

Reply: Confessions of a modest debunker

Richard Joyce

Penultimate draft of paper appearing in U. Leibowitz & N. Sinclair (eds.),
Explanation in Mathematics and Ethics (OUP, 2016) 124-145.

1. Introduction

Recent years have seen a burgeoning of discussion over so-called evolutionary debunking arguments of morality, but the fact that different debunking arguments have very different conclusions has allowed confusion to prosper. Sharon Street's well-known argument (2006), for example, has the conclusion that moral realism is probably false; thus critics often talk of "debunking arguments against *moral realism*." My own efforts, by contrast (see Joyce 2006, 2016b), aim for the conclusion that moral judgments are epistemically unjustified. That the judgments of a certain discourse lack justification, however, is consistent with these judgments being both true and mind-independently so; therefore my own debunking argument is entirely consistent with the truth of moral realism. I shall leave other would-be debunkers to fight their own street battles; in what follows I will focus on the epistemological debunking argument. The preceding chapters of part 1 of this book—by Justin Clarke-Doane, Toby Handfield, Hallvard Lillehammer, Folke Tersman, and Erik Wielenberg—provide much more that is worthy of careful consideration than I can possibly engage with here, and my reply will be neither methodical nor follow any principle of distributive justice. (The fact that a couple of the aforementioned chapters are barely referred to by me in what follows in no way indicates a dismissive opinion of them.) I hope, though, that a few crooked things can be straightened out.

An epistemological moral genealogical debunking argument purports to show that data about the genealogy of moral judgments may have an undermining impact on the justificatory status of those judgments. The kind of genealogical consideration that has been the focus of recent interest is the evolutionary perspective. The evidence suggests (the argument goes) that the human moral faculty is the product of natural selection: thinking in moral terms improved the reproductive fitness of our ancestors relative to the competition by allowing more stable and complex cooperative social structures to emerge.¹ What is noteworthy about such Darwinian explanations of human moral thinking is that they are no less plausible for the moral error theorist than for the moral realist, indicating that these explanations nowhere obviously imply or presuppose that any of our ancestors' moral judgments (or, indeed, anyone else's) were true. If (the argument continues) we have an empirically confirmed explanation of a group of beliefs that is consistent with their systematic falsehood, then our confidence in the truth of these beliefs should be dented.

¹ For much more detail, see Joyce 2006, 2013a, 2014a.

The first thing to note about this sketched argument is that the evolutionary perspective is, strictly, dispensable. Were we to explain our moral beliefs by reference to, say, developmental and socialization processes, then, so long as these processes similarly nowhere imply or presuppose that our or anyone else's moral judgments are true, the same epistemological conclusion could be drawn.

The second thing to note about the sketched general argument is that its epistemological conclusion (about confidence being "dented") is both vague and potentially very modest. This is so it can cover arguments bearing conclusions of differing strengths. Begin by noting that epistemic justification is a relative affair: Mary may be justified in believing that p while Peter is not justified in believing that p ; Mary may lose her justification for believing that p and then later regain it; and so forth. A strong epistemological debunking argument, therefore, would purport to show that even if most people's moral beliefs currently are justified, genealogical data could be brought forth that would remove that justification. An argument with even stronger ambitions would aim to show that this removal of justification would be permanent; nothing could reinstate it. A different kind of conclusion would be that nobody's moral beliefs ever have been justified. More modest arguments will allow the possibility that justification may be (re)instated. Another kind of modesty is exemplified by arguments that aim to show that although we might retain some grounds for maintaining moral belief (after accepting the genealogical data), our confidence in them should be reduced.

The recent literature against evolutionary debunking arguments of morality exhibits tendencies (a) to lump disparate views together, and (b) to take them to be more ambitious than they purport to be. Kelby Mason, for example, has explicitly but erroneously interpreted my debunking argument as striving for the conclusion that "all moral judgments are false" (2010: 775). Scott James also mistakenly construes my argument as concluding that "nothing in the world is objectively right or wrong or good or bad" (2011: 181). In this volume, Wielenberg devotes energy to rebutting an "Ontological Parsimony Argument," which has as its conclusion that no moral facts exist; but to my knowledge no one has advocated such an argument (though Wielenberg cites someone pushing it for the religious analogue).² Meanwhile, Clarke-Doane (also in this volume) takes the debunker to be aiming for the very strong conclusion that "*even assuming that our moral beliefs are (defeasibly) non-empirically justified*, learning that their contents fail to figure into the best explanation of our having them undermines them." (More on this later.) Even when an opponent distinguishes different types of debunking argument, by commencing to first rebut the aggressive versions (that is, the kinds that were never terribly plausible in the first place) and then moving on to criticize more diffident versions, he or she gives the

² I have discussed the prospects of a genealogical argument aimed at establishing a moral error theory in Joyce 2013*b*, 2016*a*, 2016*b*. In these papers I mention the possibility that Michael Ruse favors such an argument, but I do this more as an illustrative exercise than an actual attribution. For the record, I am a moral error theorist, and I do endorse the conclusions that are mentioned by Mason, James, and Wielenberg, but I would argue for those conclusions by means other than a genealogical debunking argument.

impression of the debunker retreating to an ever weaker position in a desperate bid to defend the argument. It is (as you can probably gather from my defensive tone) frustrating to one whose hopes and aspirations for the argument were fairly modest in the first place.

2. The role of realism

One of the bees in the bonnet of this chapter is that construing moral judgments as having *realistic* commitments plays little if any role in the epistemological debunking argument. It is best, then, if we address this immediately, but first I should deal with a few terminological preliminaries. I am happy to employ the characterization of *moral realism* suggested by Handfield (this volume, Ch. 4), according to which moral realism consists of two semantic claims—(i) that moral discourse should be interpreted literally and (ii) that it is truth-apt—plus two substantive claims—(iii) that at least some of the discourse is true and (iv) that this truth is mind-independent. The *error theorist* is an anti-realist who accepts (i) and (ii) but denies (iii) (and need not take a position on (iv)).³ The *success theorist* accepts (i)-(iii) (and need not take a position on (iv)); to be a success theorist, in other words, simply involves believing in moral facts to which we have access.⁴ If the success theorist also endorses (iv) then she is, of course, a realist. If she denies (iv)—if, that is, she believes in moral facts but holds that they are, in the relevant manner, mind-dependent—then we shall call her a “constructivist.”

Of the preceding five chapters of this volume, four see moral realism a crucial element of the debunking argument.⁵ This is understandable, since it is frequently Street’s “Darwinian dilemma” that is under discussion, and this dilemma explicitly raises a problem for the realist. On the other hand, all of these chapters also understand the argument in overtly epistemological terms, yet in vain does one search Street’s 2006 article for the words “unjustified” or “unwarranted” (or any equivalent) used in their epistemological sense. Street’s conclusion is that maintaining (i)-(iv) is problematic, and thus she recommends dropping (iv): she plumps for constructivism.⁶ But none of (i)-(iv) says anything about epistemic justification; that is an entirely different kind of argument.⁷ Rather than embark on the task of teasing these arguments apart, or trying to defend Street’s argument (something that I feel no urge to do), I judge it best if in what follows I stick to my more familiar ground, which is quite swampy enough to exhaust my attentions.

In order to start thinking about debunking arguments in a way that does not presuppose realism conceptually, nor target realism substantively, consider the following story:

³ Handfield is following Clarke-Doane (2012). See also Joyce (2007).

⁴ The term “success theory” comes from Sayre-McCord 1986.

⁵ The exception is Lillehammer (Ch. 6).

⁶ Elizabeth Tropman (2014) makes a plausible case that Street’s Darwinian dilemma is as much a problem for the constructivist as the realist. See also Fitzpatrick 2015 (footnote 10).

⁷ Therefore, incidentally, complaints that epistemological debunking arguments fail because they beg the question *against the moral realist* (see Fitzpatrick 2015) rather miss the mark.

Hans arrives in a foreign country, where he is confused by the local currency (call it the “loomah”). He sees the locals exchanging coins and colorful monetary bills, but Hans can’t understand the numerical characters thereon, and nor can he understand the language in order to ask. Hans chooses a stupid means of coming to a decision as to how much the various pieces of metal and paper are worth: he rolls dice. Holding a red bill, and rolling a ten, he decides the bill is worth ten loomahs. Perhaps by fluke this is true, but more probably it is not. In any case, though, Hans’ belief—“This bill is worth ten loomahs”—is, I’m sure we’ll all agree, unjustified. (Of course, we could add additional weird details to the story to produce a different result, but let’s not.) His belief is unjustified because rolling dice is not a process that tracks truths about monetary value.

Hans’ belief is debunked even though the proposition “This bill is worth ten loomahs” is not true or false by virtue of any *mind-independent* fact. Like facts about monetary value in general, its truth is constituted (in some manner) by human conventions and collective decisions. Hans knows this, the locals know this, and we know this. One can conclude that Hans’ belief lacks justification by virtue of the non-truth-tracking process of which it is a product, but this argument would nowhere presuppose *realism* about monetary value.⁸

Let us compare a moral analogue, choosing a version of metaethical constructivism: Ronald Milo’s contractarian constructivism (Milo 1995). According to this view moral facts are determined by the choices of a hypothetical idealized group of rational contractors. (The details need not detain us here.) It is not a realist theory, since the instantiation of moral facts depends on some “state of affairs [being made] the object of an intentional psychological state” (Milo 1995: 92). Consider the following story:

Hans has made a promise, but a competing consideration has arisen, such that keeping the promise will cause some harm, and he is unsure what he morally ought to do. Suppose Milo’s contractarian constructivism is true, and suppose that Hans knows this; he wants to act in accordance with the choices that would be made by a group of idealized rational contractors. But how would they choose? Hans chooses his favorite stupid means of deciding: he rolls dice. Rolling a ten, he decides that keeping the promise is morally correct. Perhaps by fluke this is true, but perhaps it is not. In any case, though, Hans’ belief—“Keeping the promise is correct”—is, I’m sure we’ll all agree, unjustified. (Of course, we could add additional weird details to the story to produce a different result, but let’s not.) His belief is unjustified because rolling dice is not a process that tracks truths about the choices that would be made by a group of idealized rational contractors.

The moral of the story is obvious: even assuming that moral realism is false (both conceptually and substantively), a moral judgment can be genealogically debunked. It will obviously be objected that an evolved moral faculty is nothing like rolling dice, but this is a matter for later discussion.

⁸ The epistemological argument, as I conceive it, does presuppose *cognitivism*. Handfield (this volume) agrees, seeing the argument as presupposing the two semantic components of realism (which of course are compatible with anti-realism). Oddly, though, as if in support of this claim he cites Guy Kahane’s contention that debunking arguments also presuppose one of realism’s *substantive* requirements: (iv) that moral truth is mind-independent. I discuss the prospects of genealogical debunking arguments that do not presuppose cognitivism in Joyce 2013b.

I have also taken the opportunity in these two stories to introduce the term “truth-tracking,” but only in a rough-and-ready kind of way. I haven’t assumed any precise theory of what it means; if a reader has such a precise theory in mind, then he or she should put it outside the backdoor while we approach the matter more directly in the next section.

3. Truth-tracking and possible worlds

I take it that most of us will accept that rolling dice is not a process that tracks truths about monetary value or the choices of idealized rational contractors. But what is it to “track truths”? A natural temptation is to answer by reference to counterfactuals. Most of the chapters to which I’m replying discuss the modal properties of *sensitivity* and *safety*, so I shall begin in like manner, though my real aim is to move away from these notions altogether. The counterfactual associated with the *sensitivity* of S’s belief that p is as follows:

SENSITIVITY: Were it false that p , S would not believe that p .

Consider another story of Hans:

Hans finds himself with the task of identifying random numbers as prime. Even numbers pose no problem, of course, and he can manage the smaller odd numbers well enough, but large odd numbers are tricky, and (for some reason) he has been given little time for calculation. So Hans chooses his favorite stupid means of deciding: rolling dice. Faced with the number 987659, and rolling a ten, he decides the number is a prime. Perhaps by fluke this is true, but more probably it is not. In any case, though, Hans’ belief—“987659 is a prime number”—is, I’m sure we’ll all agree, unjustified. (Of course, we could add additional weird details to the story to produce a different result, but let’s not.) His belief is unjustified because rolling dice is not a process that tracks truths about prime numbers.

The point of this Hans story is to illustrate how the natural temptation to understand truth-tracking by reference to SENSITIVITY leads to problems. Given that 987659 *is* a prime number, and necessarily so, what sense can we make of the counterfactual “Were it false that 987659 is prime, then Hans would not believe that 987659 is prime”? One response that I do not dismiss (though to be honest am not in a position to confidently assess) is to try to make sense of some such counterpossibles being non-vacuously true. (See Restall 1997; Vander Laan 2004; Brogaard & Salerno 2013.)

The counterfactual associated with the *safety* of S’s belief that p is, in the first instance, understood as:

SAFETY₀: Were S to believe that p , then p ,⁹

⁹ SAFETY₀ and SENSITIVITY are not equivalent contrapositives, since they are subjunctive (rather than material) conditionals.

but for various well-known reasons (see Williamson 2000; Pritchard 2005), this is better rendered as:

SAFETY: In all nearby worlds where S believes that p , then p .

A problem with appealing to SAFETY as a component of epistemic justification is that the notion of “nearby” is so vague in this context (see Baumann 2008). Handfield (this volume) recognizes this regarding the equivalent question of whether our moral beliefs might “easily” have been different; he writes that making fine-grained distinctions about whether claims of this sort are true “is not going to be fruitful.” Handfield thus embarks on the apparently more tractable task of comparing two possible worlds and asking whether one is closer to actuality than the other. Suppose that at the actual world we have norms about X which are imbued with an unusual kind of authority. At World₁ we lack X-norms altogether; at World₂ we lack X-norms imbued with this authority (but still have X-norms). Handfield argues that W₂ is closer to the actual world than W₁. His reasoning is that the closest worlds at which we lack norms altogether are ones where our ancestors did not need norms to solve problems arising from the fragility of cooperative endeavors (the threat of defection, etc.), whereas the closest worlds at which these norms are not imbued with special authority are ones where our ancestors were less cognitively flexible. He recognizes, of course, that less cognitive flexibility would be no small potatoes—“Perhaps the resulting organisms would not be readily recognisable as ‘human’”—but he nevertheless remains confident that such organisms “would be much more similar to us than organisms who evolved such that they had no norms at all regarding suffering, reciprocity, or the like.”

Well, maybe. But it’s pretty difficult to assess these matters when dealing with worlds containing counterparts who are hardly recognizable as human. Perhaps these counterparts came up with a different solution than norm-following to cooperation problems (or, rather, perhaps natural selection came up with a different solution for them). Handfield writes that “many other species have faced similar selective pressures, and many other species have something like norms to regulate conduct in these domains.” My first instinct is to deny this claim; I think *very few* other species have anything deserving the name “norms” governing their social behavior. My second instinct is to notice the vagueness of Handfield’s phrasing: “*something like* norms.” It’s difficult to know how much leeway this provides. Perhaps elephants have “something like” norms (see de Silva et al. 2011). Perhaps ants have “something like” norms. The important point, though, is that the implication of saying that something is “*like* a norm” is that it is *not* a norm. If, therefore, we allow that other species solve cooperation problems using something like norms (but which are not norms), then the same can be said of our counterpart ancestors at the closest possible worlds at which they did not need norms to solve problems arising from the fragility of cooperative endeavors; they simply solved these problems using “something like” norms instead. I’m not going to try to paint a picture of what these norm-like structures might be like (maybe Nature simply endowed them with a much bigger dose of fellow-feeling); it’s

enough to point out that it's not at all obvious that the worlds at which our counterparts employ them (rather than employ real norms) is much further away in modal space than the worlds at which our counterparts are less cognitively flexible than they actually are.

Clarke-Doane's argument (this volume) that our moral beliefs satisfy SAFETY raises similar uncertainties. He understands SAFETY to imply that that we could not easily have had different moral beliefs, and takes Street as his target debunker, citing her claim that "among our most deeply and widely held judgments, we observe many...with exactly the sort of content one would expect if the content of our evaluative judgments had been heavily influenced by selective pressures" (Street 2006: 116). The likely explanation of this observation, in Street's opinion, is moral nativism: the thesis that our faculty of moral judgment has been directly shaped by Darwinian forces. But now Clarke-Doane uses this very nativism to turn the tables on the debunking argument, because nativism implies that our moral judgments are biologically entrenched, thus "we could not have easily had different such beliefs," hence moral beliefs are safe.

I am unpersuaded by Clarke-Doane's argument. Does the nativist really maintain that our moral beliefs could not be easily different? Surely the very phenomenon of moral disagreement suggests that fairly mundane social and cultural pressures can affect moral judgments? And a large body of results from experimental psychology demonstrates how easily people's moral attitudes can be swayed by seemingly arbitrary environmental factors, such as the presence of an overflowing garbage can (Schnall et al. 2008) or bright lighting (Zhong et al. 2010). Moral variation, far from being something that the moral nativist denies, is a datum that he or she must accommodate. Clarke-Doane knows this, of course; in his 2012 paper he sensibly eschews the extreme nativist view that particular moral beliefs have been selected for, while allowing that "a more credible [nativist] view is that we were selected to have cognitive mechanisms that entail dispositions to form certain primitive belief-like [moral] representations in certain environments" (318). But how are we to reconcile this admission with the view attributed to the nativist that moral beliefs are "not easily different"? The answer seems to be Clarke-Doane's distinction between "basic" and "non-basic" moral beliefs, only the former of which are thought to be entrenched by Darwinian forces. The couple of examples of basic beliefs that he offers are "Pain is bad or pleasure is good" (Clarke-Doane 2012: 320) and "Killing our offspring is bad" (Clarke-Doane forthcoming). In order to do the job for Clarke-Doane, these beliefs have to be understood as involving *moral* goodness and *moral* badness; the modal intransigence of human *aversions* and *likings* is not at issue. But, understood as such, do these examples really stand up to scrutiny as basic moral beliefs?

Consider "Pain is morally bad and pleasure is morally good." I doubt that anyone believes this in such a coarse-grained manner. We think that some pains are worthwhile, that some are deserved, and there are many people who care about the pleasures and pains only of themselves (or in addition their friends and family, or in addition their countrymen, etc.). As for "Killing our offspring is morally bad," consider the many cultures that have allowed or encouraged infanticide. Perhaps one

could try to articulate more nuanced versions of these supposed basic moral principles, but I'll wager that one won't locate a moral belief that cannot be altered by cultural pressures with relative ease. But with how much ease? Here we face the same problematic vagueness encountered earlier when discussing Handfield's argument. Uncertainty over whether our moral beliefs could "easily" be different demands uncertainty over whether moral judgments satisfy SAFETY.

4. Truth-tracking and believable explanations

I must now confess that my heart hasn't really been in the foregoing arguments concerning SENSITIVITY and SAFETY, since I am not convinced that a straightforward modal analysis is the best way to understand the relevant sense of truth-tracking. Some people, it would seem, upon hearing the term "truth-tracking" are simply mesmerized into heading off to explore nearby possible worlds.¹⁰ But we know that such an exploration runs into immediate difficulties if the beliefs in question concern necessary truths. Gilbert Harman claims that the counterfactual test of whether *H* explains *E* "is only appropriate in certain contexts" (1986: 63); he explicitly does not apply it to mathematical beliefs, for example, and nor does he apply it to moral judgments. In the case of moral facts explaining moral judgments, Harman says that rather than an assessment of counterfactuals, "what's needed is some account of *how* the actual wrongness of [something] could help explain [someone's] disapproval of it. And we have to be able to believe in this account. We cannot just make something up" (ibid.).

Harman may be right that for some common-or-garden empirical beliefs a counterfactual understanding of truth-tracking is pertinent. However, I think it wise to be wary of any view holding that the following two propositions should be treated in a fundamentally different manner:

1. Hans' belief that the bill is worth ten loomahs is unjustified because it is based entirely on the roll of dice.
2. Hans' belief that 987659 is a prime number is unjustified because it is based entirely on the roll of dice.

The account of why Hans' belief is unjustified should be broadly the same in both cases (though an examination of counterfactuals might be relevant to supporting this account in one case and not the other). Harman's thought seems a reasonable one: in

¹⁰ I have been mesmerized myself; Clarke-Doane quotes me from 2001 making an unmistakable appeal to SENSITIVITY to back-up my moral skepticism. All I can do now is admit that enough time has passed since I wrote *The Myth of Morality* that I do brace myself for a little wince when a critic quotes a passage from that book. (Though not enough time has yet passed that I'm tempted to parrot A. J. Ayer who, when asked in later life to assess *Language, Truth, and Logic*, chuckled, "I suppose the most important of the defects was that nearly all of it was false!") I stand by the book's conclusions, and at least the spirit of its central arguments, but I now perceive many infelicities in the details—not the least of which was a tendency to vaguely discern several similar arguments and glom them together as one.

each case we lack any believable account of how the fact in question could help explain the outcome of the dice roll upon which Hans bases his belief. Note that what would count as an adequate explanatory account would not have to be *causal*. We do not, I take it, have a credible account of how 987659's being a prime number *causes* anything at all, but nevertheless we're very much in agreement that some belief-formation processes can reliably track this fact while others, like dice rolling, cannot. A corollary of putting matters in this way is that we should be suspicious of any argument for or against the claim that moral judgments track the truth that depends on the modal status of the content of those judgments. Whether the proposition "X is morally wrong" is a contingent truth or a necessary truth is not something upon which the question of whether the associated moral judgments are truth-tracking should turn.

Before proceeding I'll say something about the complaint that Harman's whole approach is misguided since it is not the job of moral facts to explain anything. Wielenberg endorses Russ Shafer-Landau's claim that the "functional role" of moral facts is not to "explain nonnormative phenomena, but rather [to] specify ideals, or standards that in some way must be met" (Shafer-Landau 2007: 323). However, I find curious the claim that a certain type of fact might have a "functional role." Compare facts about being a vacuum cleaner. Vacuum cleaners have the function of sucking up dirt; they don't have the function of explaining anything. Nevertheless, as concrete objects they *do* explain all sorts of things (the absence of dust on the carpet, why the dog is hiding under the bed, etc.). The fact that we may *use* moral facts to specify ideals and guide behavior (and so forth) hardly shows that the possibility that they don't play any explanatory roles should leave us unconcerned. (How astonished we'd be to discover that vacuum cleaners play no explanatory role at all!)

Elliott Sober voices a similar complaint against Harman's argument, but talks of the functional roles of propositions rather than of facts: "Normative ethical propositions have the job of telling us how we ought to act, not of explaining why we in fact act as we do" (Sober 2009: 141). This makes more sense to me than Shafer-Landau's (Wielenberg's) formulation, conceptually speaking, but I still question its truth. If asked the question "What is the job of propositions about vacuum cleaners?" the only answer I might tentatively hazard would be "To state truths concerning vacuum cleaners." The answers "To explain why we vacuum the house" or "To explain why we have beliefs about vacuum cleaners" would not leap to mind (and I may well even deny them). Nevertheless, this in no way diminishes my expectation that vacuum cleaners do play a robust explanatory role in the world. Similarly, were I to respond to the question "What is the job of moral propositions?" with the answer "To tell us how we ought to act," this in itself would not dampen any requirement I might place on moral properties to play a role in explaining our moral beliefs.

Let us return to Harman's question and apply it to the moral case: Do we have a believable account of how moral facts could help explain the mechanisms and forces that give rise to moral judgments? Notice the crucial presence of the word "believable." Dropping this constraint would make it the easiest thing in the world to establish that moral facts play a crucial explanatory and justifying role in our moral judgments, even assuming moral nativism. One could simply alight on any essential

aspect of the proffered evolutionary genealogy—for instance, that making moral judgments improved the probability of an ancestor’s genes being passed on to the next generation—and declare that *that* property is moral goodness (say).¹¹ This reminds me of C. L. Stevenson’s tongue-in-cheek example of an easy moral naturalism: “X is morally good” = “X is pink with yellow trimmings” (1937: 14). Stevenson’s faux-theory has much to recommend it: it allows for the existence of moral truths, moral properties, moral knowledge, and moral explanations. Yet for all its virtues this theory is obviously too silly to be taken seriously. By comparison, the theory that “X is morally good” = “X improves the probability of one passing on one’s genes” is, while maybe not quite as silly, still well beyond the pale of being taken seriously. The upshot of these observations is that the advocate of a debunking argument is certainly not claiming that moral nativism undermines moral justification *whatever the nature of moral properties (and whatever our relation to them)* may be. Believability constraints must be respected.

I’ll go out on a limb and assert dogmatically that we do *not* have a believable account of how moral facts could help explain the mechanisms and forces that give rise to moral judgments. Maybe we could come up with one, but we don’t have one now. Some people *think* that they already have believable accounts, but they’re mistaken. Is this such a surprising thought? Utilitarians, for example, believe that they can explain how moral facts relate to moral judgments, but anyone who isn’t a utilitarian (i.e., most people) thinks that they’re mistaken. Kantians also believe that they can explain how moral facts relate to moral judgments, but anyone who isn’t a Kantian (i.e., most people) thinks that they’re mistaken. And so on. It doesn’t really add anything shocking to this familiar picture to accept the view that *all* people who think that they can explain how moral facts relate to moral judgments are mistaken. Some of these theorists have a reasonable view of what the moral facts might be like but cannot account for our epistemic contact with them, while some of them can account for our epistemic contact but have an unreasonable view of what the moral facts might be like. It is the latter failing, for example, that sinks the pink-with-yellow-trimmings theory of moral goodness. It also sinks the improving-the-probability-of-passing-on-one’s-genes theory of moral goodness. And (admittedly less obviously) it sinks the maximizing-happiness theory of moral goodness. And it sinks the theory that identifies moral goodness with that which an idealized group of rational contractors would choose. And the theory that identifies moral goodness (relative to an individual) with that of which the individual’s culture approves. And it sinks quite a few other theories besides. In all these cases the description probably does succeed in picking out an actually instantiated property,¹² and in all cases there probably is a believable account of how these facts can explain our judgments concerning them, but in none of these cases does the property satisfy our desiderata for being a *moral* property. Disappointingly, perhaps, here I am not going to attempt

¹¹ And the same goes, *mutatis mutandis*, if it’s some kind of non-evolutionary moral genealogy on offer.

¹² Doubt hangs over the description concerning what hypothetical rational contractors would choose, which possibly fails to pick out any property at all. See Joyce 2011 for discussion.

to argue for this provocative declaration (I have done so so often before that it's growing repetitive); doing so would require engaging in hefty metaethical discussion. And that's the point that I really want to underscore: there is no getting away from the need for metaethics. I'll return to the machinery of the debunking argument in a moment; but first let me illustrate this point again by reference to the views of W. D. Ross, as presented by Lillehammer in this volume.

Ross provides the outline of a conditional debunking argument: he confesses that, in the absence of any positive justification-supplying considerations, knowledge about the origins of a moral belief might have an epistemologically undermining effect (1930: 14). But he doesn't endorse the conclusion of the debunking argument because he denies the antecedent: he thinks that (some) moral judgments are justified on the grounds of being *a priori* self-evidently true. This leaves one with the impression that Ross might concede that if his arguments in favor of moral intuitionism fail then the possibility of debunking is back on the table.

And fail they do. I am confident of the inadequacy (indeed, I would go so far as to say "the hopelessness") of moral intuitionism as an epistemological theory. When I reflect on the kinds of moral propositions that are supposed to be self-evidently true, I find myself not only able to doubt them, but actively thinking them false. (I am, after all, a moral error theorist.) I suppose it's possible that years of philosophy have warped my intuitions to the extent that I cannot identify self-evident moral truths as such when they're staring me in the face, but this seems to me unlikely. I don't seem to suffer such a malaise in general terms (I don't, for example, have trouble grasping the status of Descartes' *cogito*), and I don't seem to have experienced any emotional trauma or weird socialization process that might account for such an epistemic and/or moral failing on my part. And if no moral truths are self-evident to me—and there is nothing damaged in my faculties of *a priori* discernment, and I have reflected hard on the matter (as indeed I have)—then no moral truths are self-evident *tout court*.

There is much that could be said about the preceding thoughts (which I hesitate to call "an argument"), but here is not the place.¹³ The point to which I'm drawing attention is a dialectical one: the debunking argument has teeth only if certain metaethical arguments succeed. In the present case, Ross could try to defend his intuitionism; I would then attack it. Or if one didn't care for intuitionism but favored some other kind of (putatively) justification-supplying considerations, then we could argue about those instead. No one, though, thinks that genealogical empirical data alone can secure a skeptical victory; at most they battle alongside skeptical arguments of an *a priori* metaethical nature.

But does the genealogical empirical data play any essential part in the battle? If the debunking argument depends on certain metaethical theses being established in order to work, then might not these metaethical theses get the job done by themselves? Wielenberg argues that the genealogical considerations play no role at

¹³ I've argued against moral intuitionism in Joyce 2009 and 2014*b*, and numerous others have done so more diligently and at greater length.

all—that they leave the debate over moral epistemology “pretty much where it was.” In what remains I will assess this claim.

5. Belief pills and the Bishop’s defense

In the late 1990s there was some excitement in the popular media about neurological discoveries supposedly indicating the existence of a “god spot” in the human brain: a neural mechanism dedicated to religious experience. The tone of most of the discussion was that this somehow undermined theistic belief. I was living in England at the time, and I recall the Bishop of Oxford responding that, on the contrary, the presence of such a neurological mechanism is just what we should expect, since God would surely design our brains with the means to recognize Him. In honor of this obscure and dimly-recalled exchange, I propose to dub this kind of strategy against debunking arguments “the Bishop’s defense.”

Wielenberg employs such a defense against both religious debunking arguments (citing Alvin Plantinga) and against moral debunking arguments. In the latter case he refers to the possibility that a “third factor” might be responsible for both our moral judgments and the moral facts, ensuring a reliable correlation between the two. Echoing the Bishop of Oxford, one might say that the moral faculty’s evolving because it strengthened ancestral social cohesion is exactly what we should expect, since moral facts surely concern (*inter alia*, perhaps) behaviors that reinforce social cohesion.

Wielenberg constructs a Bishop’s defense against a make-believe illustrative example that I employed in my 2006 book: belief pills. Suppose you discover that one of your beliefs—let’s say, the one that Napoleon lost Waterloo—was the product of someone’s having slipped you a “belief pill”: a device that somehow (magically) made you form this belief. This discovery, I suggested (and still suggest), would and should undermine your former confidence in this belief. The belief may be true or may be false; the point of the story is to pump the intuition that the belief pill revelation would amount to your discovering that you have no grounds for leaning one way or the other. You should, then, remain agnostic about how Napoleon fared at Waterloo, at least until you gain access to some more reliable source of evidence.

But suppose (responds Wielenberg) that your taking the pill not only causes this belief but also, via an astonishing process of backward causation, interrupts the momentous events of June 18, 1815 in such a way as to influence Napoleon to lose the battle. In this case your belief, though pill-caused, would correlate reliably with the fact that it represents. This may be enough for the epistemological externalist to conclude that your belief is justified. (The internalist, though, would require that you have appropriate access to the fact that the pills work in this manner.)

I don’t think Wielenberg need invoke anything quite as *recherché* as backward causation to make his case. Instead let’s imagine that some of the people who go around slipping others belief pills are *benign* in that they give only pills that will result in *true* beliefs. Just for ease of reference, we’ll assume that other pill-pushers are *malign*: they give only pills that will result in false beliefs. (And perhaps some are

mischievous: they shake the bottle and distribute pills randomly.) A Bishop's defense in this case would be the claim that although your belief about Napoleon's losing Waterloo has turned out to be the product of a pill, it's okay because you were slipped the pill by a benign pill-pusher.

This defense, however, is only as good as the grounds you have for believing that your pill-pusher was benign. Knowing that your pill-pusher *might* have been benign is not going to quash the doubts you now have about the truth of your belief, and nor should it.¹⁴ Here probabilities make a difference, of course. If you know that 90% of pill-pushers are benign then your new attitude toward your belief that Napoleon lost Waterloo will be very different than if you know that 90% of pill-pushers are malign. Even in the former case, however, your Napoleon-lost-Waterloo belief, once held with solid confidence, would now have a question mark hanging over it.¹⁵

Since Wielenberg hasn't articulated a clear and complete picture of what the "third factor" might be, we are left with little more than the analogue of the hopeful claim: "The person who slipped you the belief pill might have been benign." Indeed he or she might have been. And indeed the factors mentioned in the evolutionary hypothesis about the genealogy of our moral beliefs *might* somehow reliably connect to a realm of moral facts. But—to invoke Harman's principle again—a mere "might" won't cut the epistemological mustard; we need to see a believable account. If ever such a believable account is forthcoming, the debunking argument will go away.

The kind of debunking argument that I favor, then, has modest ambitions; it is to create a burden of proof. Burden of proof arguments in philosophy are often tiresome (and indeed this one may be so!), but they are not unimportant, since the matter of whether a belief is justified or not may well turn on the outcome. Wielenberg doesn't accept this burden, however. He thinks that the skeptical argument is insufficient to cast doubt on beliefs that are "initially well-justified and firmly held." (Clarke-Doane also sets up the argument this way.) But it is exactly the former conjunct that is at issue, so I deny my opponent the right to assume it. The fact that many people *believe* their moral beliefs to be well-justified clinches nothing; people believe all sorts of silly things. By contrast, I don't object to the second conjunct that moral beliefs are "firmly held." If a career as a moral skeptic has taught me anything, it's how desperately unwilling people are to question their moral beliefs.

¹⁴ My mentioning "the *truth* of the belief" should not deceive the reader into thinking that we've slipped into discussing a debunking argument for the error theory. Whether our beliefs are true is an ontological claim; but whether *we have reason to doubt* whether any of them are true (whether, that is, we have some reason for suspecting that the error theory might hold) is an epistemological claim.

¹⁵ Externalist complications: Suppose your belief is the product of a belief pill, and 10% of local pill-pushers are benign and 90% malign. Is your belief the product of a reliable process? If we identify the process as "having been slipped a belief pill by a local pill-pusher" then the answer is negative, but if we identify it as "having been slipped a belief pill by a benign pill-pusher" then the answer is positive. This is an instance of the well-known "generality problem" for reliabilism (Conee & Feldman 1998), to which I know of no decent solution (and thus am inclined to conclude "So much the worse for reliabilism").

6. Epistemological conservatism and the burden of proof

There is, however, a view that would license the move, roughly speaking, from “These beliefs are firmly held” to “These beliefs are well-justified” (well, maybe not “well-justified,” but at least “*prima facie* justified”). This view is epistemological conservatism. The conservatist grants extant beliefs benefit of the doubt; they are considered “innocent until proven guilty.” Up to a point I think conservatism is all very well, but it should not be adopted blindly or unthinkingly as axiomatic epistemological bedrock; one should adopt it only if it is reasonable to think that we are pretty good believers—only if, in other words, it is reasonable to think that our beliefs are likely to be reliable indicators of the facts.¹⁶ But one doesn’t have to consider *all beliefs* as one monolithic block; we may be (and presumably are) more doxastically reliable in some domains than in others. Humans are, for example, pretty good at estimating the relative locations of nearby medium-sized dry goods, but notoriously bad at estimating certain kinds of probabilities when calculating risk. If one knows this, then it may be reasonable to grant the benefit of the doubt to beliefs of the former kind, but unreasonable to grant it to beliefs of the latter kind.

Consider an example (lifted from Christensen 1994). Suppose I am asked which has the larger population: India or the US. I find myself inclined to answer “India,” though I can no longer access any memories or other evidence in favor of this answer. Still, it’s reasonable for me to think that I have probably picked up this opinion from some reliable source, and so the mere fact that I find myself tentatively believing that India has the greater population provides that belief with some (defeasible) warrant. What grounds this warrant is a background assumption that the process that has led me to favor this belief was a (somewhat) reliable one; I assume that there exists a believable account of how India’s having a higher population than the US could help explain my coming to believe that India has the greater population, even though I cannot provide that account.

Suppose, though, that I subsequently discover that there is some doubt as to whether this process was a reliable one. If I were to learn that the process that led me to this belief is one that *implies* that the US has a higher population than India, then this, obviously, would undermine my belief in a very strong way. But doubt can arise for other reasons. If I were to learn that the process that led me to this belief is one that *is consistent with* the US having a higher population then this alone should undermine my confidence. Perhaps, for example, evidence comes forth that the reason I’m inclined to believe that India has a higher population than the US is that I was once told this by an Indian patriot who is known to wildly exaggerate such matters. Maybe on this occasion he told the truth, maybe he didn’t; in any case, this new information removes the background assumption that provided warrant for my original belief.

Something analogous very often goes on, I believe, in the case of moral beliefs. Many people unreflectively assume that their moral beliefs are likely to be good

¹⁶ See Goldman 1979: 30; Lycan 1988: 168.

indicators of the moral facts—that the moral facts help explain why they have the moral beliefs that they do—though they are unable to articulate how this could be. Moral beliefs, after all, are often accompanied by strong affective elements, such that the prospect of their being radically mistaken seems not merely improbable but *appalling*. (Even in the cool hour of metaethical reflection, for many people the feeling that the truth of the moral error theory would be simply appalling lingers in the background.) Many people, in other words, take the fact that they’re extremely reluctant to question a moral belief to provide grounds for confidence of the truth of the belief. This reluctance is likely to be strengthened if a person finds her epistemic peers also extremely unwilling to question a moral belief; she will assume (however inchoately) that surely someone, somewhere, somehow has access to the moral facts, which explains why her epistemic community, and thus she, has these beliefs. I’m not saying that people endorse a principle that underwrites this implication, merely that they often do as a matter of fact ground their confidence-level in this way. (Handfield does a good job of sketching these thoughts in his contribution to this volume; see also Joyce (2016b).)

A reliable belief-formation process is one that is incompatible with an error theory regarding the discourse that expresses the beliefs in question. Thus when one assumes that a type of belief is the product of a reliable process one assumes that an error theory for the associated discourse does not hold. If, therefore, one’s confidence in moral beliefs depends on the assumption that they are the product of a reliable process, a process in which moral facts play an explanatory role, then the confirmation of a genealogical hypothesis for those beliefs that in fact an error theorist *can* endorse should impact negatively on one’s confidence in those beliefs.¹⁷ This is why moral nativism can have an epistemological impact, for it is a plausible genealogical hypothesis that is entirely agreeable to the moral error theorist.¹⁸

The opponent of the debunking argument will respond that the genealogical hypothesis is in fact *not* compatible with an error theory; it will be claimed that although the error theorist *thinks* that she can happily endorse the hypothesis, in fact she cannot. It is like the atheist claiming to have an excellent secular explanation of religious belief (in Darwinian terms, say), and the Bishop responding that the atheist errs because this seemingly “secular” explanation in fact presupposes the existence of

¹⁷ I’m putting this in roughly internalist terms. The same arguments could be rendered, *mutatis mutandis*, in externalist terms. The need for expressing arguments in a manner that engages both the internalist and externalist has plagued my writing of this chapter. I encourage the reader to take many instances of “*mutatis mutandis*” as read.

¹⁸ This will not hold of evolutionary accounts of any type of belief, for some such accounts will presuppose the *truth* of the beliefs in question. The hypothesis that it was adaptive for our ancestors to believe that *p* may be plausible only under the assumption that *p* is in fact the case (in which case, clearly, the error theorist about *p*-discourse will be in no position to endorse the hypothesis). I have claimed in the past that this would be true of nativism regarding simple arithmetic beliefs, though Clarke-Doane (2012) disputes this. I do not have space here to argue against Clarke-Doane’s complicated argument, so I’ll take the simple way out and point out that in my work the arithmetic/mathematics example is just that—a passing illustration—and if it fails then I’ll still confidently stand by the general claim stated in the first two sentences of this footnote.

God as the creator of the universe. Here we again encounter a question of who bears the burden of proof.

There are in fact two burden-of-proof moments in this reasoning, so let's be clear. We tend to start out granting our moral judgments the benefit of doubt; we assume that they issue from a process in which the moral facts (of either a realist or constructivist nature) play some kind of explanatory role. Here (we tend to think) the error theorist shoulders the burden of proof. Now suppose empirical evidence comes forth regarding the genealogy of our moral judgments. Suppose, moreover, that this evidence and the genealogical hypothesis it supports are embraced by the error theorist. The previous assumption that our moral judgments issue from a process in which the moral facts play an explanatory role would be threatened. But (one might ask) can the error theorist *really* endorse the hypothesis? To the extent that one wishes to maintain confidence in moral beliefs, one will want to deny this. On the other hand, the error theorist *appears* to be able to endorse the genealogical hypothesis. It may seem that at this juncture an impasse has been reached. However, I think that a second burden-of-proof case can now be sustained. It seems reasonable to maintain that if one of the relevant parties posits more phenomena that are mysterious or yet-to-be-explained or hand-wavy, then that party bears a greater explanatory burden.

So: Does the error theorist, in endorsing the moral nativist hypothesis, leave any mysteries unaccounted for? As far as I can see, the answer is "No." The error theoretic perspective may be an unpopular or unattractive one, but it seems explanatorily complete in this respect. Indeed, perhaps the only mystery that the error theorist might have left us with is the question of why humans have so ubiquitously succumbed to this systematic mistake—an answer to which moral nativism supplies nicely.

By comparison, does the success theorist, in denying that the error theorist can endorse the hypothesis, leave matters unexplained? Arguably, the answer is a resounding "Yes." In order to deny the error theorist access to the genealogical hypothesis, the success theorist has to defend the place of moral facts in that hypothesis, and that is no straightforward task. (Note that I speak not of the work that a moral *realist* has to do, but a moral *success theorist*; the constructivist has plenty of metaethical work to do, too.) Just think of all the metaethical labor that would need to be undertaken to make the Bishop's defense favored by Wielenberg successful. Tersman, in his contribution to this volume, does a fine job of showing the challenges that face another "third factor" theory, advocated by David Enoch (2010). This kind of strategy, according to Tersman, must satisfy requirements that he calls "COHERENCE" and "EXPLAINING ERROR," and he expresses reasonable doubt that extant attempts have done so. Speaking more generally, success theorists must find a place in the genealogical account for items that plausibly satisfy desiderata for being *moral* properties. A mere glance at the last century of metaethics shows that this isn't easy.

The observation that success theorists have work to do doesn't imply that undertaking the labor is pointless or doomed to failure. Far from it; good luck to them. No argument currently under our consideration has as its conclusion the truth of the

error theory. Rather, what we are wondering about is whether the error theorist really can endorse the evolutionary hypothesis for moral belief. Until the case to the contrary is articulated and plausibly defended, a provisional positive answer seems reasonable. If it is reasonable to think that the versions of moral nativism currently available can be endorsed by the moral error theorist, then confirmation of such a hypothesis would be confirmation that it is reasonable to think that our moral judgments are the product of a process consistent with their falsehood. Therefore, to the extent that our confidence in our moral beliefs depends on the assumption that they are the product of a process in which the moral facts play an explanatory role, nativist hypotheses (of the kind under discussion) undermine that confidence.

7. Conclusion

If a philosopher claims to provide a positive account of how moral beliefs are justified, then the evolutionary epistemological debunking argument may not kick in. Rather, that philosophical account should be assessed on its own merits, presumably on familiar *a priori* grounds. (Perhaps, for example, the philosopher advocates, like Ross, a kind of intuitionism.) But if the account presupposes that moral facts help explain moral beliefs—rather than the account providing grounds for maintaining this—then the confirmation of a moral genealogy that the error theorist can endorse is of epistemological significance. I have suggested that arguments relying on epistemological conservatism in particular depend for their plausibility on an assumption of doxastic reliability. Some moral epistemologies explicitly rely on conservatism; others do so but less obviously.¹⁹ More than this, however, I would claim that conservatist assumptions dominate folk attitudes toward moral justification. People tend to stick with the moral beliefs with which they find themselves, resisting pressures for major revision, presuming their moral judgments to be justified due to these judgments *somehow* reflecting moral facts, supposing that the true explanation for their holding these beliefs must (somehow) exclude the moral error theorist. If, therefore, such a person acquires evidence for a genealogical explanation that does *not* exclude the error theorist, then his or her confidence in the truth of these beliefs should be dented.

Not wanting to over-reach, and aware that there are several not-quite-equivalent ways of expressing the conclusion, I shall end as I began: on the vague note of confidence in moral judgments being “dented.” An epistemological debunking argument, in other words, throws down a gauntlet. Its advocate says “Genealogical considerations have raised significant doubt about these judgments; if confidence is to be restored there is work to be done.”

¹⁹ Views that presuppose some kind of conservatist doctrine include Alvin Goldman’s maximalism, John Rawls’ reflective equilibrium, and Chisholm’s particularist methodology. See Kvanvig 1989 for discussion and references. Regarding morality more specifically, views with more-or-less central conservatist elements include Lycan 1986, Timmons 1999, Huemer 2005, and Enoch 2011.

References

- Baumann, P. 2008. "Is knowledge safe?" *American Philosophical Quarterly* 45: 19-31.
- Brogaard, B. & Salerno, J. 2013. "Remarks on counterpossibles." *Synthese* 190: 639-669.
- Christensen, D. 1994. "Conservatism in epistemology." *Noûs* 28: 69-89.
- Clarke-Doane, J. 2012. "Morality and mathematics: The evolutionary challenge." *Ethics* 313-340.
- Clarke-Doane, J. Forthcoming. "What is the Benacerraf problem?" In F. Patuat (ed.), *New Perspectives on the Philosophy of Paul Benacerraf: Truth, Objects, and Infinity*. Springer.
- Conce, E. & Feldman, R. 1998. "The generality problem for reliabilism." *Philosophical Studies* 89: 1-29.
- De Silva, S., Ranjeewa, A., & Kryazhimskiy, S. 2011. "The dynamics of social networks among female Asian elephants." *BMC Ecology* 11: 17.
- Enoch, D. 2010. "The epistemological challenge to metanormative realism: How best to understand it, and how to cope with it." *Philosophical Studies* 148: 413-38.
- Enoch, D. 2011. *Taking Morality Seriously*. Oxford University Press.
- Fitzpatrick, W. 2015. "Debunking evolutionary debunking of ethical realism." *Philosophical Studies* 172: 883-904.
- Goldman, A. 1979. "Varieties of cognitive appraisal." *Noûs* 13: 23-38.
- Harman, G. 1986. "Moral explanations of natural facts: Can moral claims be tested against moral reality?" *Southern Journal of Philosophy*, suppl. vol. 24: 57-68.
- Huemer, M. 2005. *Ethical Intuitionism*. Palgrave Macmillan.
- James, S. 2011. *Introduction to Evolutionary Ethics*. Wiley-Blackwell.
- Joyce, R. 2006. *The Evolution of Morality*. Cambridge, MA.: MIT Press.
- Joyce, R. 2007. "Moral anti-realism." *Stanford Encyclopedia of Philosophy*.
- Joyce, R. 2009. "The skeptic's tale." *Philosophy and Phenomenological Research* 78: 213-221.
- Joyce, R. 2011. "The accidental error theorist." In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics, Vol. 6*. Oxford University Press: 153-180.
- Joyce, R. 2013a. "Ethics and evolution." In H. LaFollette & I. Persson (eds.), *The Blackwell Guide to Ethical Theory*. Oxford: Blackwell: 123-147.
- Joyce, R. 2013b. "Irrealism and the genealogy of morals." *Ratio* 26: 351-372.
- Joyce, R. 2014a. "The origins of moral judgment." *Behaviour* 151: 261-278.
- Joyce, R. 2014b. "Taking moral skepticism seriously." *Philosophical Studies* 168: 843-851.
- Joyce, R. 2016a. "Evolution and moral naturalism." In K. J. Clark (ed.), *The Blackwell Companion to Naturalism*. Oxford: Blackwell.
- Joyce, R. 2016b. "Evolution, truth-tracking, and moral skepticism." In his *Essays in Moral Skepticism*. Oxford: Oxford University Press: 142-158.
- Kahane, G. 2011. "Evolutionary debunking arguments." *Noûs* 45: 103-125.
- Kvanvig, J. 1989. "Conservatism and its virtues." *Synthese* 79: 143-163.

- Lycan, W. 1986. "Moral facts and moral knowledge." *Southern Journal of Philosophy*, suppl. vol. 24: 79-94.
- Lycan, W. 1988. *Judgement and Justification*. Cambridge University Press.
- Mason, K. 2010. "Debunking arguments and the genealogy of religion and morality." *Philosophy Compass* 5: 770-778.
- Milo, R. 1995. "Contractarian constructivism." *Journal of Philosophy* 92: 181-204.
- Pritchard, D. 2005. *Epistemic Luck*. Oxford: Clarendon Press.
- Restall, G. 1997. "Ways things can't be." *Notre Dame Journal of Formal Logic* 38: 583-596.
- Ross, W. D. 1930. *The Right and the Good*. Oxford: Oxford University Press.
- Sayre-McCord, G. 1986. "The many moral realisms." *Southern Journal of Philosophy*, suppl. vol. 24. 1-22.
- Schnall, S., Haidt, J., Clore, G., & Jordan, A. 2008. "Disgust as embodied moral judgment." *Personality and Social Psychology Bulletin* 34: 1096-1109.
- Shafer-Landau, Russ. 2007. "Moral and theological realism: The explanatory argument." *Journal of Moral Philosophy* 4: 311-29.
- Sober, E. 2009. "Parsimony arguments in science and philosophy: A test case for naturalism_p." *Proceedings and Addresses of the American Philosophical Association* 83: 117-155.
- Stevenson, C. L. 1937. "The emotive meaning of ethical terms." *Mind* 46: 14-31.
- Street, S. 2006. "A Darwinian dilemma for realist theories of value." *Philosophical Studies* 127: 109-166.
- Timmons, M. 1999. *Morality Without Foundations: A Defense of Ethical Contextualism*. Oxford University Press.
- Tropman, E. 2014. "Evolutionary debunking arguments: Moral realism, constructivism, and explaining moral knowledge." *Philosophical Explorations* 17: 126-140.
- Vander Laan, D. 2004. "Counterpossibles and similarity." In F. Jackson & G. Priest (eds.), *Lewisian Themes: The Philosophy of David K. Lewis*. Oxford: Oxford University Press. 258-275.
- Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.
- Zhong, C., Bohns, V., & Gino, F. 2010. "Good lamps are the best police: Darkness increases dishonesty and self-interested behavior." *Psychological Science* 21: 311-314.