

This is a pre-publication draft of a paper forthcoming in *Episteme*; please cite the official version.

Epistemic Circularity, Reliabilism, and Transmission Failure

Patrick Bondy

Cornell University

patrbondy@gmail.com

1. Introduction

Philosophical orthodoxy has it that begging the question is a Bad Thing. Pointing out that someone has begged a question in defense of a claim is an excellent way to get that person to either offer a different argument in support of the claim or else retract the claim at issue. Given that orthodox view of begging the question, it is surprising to find so many defenses of the legitimacy of certain kinds of question-begging arguments in the literature. Russ Shafer-Landau (2009, p.193), for one, argues that “question-begging claims and arguments are ones that agents may sometimes be justified in believing.” Richard Robinson (1971, p.117), for another, memorably argued that the charge of begging the question is “nearly always a muddle, or improper, or both.”

To my mind, to give question-begging arguments is only to be dogmatic in a roundabout fashion. Of course, to take on the whole range of question-begging arguments and the literature surrounding them is too much for a single paper. The scope of this paper is therefore limited to the epistemic aspect of the sort of question-begging arguments that have come to be known as “epistemically circular.” The question of the legitimacy of epistemically circular arguments has received some serious attention in the recent literature.¹ I want to argue that epistemically

¹ See, for example, Vogel (2000, 2008), Cohen (2002, 2010), Weisberg (2010), and Kallestrup (2009; 2012) for a sampling of the literature. Discussions of epistemic circularity often go under the heading of the “problem of easy knowledge,” as epistemically circular arguments are one prominent sort of easy knowledge argument. Moorean arguments against skepticism are another.

circular arguments are always illegitimate. In what follows, I begin by briefly explaining epistemic circularity and three natural reactions to it, and I set out what I take to be the way to argue against it. The argument appeals to the notion of the failure of transmission of justification across a known entailment: the premises of epistemically circular arguments, I argue, fail to transmit justification to their conclusions. After that, I consider some reliabilist defenses of epistemically circular arguments, and I argue that those defenses do not succeed.

A couple of provisos and definitions are in order before proceeding. First of all, discussions of epistemically circular arguments, especially in the context of discussions of the problem of easy knowledge, are often carried out in terms of knowledge. The object, in such discussions, is to discover whether one can come to know a proposition on the basis of an epistemically circular argument. I am primarily interested, however, in the question of justification – that is, the question of whether one can come to be justified in believing the conclusion of an epistemically circular argument on the basis of that argument. Much of what I say will carry over to the knowledge-question, especially if we take justification or warrant to be a necessary condition for having knowledge, but knowledge is not my primary target.

Second, I take it that the correct analysis of epistemic justification will be internalist, in the sense that the reasons that justify a belief for a subject must be cognitively accessible to her, and some of the discussion which follows will no doubt appear to be too internalist to be acceptable to a reliabilist. However, the arguments in this paper are neutral with respect to the internalism/externalism divide; in particular, the arguments here are tailored to be amenable to externalist reliabilists. I will respond to the charge that I'm being too internalist when it comes up in relation to Alston's argument.

Third, much of this paper has to do with reliabilism about justification. The heart of the reliabilist position is that beliefs are justified when and because they are reliably produced (and sustained). I take that to mean that the process which produced them is a reliable one.²

Reliability, on this picture, is a property of belief-forming processes, when they produce a sufficiently high ratio of true to false beliefs. Two clarifying assumptions about the formulation of reliabilism will be helpful for keeping things clear in what follows.

(1) It is not only the actual world that determines a process's output ratio. Processes that would counterfactually have a high true-to-false belief output-ratio can fail to do so in the actual world. There are different ways to specify the relevant range of possible worlds in which a belief-forming process must be reliable, in order for it to produce justified beliefs. For the purpose of this paper, I take the relevant set of worlds to be the set of "normal" worlds, i.e. the set of worlds which are mostly the same as the way we believe the actual world to be. So if a process has a high true-to-false output ratio in worlds that are mostly the way that we believe the actual world to be, then according to this type of reliabilism, that process produces justified beliefs.³

(2) Serious reliabilists need to have a response to the generality problem, which is the problem of identifying the correct level of generality at which to describe belief-forming processes for the purpose of assessing their reliability.⁴ Because my goal here is not to knock

² The arguments in this paper could be recast to suit other versions of reliabilism, I believe, such as what Hofmann (2012) calls "reasons reliabilism". However, the recasting of the arguments will be cumbersome, and process reliabilism is the paradigmatic form of reliabilism, so it is the process reliabilist defense of epistemically circular arguments that I consider here.

³ Goldman (1986) adopted normal-worlds reliabilism, though he soon dropped it. See Goldman (1988; 2008) and Pollock and Cruz (1999) for discussion of problems with it. I think that normal-worlds reliabilism is superior to other formulations, but I will not argue for that claim here. I adopt it only as a simplifying assumption. The reliabilist defense of epistemic circularity considered in this paper could be recast to handle other relevant sets of worlds.

⁴ See Conee and Feldman (1998), Alston (1995), Comesaña (2006), Bishop (2010), and Conee (2012) for some important treatments of the generality problem.

reliabilism, but rather to argue against the legitimacy of epistemically circular arguments, I take it for granted that there is a response to the generality problem available to the reliabilist. I will assume that something like Alston's (1995) response is correct: we identify the relevant type for evaluating a token belief-forming process by the type of input that the token process takes, and the propositional content of the belief that it generates, and the way that the token process works.

2. Epistemic Circularity

2.1. What it is

An epistemically circular argument is one which (i) concludes that some belief-source is reliable, and (ii) employs premises that have been produced by that belief-source in support of that conclusion, where (iii) the truth of the premises has only been checked by the belief-source the reliability of which the argument was supposed to establish. Vogel's (2000) Roxanne case, frequently cited in discussions of epistemic circularity, illustrates the kind of argument at issue. The case goes roughly as follows. Roxanne wonders how much gas is in her gas tank. She looks at the gauge, and notes that it reads X. She therefore forms the belief that the tank is X full. Because the gauge is in fact working correctly, Roxanne's belief is true. She repeats this process many different times. Now, she never checks the reliability of the gas gauge by any other means; the only way she ever knows how much gas is in the tank is by what the gauge reads.⁵ But the gauge is in fact perfectly reliable, and Roxanne has consequently compiled a large number of true beliefs and no false beliefs as a result of looking at the gas gauge. On the basis of the gauge's track record, Roxanne proceeds to conclude that the gauge is reliable. And Roxanne's

⁵ Note that there is any number of independent checks that we can implicitly perform on our gas gauges. For example, the fact that the a gas gauge moves down toward E as one drives, and that it goes back up to F when one has filled the tank, is an indicator that the gauge is working correctly. In Vogel's case, Roxanne employs *no* independent checks on the gauge. This is perhaps difficult to imagine, unless Roxanne is a rather odd person. Perhaps we can imagine that it is a family vehicle which she never drives, but which she has a habit of turning on every day to check the fuel level. (Thanks to David Hitchcock for emphasizing how unusual realistic cases of epistemic circularity would be.)

inference is circular, in that she does not appeal to anything other than the gauge itself in order to determine the gauge's reliability.

This case involves a track-record argument: it concludes that some belief-source is reliable, on the basis of its past record of belief-outputs. Moreover, it is also an example of an epistemically circular argument, because the truth of the premises has not been checked by any independent source.⁶

Track-record arguments need not always be epistemically circular. When a track-record argument employs beliefs that have been generated by the belief-source the reliability of which is being argued for as its premises, and the truth of those beliefs has also been checked by an independent belief-source, then that can be a non-circular track-record argument. (Indeed, most track-record arguments, outside of philosophical contexts, are non-circular.) It is only when there is no independent check on the truth of the premises that a track-record argument is epistemically circular. To keep things simple, in what follows I will only be talking about epistemically circular track-record arguments.

2.2. Three reactions to epistemic circularity

There are, I believe, three initially appealing responses to epistemically circular arguments.

(1) My own first natural reaction is to appeal to intuitions about circular arguments generally: when we see circular arguments, and we see that they are in fact circular, they just look wrong. Circularity is generally a bad property for arguments and analyses to have, and we do all agree that there are clear cases where arguments are bad because of their epistemic circularity. So it's just natural to think that all epistemically circular arguments are bad.

⁶ Track-record arguments are not the only kind of epistemically circular argument. See Bergmann (2006, p.180) for an epistemically circular, non-track-record argument. To keep things simple here, I only consider track-record arguments, but the discussion will carry over to other sorts of arguments too.

However, even though intuitions are a good place to start, we need to move on to positive arguments once we find someone who denies the intuition, or who accepts the intuition but gives an argument for thinking that it's mistaken in at least some cases. We need more than intuitions that something is wrong.

(2) A second common reaction to epistemic circularity is to point to the analogy with a person trying to lift herself up by her bootstraps. If you want to lift yourself up, pulling upwards on your bootstraps is not the way to do so. You need to use something that is external to you in order to lift yourself up; using your arms to pull upwards on the very thing that is supporting them will not get you anywhere. Similarly, a belief-source cannot lift itself up using only those beliefs that it has generated. Beliefs are like our arms, in this analogy, and the belief-source is like our bodies. Extending the analogy a bit, we can say that supporting the premise beliefs in an argument for the reliability of their source, with *independent* checks on the truth of those premises, is like placing your arms on an independent support. With independent support, our arms can lift our bodies, and our beliefs can give us a reason to think that the source that generated them is a reliable one.

This analogy is very persuasive. But we might still be unsatisfied with a simple analogy between arguments and bodies. Gravity does not literally apply to arguments or belief-sources, for one thing. For another, given that there are philosophers who are aware of this analogy (the analogy is why we call the circularity problem the *bootstrapping* problem, after all), but who still argue that some epistemically circular arguments are acceptable, we need a more direct reason to think that such arguments are epistemically bad.

(3) The third reaction that I have to epistemically circular arguments is more direct. It involves appealing to the purpose of arguments as aiming to give reasons to accept their

conclusions, or to think that their conclusions are true. Good arguments are those the premises of which do provide good reason for accepting their conclusions. If that is what makes a good argument, then an epistemically circular argument cannot be good. Recall the case of Roxanne: in seeking to determine whether her gas gauge is reliable, she has relied entirely on beliefs produced by that very belief-source; her only reason for thinking that her particular beliefs about the amount of fuel in the tank are true is that the gauge said so.⁷ But when the question at hand is whether some source is reliable, beliefs or testimony generated and verified only by that source cannot *give* anyone a reason to think that that source is reliable – if there is not already a reason to think that it is a reliable source, then the mere fact that it says of itself that it is reliable generates no reason to think so. Unreliable sources can do the same thing, after all, and there are plenty of unreliable sources around. And if there is an *antecedent* reason to think that the source is reliable, then *that* is what is going to be doing the justificatory work in support of the belief in its reliability. The epistemically circular argument will add nothing to the justification that is already there for believing in the reliability of the source in question – it provides no extra epistemic “boost” for the belief in the reliability of the source.⁸

2.3. Refining the argument against epistemic circularity

That third response, I believe, is basically correct. It can be stated more precisely with the help of the following principle:

Subjective Access to Basis-Relative Safety (SABRS): On the basis of data D, belief *p* can be epistemically justified only if it's not unreasonable to think that in most close possible worlds in which the data appear the way they do, the belief is true.

⁷ Of course, Roxanne could check the truth of her gas-tank beliefs by independent means (see above, n.5), and if she did, the reasoning would not be epistemically circular. In the present case, however, she does not do so.

⁸ Jesper Kallestrup suggested this way of putting the point in discussion.

That is to say, the belief p can be justified by data D only if it's not unreasonable to think that p would be basis-relative safe, if held on the basis of D . For example, it is not unreasonable for me to think that in most close possible worlds in which I have the visual, tactile, and memorial data that I do right now, I am in fact typing at my laptop, so those data are capable of justifying my belief that I am now typing at my laptop. By contrast, it is entirely unreasonable for me to think that in most close possible worlds in which I have those same data, I am in fact typing at a desktop computer, so those data are not capable of justifying the belief that I am typing at a desktop.⁹

The SABRS principle, together with the fact that there are many unreliable belief-sources out there, entails that epistemically circular arguments do not yield justified beliefs as their conclusions. Because there are many unreliable belief-sources out there – crystal-ball gazing, tea-leaf reading, and trusting one's gut over the available evidence, for example – and because such belief-sources are capable of generating epistemically circular arguments in support of themselves, it follows that it is unreasonable to think that the belief in the reliability of a belief-source, on the basis of an epistemically circular argument, is basis-relative safe.

Two clarifications of SABRS are necessary. First, SABRS makes use of the notion of what is true in close possible worlds. This might seem troubling, because most people simply have never heard of possible worlds, so surely it will be unreasonable for most people to think anything about what goes on in close possible worlds. If so, then SABRS seems to entail a fairly widespread skepticism about evidential justification: for people who do not know anything about

⁹ At least, laptop-type data are not capable of justifying desktop-type beliefs, unless we add some significant and strange assumptions. Perhaps something like "I often feel like I am typing at a laptop when I am in fact typing at a desktop" would do the trick. I take it that there is a presumption against assumptions like these in normal contexts, so they do not affect what is reasonable to think in normal contexts.

possible worlds, no data can justify their beliefs. That will go for laypeople as well as philosophically untrained scientists. This looks bad.

But we really do not need to be worried here. The notion of a possible world is only a technical way of spelling out what the world could have been like, and that is a notion that everyone is familiar with. We think hypothetically about the world every day: what would happen, if...? We would have happened, had...? The philosopher's notion of a possible world is, at the heart of it, nothing more than a way of spelling that out. We need not take a stand on the metaphysical status of possible worlds in order to accept SABRS; whatever the status of possible worlds turns out to be, it can be reasonable for people to have beliefs about what happens in them. In other words, people do think about possible worlds; it's just that they do so by thinking of how things might be, or how things otherwise might have been.

The second clarification of SABRS is regarding the notion of reasonableness. As it stands, SABRS makes use of an unexplained notion of reasonableness, which might make it seem impossibly vague, or possibly even circular. Taking these two worries in turn, let me say, first, that SABRS is not uselessly vague even if I do not have a positive account of reasonableness in mind. If we want, we can rest with an intuitive understanding of cases where a belief is reasonable and where it is unreasonable, because there will be clear cases where it just is unreasonable to think that a belief is basis-relative safe – think of a case where there is simply a mountain of evidence against a belief p , but instead of believing p , a subject believes not- p , precisely *because* the evidence appears so strongly to support p . Surely such a belief is unreasonable. (Think of extreme conspiracy theorists, for example.) So there is an intuitive understanding of when it is reasonable to think that a belief is true in a range of possible worlds.

But I grant that leaving things at that intuitive level is unsatisfying. The following minimal, sufficient condition for unreasonableness should serve the purpose of specifying the SABRS principle:

Unreasonable: a belief is unreasonable if there is strong, undefeated¹⁰ reason for doubting it.

We can use SABRS, understanding what it is for a belief to be unreasonable in terms of Unreasonable, to yield the following argument against epistemic circularity. Let “A” be an epistemic agent and “B” be the belief that some particular belief-source is reliable.

- P1. If there were many unreliable belief-sources capable of generating self-supporting epistemically circular arguments, that would be a strong reason for doubting that in most close possible worlds in which A believes B on the basis of an epistemically circular argument, B will be true.
- P2. There exist many unreliable belief-sources capable of generating epistemically circular arguments in support of themselves.
- P3. There is a strong reason for doubting that in most close possible worlds in which A believes B on the basis of an epistemically circular argument, B will be true. (P1, P2)
- P4. The strong reason for doubting, mentioned in P3, is not defeated by any external factors.
- P5. The belief that B is true in most close possible worlds in which A believes B on the basis of an epistemically circular argument, is unreasonable. (P3, P4, Unreasonable)
- C. Holding the belief in the reliability of a belief-source, on the basis of an epistemically circular argument, is not justified. (P5, SABRS).

¹⁰ Where a defeater can be either an undercutter or an overrider.

We can use Unreasonable along with SABRS, then, to construct a clear argument against epistemic circularity. This argument is the main line of argument I mean to defend in this paper.

The other worry about SABRS's use of the concept of reasonableness, that SABRS is defining justification in circular terms, can be put to rest easily: SABRS is not part of the analysis of justification. I agree that if it was part of the analysis of justification, that would make the analysis unhelpfully circular, but it just isn't so. SABRS *is* meant as a necessary condition for a belief's being justified on the basis of some data, but that does not make it part of an analysis of justification. To see why not, consider an analogy: it is necessary for a building's being two stories tall, that it not have a third story. This is a necessary condition on what it is to be two stories tall, and it invokes the concept of a storey, but it is not part of any attempt to give an informative and accurate *analysis* of what it is to be two stories tall.

Similarly, a belief can be justified only if it is not unjustified. This is another true necessary condition on justification, which does not enter into an informative, accurate analysis of justification. In the same way, SABRS is a necessary condition on justification; the fact that it makes use of an epistemic notion, the notion of unreasonableness, does not make it a false or a useless principle. That fact only makes SABRS unhelpful for the business of conceptual analysis. But the aim of this paper is not the conceptual analysis of justification; it is to show that a certain kind of argument does not yield justification. I am free, therefore, to make use of whatever necessary conditions for justification as are suitable for the job, and that includes conditions which make use of epistemic terms.

2.4. Objections to the argument against epistemic circularity

So the main line of argument against the legitimacy of epistemically circular arguments proceeds by SABRS, with what it is to believe unreasonably spelled out by Unreasonable. In the next

section, we will look at some reliabilist defenses of certain kinds of epistemically circular arguments. First, however, there are two objections¹¹ to the main line of argument which I want to answer. Seeing why the argument survives these objections should help to clarify exactly how the argument works.

2.4.1. A rival to P1

The first objection targets P1. P1 says that if there are many unreliable belief-sources which are capable of providing epistemically circular self-supporting arguments, then we have a strong reason for doubting that beliefs formed on the basis of epistemically circular arguments in most close possible worlds are true. But surely, the objection goes, the following conditional is at least as plausible as P1:

P1*: If there are many reliable belief-sources which are capable of generating epistemically circular arguments in support of themselves, then we have a strong reason to accept that in most close possible worlds in which you form a belief in the reliability of a belief-source on the basis of such an argument, your belief is true.

Surely, there *are* many reliable belief-sources out there, so we can mount a parallel argument to the SABRS-based argument above, but this time in support of the contrary conclusion that epistemically circular arguments generate justified belief in the reliability of their target belief-sources.

But there is a response to this objection. P1* is not in fact as plausible as P1. For P1* says that if many belief-sources are reliable, and can generate epistemically circular arguments in support of themselves, then most beliefs in the reliability of belief-sources held on the basis of epistemically circular arguments, in close possible worlds, are true. The move from “many” to

¹¹ Which were helpfully suggested by a blind referee.

“most” is what causes the problem. It does not follow, from the claim that *many* reliable belief-sources are capable of generating self-supporting epistemically circular arguments, that *most* beliefs in the reliability of belief-sources held on the basis of such arguments in close possible worlds are true – for there are also many *unreliable* belief-sources out there. The existence of many reliable belief-sources and many unreliable belief-sources, each capable of generating epistemically circular arguments in support of themselves, are quite compatible. But if there are many reliable as well as unreliable sources, then it will likely not be the case that most beliefs in the reliability of a belief-source on the basis of an epistemically circular argument in close possible worlds will be true. The rival conditional to P1, then, turns out to be false.

There is no parallel objection available to P1. For P1 does not make a move from the existence of *many* unreliable belief-sources capable of generating epistemically circular arguments, to the claim that *most* beliefs in the reliability of a belief-source on the basis of an epistemically circular argument are false; it moves only to the claim that we have reason to doubt that most such claims are true. Doubting that most claims of a certain sort are true is not the same as believing that most claims of a certain sort are false.

2.4.2. Perception and hallucination

The second objection to the main argument is that it looks like it might have skeptical consequences. For if the SABRS-based argument works against epistemic circularity, then it looks like it can also work against perception. There exist, after all, many hallucinatory experiences which seem to be real perceptual experiences of the external world. So if we endorse P1, we should also endorse an analogous conditional about perception, something like: if there exist many hallucinatory experiences, then we have a strong reason to doubt that in most close possible worlds in which we form beliefs on the basis of perceptual experience, our beliefs are

true. If accepting P1 commits us to accepting this conditional, then we should accept a SABRS-based argument against the power of perception to justify perceptual beliefs. Skepticism looms.

Fortunately, accepting P1 does not commit us to accepting the conditional in the case of perceptual justification. When the question at issue is whether most belief-sources capable of generating epistemically circular arguments are reliable, the existence of many unreliable belief-sources capable of generating such arguments is sufficient to give us reason to doubt that most sources capable of generating such arguments are reliable. This is because the existence of many unreliable belief-sources capable of generating epistemically circular arguments in support of themselves gives us a good reason to doubt that the *ratio* of reliable to unreliable belief-sources which are capable of generating such arguments is very high. It is the ratio of such sources that counts, when we are considering the relative amounts of close possible worlds in which beliefs held on the basis of circular arguments are true. (P1 is, after all, a claim about what is unreasonable to believe about most beliefs of a certain sort in most close possible worlds, so it is a claim about ratios.)

In the case of perception, however, despite the fact that there do exist many hallucinatory experiences which are indistinguishable from veridical perception, there is still an overwhelmingly high ratio of veridical to hallucinatory perceptual experiences. The existence of many hallucinatory experiences does not, therefore, give us a good reason to doubt that beliefs held on the basis of perception in most close possible worlds are true. Or, if you prefer, the existence of many hallucinatory experiences does give us a reason for doubting that most beliefs based on perceptual experience in close possible worlds are true, but that reason is completely overridden by the fact that there is such an overwhelming amount of veridical perceptual experiences. But there is no analogous overriding in the case of reliable and unreliable belief-

sources, for there is not an overwhelmingly large number of reliable belief-sources out there. There are no doubt many reliable belief-sources, but not so many as to dwarf the number of unreliable ones, as in the case of perception and hallucination.

So P1 is more plausible than P1*, and accepting P1 does not commit us to skepticism about perceptual justification. With SABRS and the argument it generates against epistemic circularity now, I hope, sufficiently clear, let us consider two important defenses of epistemic circularity.

3. Two reliabilist defenses of epistemic circularity

3.1 Alston

Some of the central discussions of epistemic circularity have been given as objections to reliabilism (Vogel 2000, Fumerton 1995). The idea was that epistemic circularity is a bad thing, and that reliabilism is committed to saying that some epistemically circular arguments are legitimate. But some reliabilists, as it happens, hold that some epistemically circular arguments really are epistemically legitimate, on the very grounds that reliabilism entails their legitimacy (Alston 1993, Kornblith 2009). Vogel and Fumerton's modus tollens argument against reliabilism, one might say, is Alston and Kornblith's modus ponens argument for (some) legitimate epistemically circular arguments.

Each of those philosophers accepts that reliabilism is in fact committed to the legitimacy of at least some epistemically circular arguments. I intend to argue that there is no such commitment. The idea behind the claim that reliabilism is committed to the legitimacy of epistemic circularity, following Alston (1993), goes like this: beliefs are justified, on the

reliabilist picture, just in case they have been produced by a reliable process.¹² In a circular track-record argument, the conclusion is that a belief-forming process (or, more generally, a belief-source) is reliable, and the premises are themselves produced by the very process which features in the conclusion. So we have:

Target-process: the belief-forming process the reliability of which we are trying to establish.

Track-record process: the process of forming a belief about the reliability of a target process, on the basis of an epistemically circular track-record argument.

Now if the conclusion of *Track-record-process* is false, and *Target-process* is unreliable, then *Target-process* fails to confer reliabilist justification on its output beliefs. So it follows that the premises employed in *Track-record-process* will be unjustified. In that case, the track-record argument at hand is a bad one. On the other hand, if the conclusion of *Track-record-process* is true, and *Target-process* is in fact reliable, then *Target-process* confers reliabilist justification on its output beliefs, and so the premises employed in *Track-record-process* are justified.

So on the reliabilist picture, there is room for bad epistemically circular track-record arguments: when the *Target-process* is in fact unreliable, using the *Track-record process* in support of its reliability will involve a bad argument, because the premises employed in the *Track-record process* will lack justification. But there also seems to be room for good epistemically circular track-record arguments: when the *Target-process* is in fact reliable, its output beliefs will be justified, and so the premises which are employed in the *Track-record*

¹² There are typically further conditions added onto a reliabilist account of justification, such as the condition that there not be a different reliable process available to the subject, which would have yielded a contradictory belief if it had been used (Goldman (1986), for example, does this). Refinements such as these do not affect the basic point of the reliabilist treatment of epistemic circularity, though, so I will set them aside for now.

argument will have reliabilist justification. Provided that there is a broad enough set of such premises, they can provide good inductive support for a belief that the *Target-process* is reliable.

Of course, as Alston recognizes, it is impossible to *use* circular track-record arguments to establish the reliability of our belief-forming processes, because we cannot tell the difference between a good circular track-record argument and a bad one, without appealing to independent reasons for thinking either that the *Target-process* is reliable, or that the premises of *Track-record process* are true. And if we do that, then we're no longer relying on the circular reasoning to establish the reliability of *Target process*; we're relying on the independent sources of belief. Nevertheless, despite the fact that we cannot use circular track-record arguments to identify which of our belief-sources are reliable, according to Alston, when we *do* give circular track-record arguments in support of the reliability of belief-sources that are in fact reliable, we do attain justification for believing that those sources are reliable.

I think that we need to go further than Alston does. It should be clear that if we cannot tell the difference between good and bad arguments of a certain type, and we therefore cannot use arguments of that type to tell whether their conclusions are true, then we cannot attain justification for believing their conclusions on the basis of such arguments. Arguments are precisely the sorts of things that we give or cite in order to attain a reflective basis for thinking that some proposition is true. When we cannot use an argument for that purpose (i.e. when we cannot tell whether what we've given is an argument that in fact does provide a reason, a reflective basis, for thinking that its conclusion is true), then that is not a good argument.

Here is the place where it will likely be objected, on the reliabilist's behalf, that the principle that the indistinguishability of any difference between good and bad epistemically circular arguments entails failure of transmission of justification from premises to conclusion in

such arguments, is too internalist a principle. Reliabilists do not require that a subject be able to tell that she has justification for her beliefs, even in cases where justifications are reflective. The only question is the reliability of the process that generated the belief.

But the indistinguishability of any difference between good and bad epistemically circular arguments, and the consequent failure of basis-relative safety of beliefs produced on the basis of such arguments, is relevant on the reliabilist picture as well, because the indistinguishability of a difference makes the process of forming beliefs on the basis of such arguments an unreliable one. There are many close possible worlds in which forming conclusions on the basis of such arguments yields a false belief. Because there are many such close worlds, it follows that if we form the belief that a *Target-process* is reliable, by way of employing a *Track-record process*, it will not be reasonable to think that our belief in the reliability of the *Target-process* is basis-relative safe. And because it is not reasonable to take a *Track-record process* to yield basis-relative safe beliefs, it follows that that is not a reliable process: the true-to-false belief output-ratio of *Track-record process* in normal worlds is not sufficiently high to produce reliabilist justification for its output beliefs, even if the *Target-process* is in fact reliable in normal worlds.

So although SABRS is an internalist sort of restriction on when data may justify a belief, such a restriction is necessary in order to ensure that reflective belief-forming processes (such as the process of forming beliefs on the basis of arguments) are reliable. Because arguments can be – indeed, they very often are – given in support of false beliefs, beliefs formed on the basis of arguments must be formed in such a way that it is reasonable to think that forming beliefs on the basis of arguments of the sort in question is likely to yield a true belief. If it is not reasonable for a subject to think that, then that way of forming beliefs is too likely to yield a false output belief,

so arguments of that type will not generate (either reliabilist or other) justification for believing their conclusions.

So circular track-record arguments provide us with a clear case of transmission failure. Even if the premises in a circular track-record argument have reliabilist justification, because they were produced by a reliable process, still it will not be the case that those premises can confer, or transmit, their justification onto the conclusion of the circular track-record argument, because to draw inferences on the basis of an epistemically circular track-record argument is *itself* not a reliable process.

3.2. Kornblith

There is another way to employ reliabilist principles to defend a type of epistemically circular argument. Kornblith (2009) responds to Vogel's easy knowledge objection to reliabilism, by pointing out that there are two ways of filling in the details of Vogel's Roxanne case. On one way of filling them in, reliabilists are not committed to the legitimacy of the track-record argument; on the other, Kornblith argues, they are, but it's not a problem for them.

On the one hand, the details can be filled in so that Roxanne uses epistemically circular arguments all the time, whenever she is confronted with a question of reliability. On this way of fleshing out the case, reliabilists would not say that her use of epistemically circular arguments is legitimate, because there are many unreliable belief-sources out there, and she indiscriminately uses epistemically circular arguments in support of all of them. Her use of epistemically circular arguments is therefore unreliable.

On the other hand, the case can be filled in so that Roxanne only uses epistemically circular arguments in support of reliable belief-sources – she is somehow responsive to the reliability of belief-sources, and never gives a track-record argument in support of an unreliable

one. In that case, her use of epistemically circular arguments is legitimate, on the reliabilist account, since she only uses them in cases where the belief-source is reliable.¹³

Kornblith's defense of reliabilism from epistemic circularity and my argument against the reliabilist defense of epistemic circularity are in fact quite similar. I argued that even if the belief-source in question is reliable, and hence confers reliabilist justification on its output beliefs, an argument that appealed solely to those beliefs in support of the reliability of the belief-source would not confer reliabilist justification on its conclusion, because forming beliefs on the basis of such arguments is itself an unreliable belief-forming process. Kornblith argues in much the same way, except that he makes room for the possibility that an agent might be unconsciously sensitive to the reliability of belief-sources in such a way that she only gives track-record arguments in support of the reliable ones. In that kind of case, Kornblith thinks, her use of epistemically circular arguments is legitimate.

However, even in the type of case that Kornblith defends as legitimate, where the agent is somehow responsive to the reliability of belief-sources, the premises of that kind of argument do not generate justification for the belief in the conclusion *qua* premises in an argument. In that kind of case, the giving of the track-record argument is *part* of a process by which the agent comes to have reliabilist justification for her belief in the reliability of the belief-source, because the *giving* of track-record arguments, which are themselves unreliable, is a reliable process for the agent in question (due to her unconscious sensitivity to truth-indicative features of the case).

So although this kind of agent does achieve reliabilist justification for her belief in the reliability of her target belief source, it is not due to any epistemic virtue on the part of the type of argument that she has employed – for circular track-record arguments are not reliable, even if

¹³ I grant that this is an odd case, but it is Kornblith's, not mine.

some agents can make reliable use of them. (But how can a process be unreliable if someone makes reliable use of it? Compare: Wayne Gretzky could make very good use of a bad hockey stick. That does not make it a good hockey stick.)

This reply to Kornblith's argument – that unreliable processes can be reliably used by some people – rests on the view that belief-forming processes are not individuated at the level of each individual epistemic agent; multiple people can employ the very same belief-forming processes. If that were not so, then it would make no sense to distinguish the reliability of a process from whether a particular agent can make reliable use of the process.

To say that different people can use the same belief-forming process might sound a bit strange, at first blush – perhaps it sounds like saying that multiple people can use the same set of hands. That *would* be a strange thing to say. I can only use my own hands (except in very strange science-fiction or horror movie-type scenarios), just as you can only use yours. But bear in mind that “belief-forming process” refers to a description of a *type* of process, not an individual token process in the brain. And there is nothing strange about saying that many different people can instantiate the very same type of belief-forming process. Indeed, that is a very natural thing to say when, for example, different people knowingly draw Modus Ponens-style inferences: they are employing the *very same* type of belief-forming process.¹⁴

I do not take my argument here to be hostile to Kornblith's ultimate position. He is trying to defend reliabilism from a version of the easy knowledge problem – the problem that reliabilism entails the possibility of attaining knowledge much too easily, by circular reasoning – by showing that in some cases of the sort appealed to in motivating the easy knowledge problem,

¹⁴ At least, it is natural to think that in *normal* cases, different people employ the same process when they draw Modus Ponens inferences. Turri (2010) discusses some cases of causal deviance even in MP inferences. There is interesting work to be done on such cases, but they are not important for now. The point here is just that, absent causal deviance, people often employ the very same belief-forming processes.

like the case of Roxanne as it is typically understood, the agent is not making reliable *use* of the belief-forming process, and so reliabilists need not say that agents in such cases have justification for their beliefs. In other kinds of cases, agents can make reliable use of circular track-record arguments, and for Kornblith, that means that reliabilists should be happy to say that their beliefs are justified. My point, however, is that we can grant that the belief in the reliability of the belief-forming process, in Kornblith's modified type of case, has reliabilist justification, *without* saying that the process of forming beliefs on the basis of circular track-record arguments is a reliable one. And we can avoid saying that, because processes are not individuated at the level of particular epistemic agents – the very same processes can be employed or instantiated by different people. It follows that the fact that some particular agents can make reliable use of circular track-record arguments does not vindicate circular track-record arguments themselves.

4. Conclusion

If the arguments in this paper are correct, then even when the premises of epistemically circular arguments have justification, that justification fails to transmit to their conclusions, because it is not reasonable to take beliefs formed on the basis of epistemically circular arguments to be basis-relative safe.¹⁵ In more straightforwardly reliabilist terms, the process of forming beliefs on the basis of epistemically circular track-record arguments is not reliable. The reliabilist defenses of epistemic circularity do not work.

¹⁵ Kallestrup (2009) and Becker (2012) have offered alternative attempts to block the transmission of reliabilist justification in an epistemically circular argument. I do not have the space to give an adequate discussion of their arguments here, so I will only note that although Kallestrup and Becker offer ways for reliabilists to get out of saying that epistemic circularity is legitimate, I think that there are ways a reliabilist who was inclined to defend the legitimacy of epistemic circularity could resist their arguments. But the defense of that claim will have to wait for another paper.

Acknowledgments

In thinking about the issues discussed in this paper, I have benefitted very much from feedback by David Hitchcock, Carl Ginet, and a blind referee for *Episteme*, as well as by audiences and commentators at the McMaster Philosophy Speakers' Series, the 2013 Oxford Workshop on the Safety Condition for Knowledge and Religious Epistemology, the 2012 Copenhagen Workshop on Epistemic Circularity, and the 2011 conferences of the Canadian Philosophical Association, the Society for Exact Philosophy, and the Ontario Society for the Study of Argumentation.

References

- Alston, William (1993). *The Reliability of Sense Perception*. Ithaca/London: Cornell University Press.
- (1995). How to Think About Reliability. *Philosophical Topics*. vol.23, pp.1-29.
Reprinted in *Epistemology: An Anthology*. Ernest Sosa and Jaegwon Kim, eds. Malden, Mass: Blackwell, 2000. pp.354-371.
- Becker, Kelly (2012). Why Reliabilism Does not Permit Easy Knowledge. *Synthese*. Published online 27 Nov. 2012. DOI 10.1007/s11229-012-0222-8. Forthcoming in print.
- Bergmann, Michael (2006). *Justification Without Awareness: A Defense of Epistemic Externalism*. Oxford: Clarendon Press.
- Bishop, Michael (2010). Why the Generality Problem is Everybody's Problem. *Philosophical Studies*. vol.151, pp.285-298.
- Cohen, Stewart (2002) Basic Knowledge and the Problem of Easy Knowledge. *Philosophy and Phenomenological Research*. vol.65, no.2, pp.309-329.

- (2010). Bootstrapping, Defeasible Reasoning, and *A Priori* Justification. *Philosophical Perspectives*. vol.24: Epistemology. pp.141-159.
- Conee, Earl (2012). The Specificity of the Generality Problem. *Philosophical Studies*. Published online 10 Jan. 2012. DOI 10.1007/s11098-011-9843-x. Forthcoming in print.
- Conee, Earl, and Feldman, Richard (1998). The Generality Problem for Reliabilism. *Philosophical Studies*. vol.89, pp.1-29.
- Comesaña, Juan (2006). A Well-Founded Solution to the Generality Problem. *Philosophical Studies*. vol.129, pp.27-47.
- Fumerton, Richard. (1995). *Metaepistemology and Skepticism*. Lanham, MD: Rowman & Littlefield.
- Goldman, Alvin (1986). *Epistemology and Cognition*. Cambridge, Mass.: Harvard University Press.
- (1988). "Strong and Weak Justification," in J. Tomberlin (ed.), *Philosophical Perspectives*, Volume 13. Atascadero, CA: Ridgeview. Reprinted in A. Goldman, *Liaisons: Philosophy Meets the Cognitive and Social Sciences*, Cambridge, MA: MIT Press (1992).
- (2008). Reliabilism. *The Stanford Encyclopedia of Philosophy*. Spring 2011 Edition. Edward N. Zalta, Ed. URL = <http://plato.stanford.edu/archives/spr2011/entries/reliabilism/>.
- Hofmann, Frank (2012). Three Kinds of Reliabilism. *Philosophical Explorations*. iFirst article. DOI number: [dx.doi.org/10.1080/13869795.2013.738305](https://doi.org/10.1080/13869795.2013.738305)
- Kallestrup, Jesper (2009). Reliabilist Justification: Basic, Easy, and Brute. *Acta Analytica*. vol.24, pp.155-171.

- (2012). Bootstrap and Rollback: Generalizing Epistemic Circularity. *Synthese*. vol.189, no.2, pp.395-413.
- Kornblith, Hilary (2009). A reliabilist solution to the problem of promiscuous bootstrapping. *Analysis*. vol.69, no.2, pp.263-267.
- Pollock, John and Cruz, Joseph (1999). *Contemporary Theories of Knowledge*, 2nd edition. Lanham, MD: Rowman and Littlefield.
- Robinson, Richard (1971). Begging the question, 1971. *Analysis* 31: 4, 113-117.
- Shafer-Landau, R. (2009). A defense of categorical reasons. *Proceedings of the Aristotelian Society*. vol.109 (Part 2), pp.189-206.
- Turri, John (2010). On the Relationship between Propositional and Doxastic Justification. *Philosophy and Phenomenological Research*. vol.80, no.2, pp.312-326.
- Vogel, Jonathan (2000). Reliabilism Levelled. *The Journal of Philosophy*. vol.97, pp.602-623.
- (2008). Epistemic Bootstrapping. *Journal of Philosophy*. vol.105, pp.518-539.
- Weisberg, Jonathan (2010). Bootstrapping in General. *Philosophy and Phenomenological Research*. vol.81, pp.525-548.