different ways. Some, for example, may take themselves to have good reasons to think that God superintended the process of evolution, ensuring that our moral faculties would be reliable. (Though see Baras 2017.) It is therefore possible that a suspension of judgment will not be the most epistemically appropriate course of action for everyone. If that is right, then my core claim is best thought of as a conditional one: *if* our evidential situation is such that the epistemically appropriate response is to suspend judgment upon m^* , then we are led to a skeptical conclusion. It would be brash to describe the case for the antecedent as dispositive. But I do believe that it presents a bona fide skeptical worry to be grappled with—and one that has clear advantages over traditional debunking arguments.

7. Conclusion

Traditional evolutionary debunking arguments rest upon the assumption that a particular kind of moral history represents our *actual* history. I have argued that the historical sketches that debunkers provide are highly questionable. In addition, I have suggested that the challenges to which these how-possibly stories are susceptible may ultimately be the symptom of a deeper epistemic predicament that we face: the fact that the available evidence does not allow us to reliably adjudicate the competing hypotheses that we have so far developed.

Unfortunately, it does not follow from this that moral knowledge is home and dry; for I have also proposed that this epistemic situation may itself lay the foundations for a new kind of skeptical argument. There is of course still hope on the horizon. Perhaps as further information comes to light, we will be justified in believing our moral faculties to be the products of a non-debunking history. For the time being, though, moral knowledge seems to remain uncomfortably hostage to empirical fortune.

References

- Audi, R. (2009). *The Good in the Right*. Princeton University Press.
- Ayala, F.J. (1987). The Biological Roots of Morality. *Biology and Philosophy*, 2, 235-252.
- Baras, D. (2017). A reliability challenge to theistic Platonism. *Analysis*, 77(3), 479–487.
- Bedke, M. (2008). Ethical Intuitions: What They Are, What They Are Not, and How They Justify. *American Philosophical Quarterly*, 45(3), 253-269.
- Bekoff, M. & Pierce, J. (2009). *Wild Justice: The Moral Lives of Animals*. Chicago, IL: University of Chicago Press.

- Bergmann, M. (2005). Defeaters and Higher-Level Requirements. *The Philosophical Quarterly*, 55(220), 419-436.
- Boehm, C. (2012). *Moral origins: the evolution of virtue, altruism, and shame*. NY: Basic Books.
- Bogardus, T. (2016). Only All Naturalists Should Worry About Only One Evolutionary Debunking Argument. *Ethics*, 126(3), 636-661.
- Brosnan, K. (2011). Do the evolutionary origins of our moral beliefs undermine moral knowledge? *Biology and Philosophy*, 26(1), 51-64.
- Brosnan, S.F. & de Waal, F.B.M. (2003). Monkeys reject unequal pay. *Nature*, 425, 297-99.
- Clarke-Doane, J. (2012). Morality and Mathematics: The Evolutionary Challenge. *Ethics*, 122(2), 313-340.
- Cuneo, T. & Shafer-Landau, R. (2014) 'The moral fixed points: new directions for moral nonnaturalism.' *Philosophical Studies*, 171(3), 399-443.
- de Waal, F. (2006). *Primates and philosophers: How morality evolved*. Princeton: Princeton University Press.
- Dunaway, B. (2017). Luck: Evolutionary and Epistemic. *Episteme*, 14(4), 441–461.
- Fessler, D. & Gervais, M. (2010). From Whence the Captains of Our Lives: Ultimate and Phylogenetic Perspectives on Emotions in Humans and Other Primates. In P.M. Kappeler & J.B Silk (Eds.), *Mind the Gap: Tracing the Origins of Human Universals*. NY: Springer.
- Flack, J. & de Waal, F. (2000). 'Any Animal Whatever: Darwinian Building Blocks of Morality in Monkeys and Apes.' *Journal of Consciousness Studies*, 7(1-2), 1-29.
- Frank, R.H. (1988). *Passions within Reason: the strategic role of the emotions*. NY: W.W. Norton & Company.
- Godfrey-Smith, P. (1996). *Complexity and the function of mind in nature*. Cambridge: Cambridge University Press.
- Goldman, A. (1986). *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Graham, J., Nosek, B.A., Haidt, J., Iyer, R., Koleva, S., & Ditto, P.H. (2011). Mapping the moral domain. *Journal of Personality and Social Psychology*, 101(2), 366-385.
- Hawthorne, J. (2007). A Priority and Externalism. In S. Goldberg (Ed.), *Internalism and Externalism in Semantics and Epistemology* (201-218). Oxford: Oxford University Press.
- Hurford, J. (2007). *The Origins of Meaning: Language in the Light of Evolution*. Oxford: Oxford University Press.
- Joyce, J.M. (2005). How Probabilities Reflect Evidence. *Philosophical Perspectives*, 19(1), 153-178.
- Joyce, R. (2006). *The Evolution of Morality*. Cambridge, MA: MIT Press.

- Joyce, R. (2008). Aversions, sentiments, moral judgments, and taboos. In W. Sinnott-Armstrong (Ed.), *Moral Psychology, Volume 1: The Evolution of Morality: Adaptations and Innateness* (195-204). Cambridge, MA MIT Press.
- Joyce, R. (2016). Reply: Confessions of a modest debunker. In U.D. Leibowitz & N. Sinclair (Eds.), *Explanation in Ethics and Mathematics: Debunking and Dispensability*. Oxford: Oxford University Press.
- Kahane, G. (2011). Evolutionary Debunking Arguments. *Noûs*, 45(1), 103–125.
- Kitcher, P. (2011). *The ethical project*. Cambridge, MA: Harvard University Press.
- Korman, D.Z. (2014). Debunking Perceptual Beliefs about Ordinary Objects. *Philosophers' Imprint*, 14(13), 1-21.
- Lewontin, R.C. (1998). The evolution of cognition: Questions we will never answer. In *An invitation to cognitive science: Methods, models, and conceptual issues* (107-132). Cambridge, MA: The MIT Press.
- Nagel, T. (2012). *Mind and Cosmos: Why the Materialist Neo-Darwinian Conception of Nature is Almost Certainly False*. NY: Oxford University Press.
- Nichols, S. (2004). *Sentimental Rules: On the Natural Foundations of Moral Judgment*. Oxford/New York: Oxford University Press.
- Nozick, R. (1981). *Philosophical Explanations*. Cambridge, MA: The Belknap Press.
- Prinz, J. (2009). Against Moral Nativism. In D. Murphy, & M.A. Bishop (Eds.), *Stich and His Critics* (167-189). Oxford: Blackwell.
- Richardson, R.C. (1996). The prospects for an evolutionary psychology: human language and human reasoning. *Minds and Machines*, 6, 541–557.
- Richerson, P.J. & Boyd, R. (2005). *Not By Genes Alone: How Culture Transformed Human Evolution*. Chicago, IL: University of Chicago Press.
- Ruse, M. & Wilson E.O. (1986). Moral Philosophy as Applied Science. *Philosophy*, 61: 173-92
- Shafer-Landau, R. (2012). Evolutionary Debunking, Moral Realism and Moral Knowledge. *Journal of Ethics and Social Philosophy*, 7(1), 1-37.
- Smith, E.A. (2010). Communication and collective action: language and the evolution of human cooperation. *Evolution and Human Behavior*, 31, 231–245.
- Sterelny, K. (2003). *Thought in a hostile world: the evolution of human cognition*. Oxford: Blackwell.
- Sterelny, K. (2012a). *The Evolved Apprentice*. Cambridge, MA: MIT Press.
- Sterelny, K. (2012b). Morality's Dark Past. *Analyse & Kritik*, 34(1), 95-115.
- Sterelny, K. & Fraser, B. (2016). Evolution and Moral Realism. *British Journal for the Philosophy of Science*, 68(4), 981–1006.
- Street, S. (2006). A Darwinian Dilemma for Realist Theories of Value. *Philosophical Studies*, 127(1), 109–166.

- Street, S. (2008). Reply to Copp: Naturalism, normativity, and the varieties of realism worth worrying about. *Philosophical Issues*, 18(1), 207-228.
- Tennie, C., Call, J. & Tomasello, M. (2009). Ratcheting up the ratchet: on the evolution of cumulative culture. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1528), 2405-2415.
- Tomasello, M. (2008). *The Origins of Human Communication*. Cambridge, MA: MIT Press.
- Vauclair, C.M. & Fischer, R. (2011). Do cultural values predict individuals' moral attitudes? A cross-cultural multilevel approach. *European Journal of Social Psychology*, 41(5), 645-657.
- Vavova, K. (2014). Debunking Evolutionary Debunking. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics. Vol. 9* (76-101). Oxford: Oxford University Press.
- Wilkins, J.S. & Griffiths, P.E. (2013). Evolutionary debunking arguments in three domains: fact, value, and religion. In G. Dawes & J. Mclaurin (Eds.), *A new science of religion* (133-146). London: Routledge.