


Comments on Luke Kallberg's "Third-Factor Defenses of Moral Knowledge"

The Core Argument of the Paper, Summarized

Luke argues that moral realists' "third-factor" replies to evolutionary debunking arguments (EDAs) fail "because they predict a lack of epistemic safety for the beliefs that they defend" (1). (These are commonly held moral beliefs like "cooperation is morally good.")

Third-factor replies like Brosnan's (2011) allege that the formation of our moral beliefs may indeed be correlated with moral truths because the truth of some *third* factor raises the probability both that certain moral claims are true and that we believe them:



On Safety:

It's plausible to think that epistemic safety is "implied by knowledge" (3). Even if it's not one of the constitutive factors of knowledge, it is likely at least entailed by them. So Luke takes "total evidence against epistemic safety" to be "a defeater for knowledge" (3).

Basic idea of safety: For S to know that p, S *could not easily have believed falsely*.

SAFETY: S's belief that p via method M is safe iff in nearby worlds where S believes some similar p* via some similar M*, p* is true.¹ (3)

Luke argues that the assumption of a third factor like Brosnan's gives some evidence against the safety of our actual-world moral belief in p, e.g. "cooperation is morally good."

This is because, Luke argues, this third factor raises the probabilities in relevant nearby worlds of our holding other adaptive, similar, beliefs p* that are *false*, e.g.

- "Cooperation exclusively within our group is good"
- "Cooperation only when pleasurable is good" (4)

*What makes Luke's critique especially strong, and dialectically interesting, if it succeeds, is that it shows not just that the realist has *failed* to *defend* the existence of moral knowledge; Luke argues that third-factor replies provide some evidence, surprisingly enough, that we do *not* have moral knowledge.² This renders moot one point sometimes made by the realist in this debate: that the burden is not on them to prove we have moral knowledge but only to show defensively, in response to EDAs, that knowledge is indeed still possible after all.

¹ With a nod to Dunaway (2017), Luke emphasizes that "beliefs in nearby possible worlds can reflect upon safety even if they have different content from the belief under evaluation" (3).

² In fairness: Luke claims, modestly, that the limited evidence provided by his argument does not necessarily yield "total evidence" for a lack of safety and thereby a lack of knowledge. But it provides some evidence in this direction.

1. First, A Point of Difference: Do Third-Factor Replies Beg the Question?

Luke's View: "I doubt that question-begging is a problem here because it is the task of the debunker to formulate an argument based on plausible premises. It is unreasonable for the debunker to require her interlocutors to set aside all moral beliefs *before* engaging with the EDA" (3).

My View: I argue elsewhere that third-factor replies are indeed question-begging.³ I think the problem is not just that realists draw upon moral premises; the problem is the particular relationship in these replies between the premises and the conclusion that is supposed to be rendered *more probable* by those premises.

Why? In a context in which it's granted that certain beliefs are at least reasonable, but we have been given additional reason to question whether those moral claims are certain to be true and mind-independent, I think it is question-begging to repeat those beliefs (without also adding any *new* evidence) in an effort to deflect that worry.

2. Is Luke Right that a Third-Factor Reply Predicts a Lack of Safety for Moral Beliefs?

Luke: "In order to even allow for epistemic safety, it will need to be the case that the third factor doesn't raise the probability of relevantly similar *false* beliefs in epistemically relevant possible worlds" (4)

(And yet, he argues, the assumption of the third factor does just that.)

Two Questions:

- What is the range for how "similar" in content p^* has to be to p in order to be relevant to the safety of S's belief in p ?
- In nearby worlds where S believes p^* , why is Luke able to assume that p^* is *false* after all (in at least enough of those worlds), thereby rendering S's actual belief in p unsafe? The third factor raises the probability of p . Yet how *much* does the probability that p^* is false have to be raised for us to be confident p^* is false? What are our priors??

Suppose, as Brosnan (2011) invites us to do, it's true in this world (and in most? all? nearby worlds) that "it helps our group for us to believe that cooperation is morally good." Does this third-factor supposition render the target moral belief in p , that "cooperation is morally good" *unsafe*?

Well, let's start checking nearby worlds. In one, S forms the belief that p^* , that "cooperation only when pleasurable is good," (Luke's example) on the basis of method M^* , which is presumably forming beliefs using some method that has been adaptive for S's species. Is p^* false in that world, rendering S's belief (in this world) in p unsafe? (I'm just not sure!!)

In a dialectical context where we *could* assume it's true in all worlds e.g. that "cooperation is morally good" and false in all worlds that "cooperation only when pleasurable is good," then a third-factor reply would hardly seem necessary or appropriate. We would already be assuming a (pretty strong!) form of moral realism according to which moral principles are necessarily true. If we were allowed to assume that, then a third-factor reply definitely would be question-begging, especially as a way to provide comparative support for moral realism over views that take normative facts to be mind-dependent (and thus to vary in other nearby worlds in tandem with attitudes).

³ Michelle M. Dyke, "Bad bootstrapping: the problem with third-factor replies to the Darwinian Dilemma for moral realism," *Philosophical Studies*, Online first May 08 2019, <https://doi.org/10.1007/s11098-019-01301-4>