

What's the matter with epistemic circularity?

David James Barnett

Published online: 19 December 2013
© Springer Science+Business Media Dordrecht 2013

Abstract If the reliability of a source of testimony is open to question, it seems epistemically illegitimate to verify the source's reliability by appealing to that source's own testimony. Is this because it is illegitimate to trust a questionable source's testimony on any matter whatsoever? Or is there a distinctive problem with appealing to the source's testimony on the matter of that source's own reliability? After distinguishing between two kinds of epistemically illegitimate circularity—bootstrapping and self-verification—I argue for a qualified version of the claim that there is nothing especially illegitimate about using a questionable source to evaluate its own reliability. Instead, it is illegitimate to appeal to a questionable source's testimony on any matter whatsoever, with the matter of the source's own reliability serving only as a special case.

Keywords Epistemic circularity · Bootstrapping · Skepticism · Closure · Uniqueness thesis

1 Two kinds of circularity

Consider Roxanne.¹ Although the fuel gauge in her car is reliable, Roxanne neither knows nor has justification to believe that it is reliable. Aside from the gauge, Roxanne has no source of information concerning the fuel level in the tank—she cannot see into the tank, cannot remember whether she filled it recently, and so on.

¹ The example is originally due to Vogel (2000). See also Alston (1986), Brueckner (forthcoming), Cohen (2002, 2005, 2010), Kornblith (2009), (Vogel 2008), (Weisberg 2010, —2012).

D. J. Barnett (✉)
Department of Philosophy, Union College, 807 Union St, Schenectady, NY 12308, USA
e-mail: davidjamesbarnett@gmail.com
URL: <http://www.davidjamesbar.net>

Nonetheless, when Roxanne sees on Day One that the gauge reads ‘full’, she believes that the tank is full. And since she believes both that the gauge reads ‘full’ and that the tank is full, Roxanne infers that the gauge’s reading is correct. Again on Day Two, Roxanne sees that the gauge reads ‘1/4’, believes that her tank is 1/4 full, and infers that the gauge’s reading is correct. After continuing in this fashion for eight additional days, at the end of Day Ten Roxanne believes that her fuel gauge has a flawless track record, and she concludes on this basis that the gauge is reliable.

There is something the matter with Roxanne’s **bootstrapping** procedure, such that it can confer neither knowledge nor justification to believe that her fuel gauge is reliable. Put roughly, Roxanne’s procedure attempts to vindicate the reliability of her fuel gauge in a way that requires Roxanne already to trust its testimony (or ‘testimony’). In the discussion following Jonathan Vogel (2000) introduction of the example, it has been widely agreed that there is something the matter with such a procedure, although as we will see there is considerable room for disagreement about what that something is.

Roxanne’s bootstrapping procedure is not the only way that one might evaluate a source’s reliability by trusting that source’s own testimony. There is a more direct route as well. Consider Raymond. A visitor to the Island of Knights, Knaves, and Fools,² Raymond wishes to determine whether the source before him is a knight. Raymond knows that if asked a yes or no question, a knight will answer correctly, a knave will answer incorrectly, and a fool will select an answer at random. But Raymond has no evidence concerning whether the source before him is a knight, a knave, or a fool. So Raymond asks the source himself whether he is a knight. When the source responds that he is, Raymond trusts his source’s testimony, and believes that the source is a knight.

There is something the matter with Raymond’s procedure as well. Like Roxanne’s procedure for verifying the reliability of her fuel gauge, Raymond’s procedure involves trusting a source’s testimony in an attempt to evaluate the reliability of that very source. When a source’s testimony speaks directly in favor of its own reliability in the way Raymond’s does, we can say that the source’s testimony is **self-verifying**. By extension, we can call Raymond’s procedure of trusting the self-verifying testimony of his source a **self-verification** procedure.

Both Roxanne’s bootstrapping procedure and Raymond’s self-verification procedure are **epistemically circular**, in the (rough) sense that they attempt to verify the reliability of a given source by trusting the source’s own testimony. There may be other kinds of epistemic circularity as well.³ Despite their further differences, these epistemically circular procedures are both epistemically defective, and can confer

² Examples involving knights, knaves, and fools are loosely adapted from Smullyan (1978). Other examples of self-verifying and self-undermining sources can be found in Bergmann (2004), DeRose (1992), Elga and Egan (2005), Fumerton (1995), and Reid (1983, p. 276).

³ A salient possibility here is **rule circularity**, which occurs when one vindicates the reliability of an inference rule in part by reasoning in accordance with that inference rule. In contrast to bootstrapping and self-verification procedures for vindicating the reliability of a source of information, it is an open question whether rule-circular vindications of inference rules are epistemically defective. See, e.g., Boghossian (2001), Dogramaci (2010), Van Cleve (1979), and (Vogel 2008).

neither knowledge nor justification. But what, more specifically, is the source of the epistemic defect? That is, *what's the matter with epistemic circularity?*

My aim here is to contrast what I call 'reductionist' and 'antireductionist' answers to this question, to defend a qualified version of the reductionist answer, and to sketch some implications for an important skeptical challenge to beliefs about the reliability of one's own cognitive faculties. The plan is as follows. In Sect. 2, I will explain the nature and significance of the contrast between reductionist and antireductionist theories of epistemic circularity. In Sect. 2, I will explain and reinforce a prominent argument that a common form of antireductionism is untenable. In Sects. 4 to 6, I will defend a reductionist account of epistemic circularity. In Sect. 7, I will highlight some pessimistic lessons of our discussion for the prospects of one kind of response to a skeptical challenge to our beliefs. Finally, in Sect. 8 I will sketch an alternative way forward.

2 Reductionism and antireductionism

Although Roxanne's and Raymond's sources are objectively reliable, given only the evidence initially available to Roxanne and to Raymond the reliability of their sources is open to question. When Roxanne and Raymond go on to conclude that their sources are reliable by trusting the testimony of those very sources, few theorists would be willing to license their conclusions as justified.⁴ Nevertheless, many theorists do wish to grant that in some circumstances an agent can be justified in trusting the testimony of a source whose reliability is open to question. Call these theorists **credulists** about testimony (henceforth: credulists).⁵ The credulist claims that Roxanne—depending, perhaps, on the specification of further details of her case—could be justified in trusting her fuel gauge's 'full' reading even in the absence of justification to believe that her gauge is reliable. The credulist furthermore claims that Raymond—depending again on further details of his case—could be justified in trusting the testimony of an unverified source on matters other than that source's own reliability. (Roughly speaking, to **trust** a source's testimony that *p* is to believe that *p* on the basis of the fact that the source says that *p*, and in the absence of other good reasons to believe that *p*.)

It has seemed to many theorists, who we can call **incredulists**, that credulists are wrong to license an agent in trusting the testimony of a source who is, so far as the agent initially can tell, unreliable. For in trusting a source's testimony, the incredulist will say, an agent in some important sense treats the source's testimony as a guide to the truth. If the agent treats the source's testimony as a guide to the truth in this sense, and yet does not believe that the source is a reliable guide to the truth, then there appears to be a conflict between the agent's explicit attitudes about the source's reliability and the belief-forming practices she employs in trusting the

⁴ For one of the few, see Van Cleve (2003). See also Bergmann (2004), which offers a more qualified endorsement of some instances of bootstrapping.

⁵ See Pryor (forthcoming) for a distinct but related use of the term 'credulism'.

source's testimony. The incredulist's guiding idea, which we will clarify and reinforce as we go along, is that:

CREDIBILITY REQUIRES APPARENT RELIABILITY: An agent is justified in trusting a source's testimony only if the agent has prior justification to believe that the source is reliable.⁶

So much for credulism and incredulism about testimony. Turn now to reductionism and antireductionism about epistemic circularity.

The **reductionist** about epistemic circularity—not to be confused with a reductionist about testimony⁷—is a theorist who both accepts the principle that CREDIBILITY REQUIRES APPARENT RELIABILITY and thinks that the illegitimacy of Roxanne's and Raymond's bootstrapping and self-verification procedures can be reduced to their violation of this principle. Here's how. Under the assumption that CREDIBILITY REQUIRES APPARENT RELIABILITY, when Roxanne sees on Day One that the fuel gauge reads 'full', she cannot justifiably believe that the tank is full. When she goes on to conclude that the gauge is reliable on the basis of this and other similarly unjustified premises, her belief in this conclusion is unjustified because she was unjustified in believing the gauge's readings in the first place. And when Raymond hears his source claim to be a knight, he is not justified in believing that the source is a knight because he cannot justifiably believe the source's testimony about *anything* without already having justification to believe that the source is a (reliable) knight. There are subtleties here that will need to be addressed in more detail as we go on. But for now, the reductionist's guiding idea is simply that one cannot be justified in verifying a source's reliability by believing the source's own testimony simply because one cannot be justified in trusting the source's testimony unless one already is justified in believing that the source is reliable.

In contrast, the **antireductionist** denies that the illegitimacy of Roxanne's and Raymond's procedures can be accounted for simply by appeal to CREDIBILITY REQUIRES APPARENT RELIABILITY. One kind of antireductionist accepts CREDIBILITY REQUIRES APPARENT RELIABILITY but denies that it fully explains the illegitimacy of Roxanne and Raymond's epistemically circular procedures. Such an antireductionist can grant that Roxanne and Raymond both are guilty of violating CREDIBILITY REQUIRES APPARENT RELIABILITY, so long as she maintains that there is a further problem with epistemic circularity over and above what the reductionist can account for. Although we will have more to say about this sort of view in what follows, our focus will be a second, more common route to antireductionism that simply denies the principle that CREDIBILITY REQUIRES APPARENT RELIABILITY. A theorist who wishes to deny this principle and yet avoid licensing Roxanne's and Raymond's procedures must say that

⁶ A similar principle is discussed in Vogel (2008). See also Cohen (2002) and Van Cleve (2003) for discussion of a corresponding principle concerning knowledge rather than justified belief. Note that 'prior' refers to epistemic rather than temporal priority. See Pryor (2000, pp. 524–525) for a prominent discussion of the distinction. Without the requirement that one's justification for believing the source is reliable be epistemically prior to one's justification for believing that the source is reliable, CREDIBILITY REQUIRES APPARENT RELIABILITY would be consistent with a view like Zalabardo (2005), which allows the source's testimony itself to provide one's justification to believe that the source is reliable.

⁷ See Adler (2012) for a helpful review of recent work on the epistemology of testimony.

their illegitimacy lies elsewhere.⁸ For this reason, any theorist who wishes to license one in believing the testimony of an unverified source is committed to antireductionism at the pain of licensing epistemic circularity.

The antireductionist owes us a story about what the matter is with epistemic circularity, if it is not (merely) what the reductionist alleges. A number of proposals have appeared in the recent literature.⁹ Rather than considering these proposals piecemeal, we will instead confine our attention to considerations which do not turn on the peculiarities of particular proposals.

A few further points of clarification are in order. First, in contrasting reductionism and antireductionism we have been taking it for granted that there is something the matter with epistemic circularity. For this reason, there is room for a further view that simply accepts the legitimacy of Roxanne's bootstrapping and Raymond's self-verification procedures. This is a *prima facie* unattractive view, although for all we will say here it could be that we are forced to accept such a view in order to avoid implausible skeptical results, as Van Cleve (2003) has argued.

Second, Roxanne's and Raymond's self-verification procedures appear to be epistemically illegitimate in at least two respects. The first is that neither procedure confers *knowledge* that the source in question is reliable. The second is that neither procedure confers *justification* to believe that the source is reliable. Because justification is plausibly a necessary condition for knowledge, it is natural to assume that Roxanne's and Raymond's lack of justification explains their lack of knowledge. But it is possible for this assumption to be questioned. Since justification is not sufficient to turn a true belief into knowledge, it may be that Roxanne's and Raymond's epistemically circular procedures violate some further necessary condition for knowledge (such as a safety condition) in addition to their violating the justification condition. While advocating a reductionist view about the failure of epistemically circular procedures to confer justification, my intention is to leave it as an open question whether a corresponding position is tenable with respect to their failure to confer knowledge. For all I will say here, it might be that Roxanne's and Raymond's failures to obtain knowledge are overdetermined because they violate more than merely the justification condition.

Third, we have so far been taking it for granted that any reductionist about bootstrapping also must be a reductionist about self-verification, and vice versa. This too is questionable, and we will consider in Sect. 6 the possibility of severing reductionism about self-verification from reductionism about bootstrapping.

3 Transmission, closure, and the existence thesis

The paradigmatic (credulist) antireductionist about bootstrapping wishes to license as justified Roxanne's belief that her tank is full on Day One (as well as

⁸ Examples of antireductionist views include those advanced by Roush (2005), Titelbaum (2010), Vogel (2008), and Weisberg (2010). See Weisberg (2012, Sect. 2) for a helpful review.

⁹ See, e.g., Weisberg (2010). No Feedback principle, Kallestrup (forthcoming) GEC, and Titelbaum (2010) and Pryor (MS, Sect. VII) ban on No-Lose Investigations.

corresponding beliefs about the level of fuel on other days), and yet to prohibit her subsequent belief on Day Ten that her gauge is reliable. The antireductionist therefore accepts that

- (1) Roxanne's belief that the tank is full on Day One is justified.

If (1) is accepted, then the unwanted conclusion that Roxanne is justified in believing that the gauge is reliable follows from three additional premises, each of which has been contested by antireductionists:

- (2) If Roxanne's belief that the tank is full on Day One is justified, then her belief that the gauge was correct on Day One is justified.
 (3) If Roxanne's belief that the gauge was correct on Day One is justified, then her belief that the gauge made no errors on Days One through Ten is justified.
 (4) If Roxanne's belief that the gauge made no errors on Days One through Ten is justified, then her belief that the gauge is reliable is justified.

From (1) to (4), it follows that

- (5) Roxanne's belief that the gauge is reliable is justified.

Since the antireductionist wishes to accept (1) and deny (5), she must deny one or more of the premises (2)–(4). And this is difficult to do. Despite the vigorous debate these matters have received in the literature,¹⁰ I think the difficulty of maintaining the antireductionist's positions have yet to be fully appreciated. It will help to review some familiar difficulties confronting a theorist who wishes to deny (2). We then will be in a better position to see why similar difficulties confront those who wish to deny (3) and (4).

Consider Roxanne's position when she learns that the gauge reads 'full' and comes to believe that the tank is full. According to (1), Roxanne is justified in her belief that the tank is full. And she surely is justified in her belief that the gauge reads 'full'. Since the fact that the gauge reads 'full' and the tank is full entails that the gauge's reading is correct, any antireductionist who wishes to accept (1) and yet deny (2) must also deny that

DEDUCTIVE TRANSMISSION: If one justifiably believes that *p*, and if one can tell that *p* entails *q*, then one is justified in inferring from *p* that *q*.

In denying DEDUCTIVE TRANSMISSION, the antireductionist is in good company.¹¹ For DEDUCTIVE TRANSMISSION faces uncontroversial counterexamples in which an agent (arguably but plausibly) is in no position justifiably to believe a premise of an inference unless she already is justified in believing its conclusion. For example, even though the fact that a wall is red entails that it is not white with red light shining upon it, one cannot justifiably infer from the fact that it is red that it is not

¹⁰ See note 8 above.

¹¹ See, e.g., Pryor (MS) and Wright (2004) for recent discussions of transmission failure.

white with red light shining upon it if one's only way of knowing that it is red is insensitive to the difference between a red wall and a white wall illuminated by red light. To be sure, a theorist who wishes to deny (2) owes us a story about why Roxanne's case is a case of transmission failure.¹² But given the challenges to DEDUCTIVE TRANSMISSION from other quarters, there is at least some room here for such a theorist to negotiate.

Although the denial of DEDUCTIVE TRANSMISSION places the proponent of (2) in good company, there is a weaker and less dispensable principle that the proponent of (2) also must deny:

DEDUCTIVE CLOSURE: If one is justified in believing that *p*, and one can tell that *p* entails *q*, then one has justification to believe that *q*.

The difference between these two principles is subtle. DEDUCTIVE TRANSMISSION concerns the *source* of one's justification to believe a conclusion that is entailed by a premise one justifiably believes. It says not only that one always is in a position to justifiably believe the deductive consequences of one's existing justified beliefs, but moreover that one is in a position justifiably to believe these consequences by inferring them from the justified belief that entails them. DEDUCTIVE CLOSURE merely requires one to have some justification or other for believing the deductive consequences of one's other justified beliefs. So far as DEDUCTIVE CLOSURE is concerned, this justification might come from some other source. It is therefore open to the denier of DEDUCTIVE TRANSMISSION to accept DEDUCTIVE CLOSURE. For she might say that whenever one is justified in believing *p* and yet unjustified in inferring from *p* the deductive consequence *q*, one must be justified in believing *q* on other grounds. Indeed, this is an arguable but plausible diagnosis of the uncontroversial cases of transmission failure. For even though one cannot infer that the wall is not white with red light shining upon it from the perceptually justified premise that the wall is red, it is plausible that one must already be justified on other grounds in believing that the wall is not white with red light shining on it to be perceptually justified in believing that it is red in the first place.

DEDUCTIVE CLOSURE, in contrast, merely imposes a kind of coherence constraint on justified belief. To see the difficulty of denying such a constraint, suppose that I justifiably believe that *p* and then consider its deductive consequence *q*. If I can tell that *q* must be true if *p* is true, then it seems I cannot justifiably withhold belief from *q* and instead take a different doxastic attitude like disbelieving *q* or remaining agnostic about *q*. For again, I believe that *p* is true and I can see that this means that *q* must be true as well. Since I cannot justifiably take any doxastic attitude other than belief to *q*, I must be justified in taking the attitude of belief to *q* instead. After all, I have to adopt *some* attitude to *q*, even if it is agnosticism or uncertainty, and if the attitude I adopt is not something other than belief then it will have to be belief.

In sketching this rough motivation for DEDUCTIVE CLOSURE, I have appealed to the principle that if one cannot justifiably adopt an attitude other than belief to a

¹² See Kallestrup ([forthcoming](#)) for a recent attempt.

proposition, then one can justifiably adopt the attitude of belief. We can state this principle more generally as follows:

EXISTENCE: Given one's total evidence, there is at least one justified doxastic attitude that one can take to any proposition.

It will help to compare EXISTENCE to the more familiar UNIQUENESS thesis, which says that there is *exactly* one justified doxastic attitude that an agent can take to any given proposition. The UNIQUENESS thesis stands opposed to the permissivist's claim that epistemic norms are at least somewhat permissive, licensing as justified more than one possible doxastic attitude in at least some cases. The weaker EXISTENCE thesis says only that there is *at least* one justified doxastic attitude, and leaves it open whether there ever is more than one justified attitude. EXISTENCE thus stands opposed only to the *nihilist* claim that sometimes all doxastic attitudes are forbidden—that in some cases an agent is unjustified no matter what attitude she takes to *p*.

EXISTENCE is open to question, and opponents of DEDUCTIVE CLOSURE may wish to deny it.¹³ What seems less appealing is to deny DEDUCTIVE CLOSURE and yet hold on to EXISTENCE. For in any alleged violation of closure, an agent must justifiably believe that *p*, see that *p* entails *q*, and yet fail to be justified in believing that *q*. By EXISTENCE, such an agent must then be justified in withholding belief from *q*, even though he can see that *q* must be true if another thing that he believes is true. It is difficult to see how such an attitude could be justified.

We have just considered the costs of rejecting (2), which include the rejection of DEDUCTIVE CLOSURE and the related EXISTENCE thesis. Before turning to (3), which we will be in a better position to evaluate shortly, we can first turn our attention to (4), the claim that Roxanne is justified in believing that the gauge is reliable if she is justified in believing that the gauge made no errors. A theorist who rejects (4) must grant that Roxanne is justified in believing that her gauge has made no errors on Days One through Ten, and yet deny that she can be justified in believing that her gauge is generally reliable.¹⁴ In doing so, this theorist licenses as justified the deductive steps of Roxanne's bootstrapping procedure, only then to say that her procedure fails at its final inductive step. So the theorist who denies (4) also must deny the general principle that

¹³ Avnur (2012) says that in special cases one can be *committed* to believing the consequences of one's other beliefs without thereby being justified. It is natural to think that if one is committed to believing that *p*, then one is unjustified in adopting an attitude other than belief to *p*—in which case Avnur can be read as denying EXISTENCE. Jim Pryor has told me in conversation that he wishes to deny DEDUCTIVE CLOSURE by denying EXISTENCE.

¹⁴ Vogel (2008) accepts (4) in the case of Roxanne, but denies a corresponding premises for cases involving basic inferential rules in the place of fuel gauge reading. Pryor (MS) does the same for perceptual cases. Weisberg's (2010) account of bootstrapping lends itself most naturally to the denial of (4), although in response to an objection from White he allows for the possibility of denying (2) and with it DEDUCTIVE CLOSURE instead (see pp. 20–21). Kallestrup (forthcoming) denies a premise like (4) which concerns knowledge rather than justification, as does Kornblith (2009) for typical cases of bootstrapping.

INDUCTIVE TRANSMISSION: If one justifiably believes that p , and if one can tell that p is the premise of an inductively strong argument for q , then one is justified in inferring from p that q .

It is a delicate matter whether the uncontroversial cases of deductive transmission failure, in which one plausibly must be antecedently justified in believing a deductive consequence of p in order to be justified in believing that p , can be adapted to generate failures of inductive transmission as well. If these cases can be generalized, then the denier of (4) might find himself in good company in denying the affiliated INDUCTIVE TRANSMISSION principle. But for our purposes it is not necessary to settle these difficult matters. We need only observe that the theorist who wishes to deny (4) must deny not only INDUCTIVE TRANSMISSION but also the weaker principle that

INDUCTIVE CLOSURE: If one justifiably believes that p , and if one can tell that p is the premise of an inductively strong argument for q , then one has justification to believe that q .¹⁵

INDUCTIVE CLOSURE imposes a kind of probabilistic coherence constraint on justified belief. And as with the corresponding logical coherence constraint imposed by DEDUCTIVE CLOSURE, it is difficult to reject the constraint imposed by INDUCTIVE CLOSURE without also rejecting the appealing EXISTENCE thesis. For according to EXISTENCE, Roxanne must be justified in believing that the gauge is reliable if she is not justified in withholding belief and adopting another attitude instead. And if Roxanne justifiably believes that the gauge has made no errors on Days One through Ten, as the theorist who wishes to avoid bootstrapping by denying (4) claims, then it is difficult to see how she could justifiably withhold belief that the gauge is reliable. Indeed, it seems that Roxanne could no more justifiably withhold belief that the gauge is reliable than she could withhold belief from a deductive consequence of her belief that the gauge has made no errors.¹⁶

To reinforce this point, consider what Roxanne's attitude must be to the possibility that her gauge is not merely unreliable but moreover **anti-reliable**. Unlike a merely unreliable fuel gauge, which gives readings at random, an anti-reliable fuel gauge is miscalibrated so as to give systematically incorrect readings. A

¹⁵ Although it does not affect the main thread of our discussion, it is arguable that INDUCTIVE CLOSURE but not DEDUCTIVE CLOSURE must be amended to include a 'no defeaters' clause. INDUCTIVE CLOSURE requires such an amendment because one might have a defeater for the conclusion of an inductive argument that does not defeat the premises. For example, I might have it on good authority that there is a black ball among those in an urn, and thus have a defeater for the conclusion that all the balls in the urn are white. Yet I still might know and be justified in believing that each of a large number of balls drawn so far has been white. In contrast, it is less obvious that DEDUCTIVE CLOSURE requires a 'no defeaters' clause, because one might think that any defeater for the conclusion of a single-premise deductive argument must also defeat the premise. (See Schechter 2013, however, for a dissenting view.) Even so, I think this potential difference between the inductive and deductive closure principles can be safely ignored, since no defeaters are present in the cases that concern us. (I furthermore take the addition of a 'no defeaters' clause to leave unaffected the motivation of DEDUCTIVE CLOSURE by the EXISTENCE thesis).

¹⁶ I take this to be the guiding idea underlying rough and ready objections from White (2006, pp. 546–547) and Cohen (2010, p. 144).

perfectly anti-reliable fuel gauge is one that gives incorrect readings without exception, while a **stunningly** anti-reliable gauge gives an incorrect reading 99.99999 % of the time, making the chance of its giving an correct reading on a given occasion 10 million to 1. Now it is a consequence of DEDUCTIVE CLOSURE that if Roxanne justifiably believes that the tank is full on Day One and that the gauge reads ‘full’, then she must also be justified in believing that the gauge is not perfectly anti-reliable. And as we have seen, this consequence of DEDUCTIVE CLOSURE is reinforced by EXISTENCE, since it seems implausible that Roxanne could justifiably withhold belief from the proposition that the gauge is not perfectly anti-reliable when this proposition follows deductively from others that she believes. For similar reasons, EXISTENCE also reinforces a probabilistic coherence constraint that says that Roxanne must be justified in believing that the gauge is not stunningly reliable either. For it seems equally implausible that Roxanne could withhold belief from the proposition that the gauge is not stunningly anti-reliable when she believes that its reading is correct. (To put this in perspective, suppose Roxanne were to run one trial every day for the rest of her life. In order to have even a 50 % chance of a 99.99999 % anti-reliable gauge giving a single correct reading at some point, Roxanne will have to live for 19,000 years.¹⁷) Again by EXISTENCE, if she cannot justifiably withhold belief, then she can justifiably believe that the gauge is not stunningly anti-reliable.

The lessons we have learned for stunning anti-reliability carry over to mere unreliability across multiple trials. Assuming conservatively that there are five possible readings for her fuel gauge—i.e., ‘empty’, ‘1/4’, ‘1/2’, ‘3/4’, and ‘full’—the odds that a merely unreliable fuel gauge will give correct readings on 10 days in a row are roughly the same as 99.99999 % anti-reliable giving a correct reading on a single day.¹⁸ So Roxanne surely could not be justified in remaining unconvinced of the fuel gauge’s reliability if she believes that it has produced no errors. If we accept EXISTENCE, therefore, we must accept that Roxanne is justified in believing that her gauge is reliable if she is justified in believing that it made no errors. That is, if we accept EXISTENCE, we must also accept (4).

It is time to return to (3), which says that if Roxanne is justified in believing that the gauge’s reading is correct on Day One, then she is justified in believing that the gauge made no errors on Days One through Ten. A credulist who wishes to deny (3) cannot plausibly claim that there is a particular day for which Roxanne is unjustified in believing that the gauge’s reading is correct on that day, for it cannot plausibly be claimed that she is justified on Day One but not, e.g., on Day Six. A more plausible strategy for the antireductionist is to grant that Roxanne is justified in believing, for each particular day, that the gauge was correct on that day, but then to deny that she is justified in believing the conjunction of all these particular beliefs. This strategy may appear *prima facie* to have some promise. For even if we grant a single-premise closure principle like DEDUCTIVE CLOSURE, which says that one must be justified in believing

¹⁷ In 19,000 years, Roxanne will run 6,939,750 trials. The chance of a stunningly unreliable gauge giving incorrect readings on every trial will be 99.999999 %^{6,939,750}, or just under 50 %.

¹⁸ If the chance of a correct reading from an unreliable gauge is 1/5, then the chance of 10 correct readings in a row will be $(1/5)^{10} = 99.99998976$ %.

the consequence of a single premise that one justifiably believes, there are well-known problems with a corresponding multi-premise closure principle that says one must be justified in believing the consequence of a potentially large number of premises each of which one justifiably believes. Put roughly, violations of multi-premise closure principles are possible because the slight risk of each premise being mistaken can accumulate into a substantial risk of their deductive consequence being mistaken. For example, one might be justified in believing of each claim made in a book that that claim is true and yet not be justified in believing the deductive consequence that every claim made in the book is true.¹⁹ Appealing to the failure of multi-premise closure principles might therefore appear a promising strategy for the credulist who wishes to avoid commitment to bootstrapping by denying (3).²⁰

This *prima facie* promise does not withstand scrutiny, however. It takes a large number of trials for the risk of a reliable gauge making an error to accumulate, but as we have seen it does not take long for the likelihood of an unreliable gauge succeeding to get very low. There are many cases in between where probabilistic coherence demands both that Roxanne believe that the gauge made no errors if she believes the particular readings and that she believe that the gauge is reliable if she believes that it has made no errors.²¹

We have seen the difficulty of denying (2)–(4) in an attempt by the credulist to avoid commitment to bootstrapping. The upshot is that we either must say that her belief that the fuel gauge is reliable *is* justified, or that her beliefs in its particular readings are *not* justified. I take the latter option, which favors CREDIBILITY REQUIRES APPARENT RELIABILITY, clearly to be preferable. For CREDIBILITY REQUIRES APPARENT RELIABILITY has considerable intuitive appeal, even setting aside worries about illegitimate bootstrapping. Since an unreliable source is somewhat likely to give incorrect testimony about a given proposition *p*, absent further evidence one intuitively cannot have justification to believe the proposition that even if a given source is unreliable, that source nevertheless is correct about *p*. So when you lack justification to believe the source is reliable, it is difficult to see how you could have justification to believe that it is not an unreliable source that is incorrect about *p*.

4 Incremental bootstrapping and expected reliability

We have just seen the difficulty of the antireductionist's position. The common antireductionist wishes to deny that CREDIBILITY REQUIRES APPARENT RELIABILITY, and

¹⁹ See Christensen (2004) for a discussion with which I am largely sympathetic.

²⁰ A proposal in this vein is considered, but not endorsed, by Weisberg (2010, pp. 6–7).

²¹ Supposing (conservatively!) that a reliable gauge stands a 99.9 % chance of giving a correct reading on a single trial, the chance of a reliable gauge giving 10 correct readings on 10 consecutive trials is $99.9\%^{10} = 99\%$. If Roxanne's believing a gauge's reading means being as confident as one should be for a gauge that is known to be reliable, then probabilistic coherence requires a 99 % degree of confidence that the gauge has made no errors. It is plausible that if one is required to be 99 % confident then one is required to believe, but at any rate it is no more appealing to say that Roxanne is justified in being highly confident that the gauge has made no errors and is reliable than it is to say that she is justified in believing these things.

license as justified Roxanne's trust in the particular readings of her fuel gauge. Yet the antireductionist also wishes to deny that Roxanne is justified in subsequently concluding that her gauge is reliable on the basis of the beliefs she forms by trusting the gauge. As we have just seen, this is a difficult set of positions to maintain.

It is time now to consider the viability of reductionism, the view that Roxanne's and Raymond's procedures are defective simply because both involve believing the testimony of a source in the absence of justification to believe that that source is reliable. In this section, we will consider the feasibility of reductionism about bootstrapping procedures like Roxanne's. After building on our discussion in Sect. 5, we will then turn in Sect. 6 to consider the feasibility of a reductionist account of self-verification procedures like Raymond's.

Because the reductionist denies that Roxanne is justified in believing the individual readings of her gauge, the reductionist is not committed to licensing as justified Roxanne's bootstrapping procedure. We should not take this to mean, however, that the reductionist faces no trouble from bootstrapping. Since the reductionist proposes an explanation of the illegitimacy of Roxanne's procedure, his task is not only to avoid commitment to the falsehood that Roxanne is justified, but moreover to explain why she is unjustified. And here, the reductionist faces serious challenges of his own.

In appealing to CREDIBILITY REQUIRES APPARENT RELIABILITY to handle the case of Roxanne, the reductionist's guiding idea is that Roxanne's bootstrapping arises from a rational tension between her (justified) uncertainty about the gauge's reliability and her trust in its particular readings. Because she treats the gauge as reliable when she believes its individual readings, her uncertainty about its general reliability is unstable. As she continues to trust its particular readings, she cannot, on pain of irrationality, remain uncertain of the gauge's reliability. And thus, a theorist who rejects CREDIBILITY REQUIRES APPARENT RELIABILITY, and says that Roxanne's trust can be justified even when she initially is justifiably uncertain of its reliability, is saddled with the implausible result that merely by trusting its readings she gains justification to believe that the gauge is reliable.

Although I believe that this guiding idea is ultimately correct, as it stands the reductionist account of bootstrapping is incomplete. This is because CREDIBILITY REQUIRES APPARENT RELIABILITY is too weak to handle a class of variant cases of bootstrapping. After explaining why CREDIBILITY REQUIRES APPARENT RELIABILITY is too weak to handle the variant cases, I will propose a stronger principle that fares better.

The principle CREDIBILITY REQUIRES APPARENT RELIABILITY says that an agent is justified in trusting a source's testimony only if the agent has justification to believe that the source is reliable, where trusting a source's testimony is understood to involve coming to believe what a source says because the source says it. The principle therefore is silent about an important class of cases in which an agent's attitudes change in light of a source's testimony, but in which this change falls short of the agent's coming to hold a new belief. In these cases of **incremental bootstrapping**, one can still illegitimately increase one's estimation of a source's reliability by a procedure similar to Roxanne's, but without violating CREDIBILITY

REQUIRES APPARENT RELIABILITY's prohibition on forming new beliefs on the basis of an unverified source's testimony.

For an example of incremental bootstrapping, consider Alice, a visitor to the Island of Knights and Knaves. Alice knows that the source before her is either a (perfectly reliable) knight or a (perfectly anti-reliable) knave, but she has no evidence favoring one possibility over the other. Alice furthermore has no evidence concerning whether it will rain, and she is no more (or less) confident than not that it will rain. Nevertheless, when she asks the source whether it will rain and is told that it will, Alice becomes more confident than not that it will rain, but without going so far as to believe that it will rain. Because Alice knows that her source is either a knight who speaks the truth or a knave who speaks falsely, she becomes more confident than not that her source is a knight, but without going so far as to believe that her source is a knight.

Alice's procedure is illegitimate in the same way that Roxanne's is. But because Alice's confidence in her source's testimony falls short of belief, she does not violate CREDIBILITY REQUIRES APPARENT RELIABILITY, which merely bars her from believing the source's testimony. Since CREDIBILITY REQUIRES APPARENT RELIABILITY is too weak to handle this case of bootstrapping, a stronger principle is needed.

A related kind of incremental bootstrapping that fails to violate CREDIBILITY REQUIRES APPARENT RELIABILITY is found in the case of Charles, which I adapt from Jonathan Weisberg's *Charlie*.²² Charles knows that his barometer is either merely reliable or perfectly reliable, and he has no reason to consider one possibility more likely than the other. (The only difference between a barometer and a fuel gauge is that a barometer, we can suppose, has only two possible readings, 'high' and 'low'. This supposition will simplify our work later on.) When Charles sees on Day One that the barometer reads 'low', he not only believes but moreover becomes psychologically certain that the barometric pressure is low. On this basis, he is certain that that on Day One the barometer reads 'low' and the pressure is low. Charles then repeats this process for ninety-nine additional days. At the end of Day One Hundred, Charles is certain that the barometer has a perfect track record for one hundred days in a row, and he infers that the barometer is not only reliable but perfectly reliable.

Once again, we are confronted with a procedure that is illegitimate in the same way that Roxanne's is. But because Charles knows from the outset that his gauge is reliable, he is no more in violation of CREDIBILITY REQUIRES APPARENT RELIABILITY than Alice is. Once again, a stronger principle is needed.

Although CREDIBILITY REQUIRES APPARENT RELIABILITY is too weak to handle cases of incremental bootstrapping like Alice's and Charlie's, there is a principle strong enough to handle such cases that can plausibly be seen as an extension of the reductionist's guiding idea. Let me explain.

Notice first that although you cannot justifiably trust an unverified source's testimony, this does not mean you should be entirely indifferent to such a source's

²² The example has been adapted slightly from Weisberg (2010). As Weisberg notes, a similar point is discussed in White (2006), who attributes it to Greg Epstein, Matt Kotzen, and Nico Silins.

testimony—allowing it no influence at all on your credences. Indifference would be appropriate if you knew that the source is unreliable, in the (stipulative) sense that it is as likely to be incorrect as correct. But when you are uncertain about a source’s reliability, you cannot justifiably disregard its testimony altogether. Just as belief is unjustified without justification to believe your source is reliable, total indifference is unjustified without justification to believe the source is unreliable.

If both belief and indifference are unjustified for an unverified source’s testimony, what degree of confidence, or credence, would be justified? Consider a simple case, in which you know some source either is perfectly reliable or is just taking random guesses, although you are not certain which it is. When you first ask the source whether p , how confident should you be that the source will return with a correct answer?

Whatever your confidence is that the source is perfectly reliable, you should be equally confident that it is both perfectly reliable and correct about whether p is true. And no matter how confident you are that the source is unreliable, your credence should be evenly split between the possibility that it is both unreliable and correct and the possibility that it is both unreliable and incorrect. In other words, your justified credence that the source is perfectly reliable and correct should equal 100 % of your credence that it is 100 % reliable. And your justified credence that the source is unreliable and correct should equal 50 % of your credence that it is 50 % reliable.

Suppose for example that your credence is evenly split between these two possibilities with respect to your source’s reliability—giving you a credence of 1/2 that the source is 100 % reliable and a credence of 1/2 that it is 50 % reliable. If so, your credence that the source is reliable and correct about p should equal (1/2)(100 %), and your credence that the source is unreliable and correct should equal (1/2)(50 %), for an overall credence of

$$\Pr(\text{Rel} = 100\%)(100\%) + \Pr(\text{Rel} = 50\%)(50\%) = 3/4.$$

This informal probabilistic argument can be generalized to trickier cases, where you are open to more than two hypotheses concerning your source’s reliability. Even in the trickier cases, the hypotheses concerning its reliability will form a partition, and your credence that it will deliver correct testimony about p conditional on any one reliability hypothesis should equal the source’s reliability under that hypothesis. So if the basic style of probabilistic argument is accepted, it is trivial to show that where $\text{Rel} = n\%$ is the hypothesis that the source is $n\%$ reliable, your credence that the source’s testimony will be correct should equal the following:

$$ER = \text{df.} \sum n \Pr(\text{Rel} = n\%)(n\%).$$

This value, which we can join Roger White in calling the source’s **expected reliability (ER)**,²³ represents your justified expectation that the source will give the correct answer concerning p . When you know with certainty how reliable your

²³ The name ‘expected reliability’ was arrived at independently by the author and White (2009), due to its similarity to ‘expected utility’.

source is, its expected reliability will equal its actual reliability. But when you are uncertain of its actual reliability, as in the simple case we are considering, your source's expected reliability will reflect this uncertainty, taking a weighted average of the different possibilities. For example, when your credence is split between $Rel = 100\%$, $Rel = 50\%$, and $Rel = 0\%$, the source's expected reliability will be $1/2$, and its testimony therefore will give you no reason to change your credences.

When our source with expected reliability of $3/4$ goes on to claim that p is true, what credence that p are you justified in adopting on the basis of its testimony? Since initially you are justified in a credence of $3/4$ that the source's testimony will be correct, and since you know now that the source's testimony can be correct only if p , a credence of $3/4$ that p is justified unless the fact that the source claimed that p is itself a reason to revise your credence that the source's testimony is correct—i.e., just in case you have no reason to consider a given source's testimony that p more (or less) likely to be true than that source's testimony that not- p . A credence of $3/4$ is therefore justified so long as the source's testimony is **neutral** in this sense. More formally, where $S(p)$ means that the source says that p , neutrality consists in the source's testimony meeting the following condition²⁴:

$$\text{NEUTRALITY CONDITION : } \Pr(p | S(p)) = \Pr(\neg p | S(\neg p)).$$

In the simple case we have just considered, so long as the NEUTRALITY CONDITION is met your justified credence in the source's subsequent testimony that p equals the source's prior expected reliability of $3/4$. If the informal line of reasoning supporting this result is accepted, its generalization is straightforward. Let $Pr(p)$ denote the epistemic probability of p for you, which we treat as equivalent to your ideally rational or justified credence that p . For example, if you know that a coin either is double-headed or double-tailed, but you do not know which, the epistemic probability for you that it will land heads might be $1/2$. We can take the claim that a source is $n\%$ reliable to mean the source's testimony stands an objective chance of $n\%$ of being true.²⁵

²⁴ Proof: When you conditionalize on the fact that the source said that p , your justified posterior credence that the source's testimony about p is correct will equal your justified prior credence iff

$$(i) \Pr((S(p) \wedge p) \vee (S(\neg p) \wedge \neg p) | S(p)) = \Pr((S(p) \wedge p) \vee (S(\neg p) \wedge \neg p)).$$

In the artificial cases under consideration, if your source does not say that p then it instead says that not- p . So iff (i),

$$(ii) \Pr((S(p) \wedge p) \vee (S(\neg p) \wedge \neg p) | S(p)) = \Pr((S(p) \wedge p) \vee (S(\neg p) \wedge \neg p) | S(\neg p)).$$

Iff(ii),

$$(iii) \Pr((S(p) \wedge p) | S(p)) = \Pr((S(\neg p) \wedge \neg p) | S(\neg p)),$$

or equivalently,

$$(iv) \Pr(p | S(p)) = \Pr(\neg p | S(\neg p)).$$

²⁵ Even in a deterministic world, there surely is an important sense in which a double-headed coin stands a 100% chance of landing heads, while fair coin stands a 50% chance of landing heads. Arguably but plausibly, this can be understood in terms of the robustness of the coin's landing heads in the face of minor variations in the initial conditions of the coin toss.

We will show that for a neutral proposition, the source's testimony justifies credence equal to the source's expected reliability. Start with a case where you know that a given source is $n\%$ reliable. Absent other relevant evidence, your epistemic probability that its testimony about p will be correct will equal $n\%$ by Lewis's Principal Principle.²⁶ The proposition that the source's testimony is correct is equivalent to the proposition that either it says that p and p is true, or it says that not- p and p is false. It follows that for any value of n

$$(6) \quad \Pr((S(p)) \wedge p \vee (S(\neg p) \wedge \neg p) \mid \text{Rel} = n\%) = n\%.$$

In cases where you do not know the objective reliability of your source, the various possibilities for its reliability will form a partition, and so

$$(7) \quad \sum n \Pr(\text{Rel} = n\%) \Pr((S(p) \wedge p) \vee (S(\neg p) \wedge \neg p) \mid \text{Rel} = n\%) \\ = \Pr((S(p) \wedge p) \vee (S(\neg p) \wedge \neg p)).$$

From (6) and (7), it follows that

$$(8) \quad \Pr((S(p) \wedge p) \vee (S(\neg p) \wedge \neg p)) = ER.$$

And from (8) and the NEUTRALITY CONDITION, it follows that²⁷

$$(9) \quad \Pr(p \mid S(p)) = ER.$$

Let's take stock. In Sect. 2 we introduced the principle that CREDIBILITY REQUIRES APPARENT RELIABILITY. Although our discussion is consistent with this principle, we have found that it is too weak to handle cases of incremental bootstrapping. We now are in a position to introduce a stronger principle that handles such cases. The principle is that

CREDIBILITY EQUALS EXPECTED RELIABILITY: An agent is justified in adopting a credence equal to the prior expected reliability of a source's neutral testimony.

CREDIBILITY EQUALS EXPECTED RELIABILITY entails CREDIBILITY REQUIRES APPARENT RELIABILITY, given plausible but controversial assumptions about the relationship between credence and all-out belief. Assuming that belief requires high credence, CREDIBILITY EQUALS EXPECTED RELIABILITY entails that an agent is justified in believing a source only if that source's expected reliability is high. Because a source's expected reliability can be high only if one is justified in having a high credence that the source is reliable, it follows that CREDIBILITY EQUALS EXPECTED RELIABILITY entails CREDIBILITY REQUIRES APPARENT RELIABILITY on the plausible but controversial assumption that high credence is sufficient for belief.

²⁶ Lewis (1980). A similar observation is made by White (2009).

²⁷ See note 24 above.

With this in mind, observe that the bootstrapping problems lodged against belief in a source's testimony can be seen as a special case of a more general problem, which arises from any credence exceeding the source's expected reliability. Look at it this way. The more questions you ask the source, the less likely it becomes that the source's track record will significantly deviate from its overall reliability. So for a source that is known to be either 50 or 100 % reliable, after a sufficiently long series of questions only two outcomes are remotely likely: first, that the source is 100 % reliable and gave correct answers to 100 % of your questions, and second, that it is 50 % reliable and gave correct answers to approximately 50 %. Unless you are inconsistent or arbitrarily more confident of some answers than others, your credence in each answer will equal 100 % of your credence that 100 % of the answers are correct plus 50 % of your credence that 50 % of the answers are correct. In order to have credence greater than 3/4 in each answer, therefore, you must be more confident that 100 % of the answers are correct than that 50 % are, and consequently more confident than you were initially that your source is reliable.²⁸

This observation is the centerpiece of a reductionist account of cases of incremental bootstrapping—cases that involve epistemically illegitimate bootstrapping that does not generate the kind of all out belief in the testimony of one's source that CREDIBILITY REQUIRES APPARENT RELIABILITY prohibits. Recall Alice, who knows that the source before her is either a knight or a knave, but who has no evidence supporting one possibility over the other. When her source claims that it will rain, Alice becomes more confident than not that it will rain without going so far as to believe that it will. As we have seen, CREDIBILITY REQUIRES APPARENT RELIABILITY is too weak to prohibit Alice's moderate confidence that it will rain, even though this moderate confidence is plausibly sufficient to generate illegitimate bootstrapping.

The stronger principle that CREDIBILITY EQUALS EXPECTED RELIABILITY fares better. For someone in Alice's situation, the source's expected reliability is $\Pr(\text{Rel} = 100\%)(100\%) + \Pr(\text{Rel} = 0\%)(0\%) = \Pr(\text{Rel} = 100\%) = 1/2$. So when Alice's source claims that it will rain, CREDIBILITY EQUALS EXPECTED RELIABILITY dictates that she retain a credence of 1/2 that it will rain, rather than become more confident than not that it will rain. So unlike the weaker principle that CREDIBILITY REQUIRES APPARENT RELIABILITY, the stronger CREDIBILITY EQUALS EXPECTED RELIABILITY rules this step in Alice's procedure unjustified.

Turn now to Charles. Charles initially knows that his source either is merely reliable or is perfectly reliable, but he has no evidence favoring one possibility over the other. When Charles nevertheless becomes psychologically certain of the barometer's readings on each of the subsequent one hundred days, he concludes that his barometer has a flawless track record and must be perfectly reliable. Although Charles' procedure plainly is illegitimate, he does not violate CREDIBILITY REQUIRES APPARENT RELIABILITY, which merely bars belief in a source's testimony in the absence of justification to believe that it is reliable.

²⁸ A similar observation is made by White (2009, p. 243).

As before, the stronger CREDIBILITY EQUALS EXPECTED RELIABILITY fares better. Assume for simplicity that a merely reliable gauge is 95 % reliable. If so, the expected reliability of Charles' barometer is $\Pr(\text{Rel} = 100\%)(100\%) + \Pr(\text{Rel} = 95\%)(95\%) = (1/2)(100\%) + (1/2)(95\%) = 97.5\%$. So when the barometer reads 'low', CREDIBILITY EQUALS EXPECTED RELIABILITY dictates that Charles adopt a credence of 97.5 % that the pressure is low. It is plausible that a credence this high is sufficient for belief, and it is correspondingly plausible that Charles is justified in merely believing that the pressure is low on the basis of the barometer's reading. But when Charles not only believes but becomes psychologically certain that the pressure is low, he violates CREDIBILITY EQUALS EXPECTED RELIABILITY. So unlike the weaker principle that CREDIBILITY REQUIRES APPARENT RELIABILITY, the stronger principle that CREDIBILITY EQUALS EXPECTED RELIABILITY rules this step in Charles' bootstrapping procedure unjustified.

5 The neutrality condition: parity and lack of bias

Our discussion of bootstrapping has served to motivate the principle that CREDIBILITY EQUALS EXPECTED RELIABILITY. We will shortly consider the import of this principle for a different kind of epistemically circular procedure—the self-verification procedure employed by Raymond, who believes that his source is reliable on the basis that the source himself claims to be reliable. But before doing so, we must pause to consider some ways in which the credibility of a source's testimony can depart from that source's expected reliability.

Recall that CREDIBILITY EQUALS EXPECTED RELIABILITY allows the credibility of a source's testimony that p to depart from its expected reliability whenever p violates the NEUTRALITY CONDITION—the condition that one has no reason to consider the source's testimony that p more (or less) likely to be true than that source's testimony that not- p . There are two ways in which violation of the NEUTRALITY CONDITION can occur. First, you might simply have reason to place greater confidence in p than in not- p to begin with. A case like this can violate the NEUTRALITY CONDITION because it violates the condition that

$$\text{PARITY CONDITION : } \Pr(p) = \Pr(\neg p).$$

A trickier kind of violation occurs in cases where a source's likelihood of error is unevenly distributed between erroneously saying that p and erroneously saying that not- p . The most straightforward cases of this kind involve sources that exhibit an **individual bias** towards a particular answer, such as a medical test that has a higher tendency towards false positives than towards false negatives. In cases of this kind, the NEUTRALITY CONDITION can be violated because there is a violation of the following:

$$\text{LACK OF BIAS CONDITION : } \Pr(S(p) | p) = \Pr(S(\neg p) | \neg p).$$

It can be proved that the NEUTRALITY CONDITION can be violated only when either the PARITY CONDITION or the LACK OF BIAS CONDITION is violated.²⁹ We will have more to say about cases that violate the LACK OF BIAS CONDITION in Sect. 6. For now, we can ask what credence in p is justified in cases where the LACK OF BIAS CONDITION is satisfied but the PARITY CONDITION is not. By Bayes' Theorem,

$$(10) \quad \Pr(p | S(p)) = \frac{\Pr(S(p) | p) \Pr(p)}{\Pr(S(p) | p) \Pr(p) + \Pr(S(p) | \neg p) \Pr(\neg p)}.$$

From the LACK OF BIAS CONDITION it follows both that³⁰

²⁹ Proof: Assume both the PARITY CONDITION and the LACK OF BIAS CONDITION. Because the possibilities that the source says p and that the source says not- p are exhaustive,

$$(v) \quad \Pr(S(p) | p) = 1 - \Pr(S(\neg p) | p).$$

and

$$(vi) \quad \Pr(S(\neg p) | \neg p) = 1 - \Pr(S(p) | \neg p).$$

It follows from (v), (vi), and the LACK OF BIAS CONDITION that

$$(vii) \quad \Pr(S(p) | \neg p) = \Pr(S(\neg p) | p).$$

From (vii), the PARITY CONDITION, the LACK OF BIAS CONDITION, and the elementary theorems

$$(viii) \quad \Pr(S(p)) = \Pr(S(p) | p) \Pr(p) + \Pr(S(p) | \neg p) \Pr(\neg p)$$

and

$$(ix) \quad \Pr(S(\neg p)) = \Pr(S(\neg p) | p) \Pr(p) + \Pr(S(\neg p) | \neg p) \Pr(\neg p),$$

it follows that

$$(x) \quad \Pr(S(p)) = \Pr(S(\neg p)).$$

And now we are a short step from the NEUTRALITY CONDITION. From the LACK OF BIAS CONDITION and the definition of conditional probability, we have

$$(xi) \quad \frac{\Pr(S(p) \wedge p)}{\Pr(p)} = \frac{\Pr(S(\neg p) \wedge \neg p)}{\Pr(\neg p)}.$$

From the PARITY CONDITION and (v), we have

$$(xii) \quad \Pr(S(p) \wedge p) = \Pr(S(\neg p) \wedge \neg p).$$

Finally, from (x) and (xii), it follows that

$$(xiii) \quad \frac{\Pr(S(p) \wedge p)}{\Pr(S(p))} = \frac{\Pr(S(\neg p) \wedge \neg p)}{\Pr(S(\neg p))}$$

which is equivalent to the NEUTRALITY CONDITION given the definition of conditional probability.

³⁰ Proof: It is an elementary theorem of the probability calculus that

$$(xiv) \quad \Pr(S(p) \wedge p) + \Pr(S(\neg p) \wedge \neg p) = \Pr((S(p) \wedge p) \vee \Pr(S(\neg p) \wedge \neg p)).$$

By algebra, we have:

$$(xv) \quad \frac{\Pr(S(p) \wedge p)}{\Pr(p)} \Pr(p) + \frac{\Pr(S(\neg p) \wedge \neg p)}{\Pr(\neg p)} \Pr(\neg p) = \Pr((S(p) \wedge p) \vee \Pr(S(\neg p) \wedge \neg p)).$$

From the definition of conditional probability, it follows that

$$(11) \quad \Pr(S(p) | p) = \Pr((S(p) \wedge p) \vee (S(\neg p) \wedge \neg p)).$$

and³¹

$$(12) \quad \Pr(S(p) | \neg p) = [1 - \Pr((S(p) \wedge p) \vee (S(\neg p) \wedge \neg p))].$$

And so from (8) and (10)–(12), we have

$$(13) \quad \Pr(p | S(p)) = \frac{(ER) \Pr(p)}{(ER) \Pr(p) + [1 - ER] \Pr(\neg p)}.$$

6 Self-verification and selection bias

It is time now to turn to the self-verification procedure employed by Raymond, the visitor to the Island of Knights, Knaves, and Fools. Much like Roxanne, Raymond comes to believe that the source before him is reliable by believing the source’s own testimony. But Raymond takes a more direct route to this conclusion than Roxanne does. For unlike Roxanne’s fuel gauge, whose testimony is limited to the level of fuel in the tank, the inhabitants of Raymond’s island are able to make claims about their own reliability. When Raymond asks the source before him ‘Are you a knight?’, the source responds that he is indeed a (perfectly reliable) knight, and Raymond believes his source’s claim.

Raymond’s procedure is epistemically illegitimate, and the reductionist has a straightforward explanation of why this is so. According to the reductionist, Raymond cannot justifiably believe a source’s testimony about *anything* without having prior justification to believe that the source is reliable. And so, he cannot

Footnote 30 continued

$$(xvi) \quad \Pr(S(p) | p) \Pr(p) + \Pr(S(\neg p) | \neg p) \Pr(\neg p) = \Pr((S(p) \wedge p) \vee (S(\neg p) \wedge \neg p)).$$

From (xvi) and the LACK OF BIAS CONDITION, it follows that

$$(xvii) \quad \Pr(S(p) | p) \Pr(p) + \Pr(S(p) | p) \Pr(\neg p) = \Pr((S(p) \wedge p) \vee (S(\neg p) \wedge \neg p)),$$

and therefore that

$$(xviii) \quad \Pr(S(p) | p) \Pr(p) + \Pr(S(p) | p)[1 - \Pr(p)] = \Pr((S(p) \wedge p) \vee (S(\neg p) \wedge \neg p)).$$

Simplifying terms then gives us (11).

³¹ Proof: Because the possibilities that the source says p and that the source says not-p are exhaustive,

$$(xix) \quad \Pr(S(\neg p) | \neg p) = [1 - \Pr(S(p) | \neg p)].$$

From (11) and the LACK OF BIAS CONDITION it follows that

$$(xx) \quad [1 - \Pr(S(\neg p) | \neg p)] = [1 - \Pr((S(p) \wedge p) \vee (S(\neg p) \wedge \neg p))].$$

We now have (12) from (xix) and (xx).

believe his source's testimony that he is reliable unless Raymond already has justification to believe that the source is reliable.

Although the reductionist is able to handle a simple case like this, we saw in our discussion of bootstrapping that further trouble arose once we broadened our focus to include cases of incremental bootstrapping. So it is natural now to ask how reductionism about self-verification fares once we introduce the graded conception of credibility, and ask to what degree one is justified in trusting an unverified source's self-verifying testimony. Is self-verifying testimony just as credible as testimony about other matters? It turns out that the answer is often No, but for reasons that do not undermine the reductionist's explanatory aims. Let me explain.

One incidental problem with the credibility of self-verifying testimony is **individual bias**. For example, when evaluating our own abilities, we human beings may tend to look at things through rose-tinted glasses. And this can make us less reliable concerning our own reliability than we are concerning other matters, perhaps including the reliability of other people.³² Here I will set aside individual biases of this kind, and assume that whatever a source's reliability is in general, it is equally reliable with respect to whether it is reliable.

It is tempting to think that once we take individual bias off the table, there can be no further problem with the epistemic credibility of a source's self-verifying testimony, above and beyond its credibility about other matters. For illustration, consider Carol, a visitor to the Island of Knights and Fools. Suppose that Carol knows that the source before her, source A, either is a perfectly reliable knight or a merely unreliable fool, and that Carol has no reason to consider one possibility more likely than the other. Suppose further that Carol knows of another source, source B, who also is known to be a knight or a fool, with neither possibility more likely than the other. Carol's only source of information concerning whether A is a knight is A himself, and her only source concerning B also is A. When Carol asks source A whether source B is a knight, A claims that source B is a knight. And when Carol then goes on to ask source A whether A himself is a knight, A claims that he is a knight.

There is a *prima facie* appealing argument that A's claim that he is a knight must be just as credible as was his prior claim that B is a knight. For Carol knows that if A is a knight, then A will be perfectly reliable concerning both his own reliability and B's. Carol furthermore knows that if A is a fool, then he will stand a 50 % chance of speaking the truth about his own reliability and a 50 % chance of speaking the truth about B's reliability. So Carol knows that no matter how reliable A is in general, he is just as reliable concerning his own reliability as he is concerning B's reliability. From these considerations it is tempting to conclude that A's self-verifying testimony is epistemically just as credible as was his testimony that B is reliable.

³² For an overview of some findings from the social psychology literature along with philosophical reflection on those findings, see Elga (2005).

But this *prima facie* appealing argument is unsound. To see why it must be, first observe that, where b is the proposition that B is a knight and $A(b)$ is that A claims that b ,³³

$$\begin{aligned}\Pr(b | A(b)) &= \frac{\Pr(b | A(b)) \Pr(b)}{\Pr(b | A(b)) \Pr(b) + \Pr(A(b) | \neg b) \Pr(\neg b)} = \frac{75\% (1/2)}{75\% (1/2) + 25\% (1/2)} \\ &= 3/4.\end{aligned}$$

No surprises here. Since Carol has no reason to consider it more likely that B is a knight than a fool (or vice versa), the PARITY CONDITION is satisfied. And since A is no likelier to falsely claim that B is a knight than he is to falsely claim otherwise, the LACK OF BIAS CONDITION also is satisfied. So it comes as no surprise that her credence that B is a knight given A's testimony equals A's prior expected reliability of 3/4.

Consider now the credibility of A's self-verifying testimony that A himself is a knight. Where a is the proposition that A is a knight, by Bayes' Theorem we have:

$$\begin{aligned}\Pr(a | A(a)) &= \frac{\Pr(A(a) | a) \Pr(a)}{\Pr(A(a) | a) \Pr(a) + \Pr(A(a) | \neg a) \Pr(\neg a)} \\ &= \frac{100\% (1/2)}{100\% (1/2) + 50\% (1/2)} = 2/3.\end{aligned}$$

It turns out that in this simple case, A's self-verifying testimony that he is a knight is epistemically less credible than is his testimony that B is a knight—even though it is known that no matter how objectively reliable A is, he is as reliable about his own reliability as he is about B's reliability. This is because even when we take individual bias off the table, there is still another form of bias at work in cases of self-verification: **selection bias**. To illustrate, suppose that Carol continues her inquiries on the Island of Knights and Fools. As she goes on to ask the island's inhabitants to evaluate the reliability of themselves and of others, all of the errors will be found among the fools' testimony. When one of these errors concerns some *other* source's reliability, sometimes it will take the form of erroneous 'other-verifying' testimony, saying of another fool that he is a knight. But sometimes it will be erroneous 'other-undermining' testimony, saying of a knight that he is a fool. In contrast, when a fool's erroneous testimony concerns its own reliability, it must be an erroneous self-verification, saying of himself that he is a knight. So it is true that with individual bias off the table, the overall rate of error will be the same for self-assessments and for other-assessments. But when it comes to self-assessments, all the errors will take the form of erroneous self-verifications, making self-verifying testimony epistemically less credible than other-verifying testimony. (Conversely, self-undermining testimony will be *more* credible than other-undermining testimony.)

³³ Two minor notes are in order. First, I assume that A's and B's reliability are epistemically independent, in the sense that one of the source's being a knight would not on its own amount to evidence concerning whether the other is a knight. Second, when A goes on to claim that A himself is a knight, his expected reliability will change, and so too will the credibility of his prior testimony that B is a knight.

The existence of selection bias is consistent with CREDIBILITY EQUALS EXPECTED RELIABILITY. Because the LACK OF BIAS CONDITION is violated in cases of selection bias, so too is the NEUTRALITY CONDITION. So CREDIBILITY EQUALS EXPECTED RELIABILITY does not imply that one is justified in adopting a credence in a self-verifying source's testimony equal to that source's (prior) expected reliability.

Nevertheless, it might be worried that selection bias spells trouble for the reductionist's explanatory aims. Because the reductionist claims that the illegitimacy of circular procedures like Raymond's and Roxanne's can be explained by the illegitimacy of trusting a questionable source's testimony on any matter whatsoever, it is natural to take the reductionist to be committed to a source's self-verifying testimony always being as credible as its testimony on other matters. I think it is a mistake, however, to view selection bias as a marker of vicious circularity. Selection bias is better seen as a confounding factor, whose strength and direction vary with the types of sources in the population, and with their respective frequencies. More specifically, where you know that a source's reliability either equals r or equals some lower value of r minus x , and where p is an arbitrary proposition such that $\Pr(p) = \Pr(\text{Rel} = r)$, the source's self-verifying testimony that its reliability is r (rather than r minus x) will be epistemically less credible than its testimony that p whenever

$$\Pr(\text{Rel} = r) < \frac{x + 1 - r}{x + 1}.$$

I will leave the proof as a take home exercise. Here we need only note the upshot that a source's self-verifying testimony can be *more* credible than its other testimony if the source is highly credible to begin with. For example, for a source that is known to be either 50 or 100 % reliable, selection bias will become beneficial rather than harmful to the credibility of the source's self-verifying testimony whenever the source's prior expected reliability is greater than 5/6.

7 Epistemic circularity and skepticism

One reason why it is important to understand the phenomenon of epistemic circularity is that it lies at the heart of a traditional skeptical challenge to our beliefs about the reliability of our own cognitive faculties. According to this skeptical challenge, you neither can know nor justifiably believe that your cognitive faculties of perception, memory, and reasoning are generally reliable, since any attempt to verify their reliability inevitably will employ those very faculties. To verify the reliability of perception, for example, you will have to appeal to perceptual evidence. And to verify the reliability of your reasoning faculties, you must employ reasoning. Employing one's own cognitive faculties in order to verify their reliability, the skeptic alleges, is no different in principle from the illegitimately circular procedures used by Roxanne and Raymond, and is thus no better suited to confer knowledge or justification than their procedures are.³⁴

³⁴ Although I do not wish to argue the historical point here, I take this skeptical challenge to be closely related to the traditional problem of the criterion. I thank Andrew Cling for pressing me to make this point explicit.

It might be worried that what we have said so far has severely constrained our options for responding to the skeptic. Concerning sources as varied as the testimony of other people and the readings of fuel gauges, we have endorsed an **incredulism** that says that you must have prior justification to believe that these sources are reliable in order to be justified in trusting what they tell you. It might be worried that there are no grounds on which we could resist embracing a corresponding incredulism about our own cognitive faculties. And if so, the worry goes, skepticism looms.

But our situation is not as bleak as it might at first appear. There are a number of anti-skeptical strategies consistent with everything we have said so far. Some strategies might hold that our justification to believe that our cognitive faculties are reliable does not derive from procedures that involve the use of those faculties. One example is the prominent proposal that such beliefs are in some sense justified by default.³⁵ Another proposal, at least for the special case of perception, says that the best explanation of the coherence of our perceptual experiences is that our perceptual faculties are generally reliable.³⁶ These proposals merit a more thorough examination than I can give them here. I wish to focus instead on strategies that try to make room for procedures that involve the use of one's own cognitive faculties to verify their reliability, and in doing so distinguish these procedures from the uncontroversially illegitimate procedures that we have discussed so far. In Sect. 8 below, I will consider what I see as a promising proposal that turns on a distinction between one's 'internal' cognitive faculties and mere 'external' source's of testimony, affirming incredulism about the latter but not the former. But first I will consider in the present section a proposal that turns on the distinction between bootstrapping procedures like Roxanne's and self-verification procedures like Raymond's, and I will argue that a proposal along these lines is not as promising as it might initially appear.

It is sometimes mistakenly assumed that the only way we might use our cognitive faculties to evaluate their own reliability is to engage in a bootstrapping procedure similar to Roxanne's. If so, it can be difficult to see how our justification to believe that our faculties are reliable could derive from a procedure employing those very faculties. In cases of bootstrapping, one tries to assess a source's reliability based on its testimony on matters otherwise unrelated to its reliability, such as how much gas there is in the tank. Whether there is gas in the tank has nothing directly to do with

³⁵ See, e.g., Wright (2004).

³⁶ For a recent discussion, see Vogel (2005). An IBE response to perceptual skepticism along the lines suggested by Vogel could potentially be generalized to respond to skepticism about external sources of testimony, such as other people or fuel gauges. However, I think this strategy is less promising with respect to our faculties of memory and reasoning. For in the case of memory, one cannot judge that one's apparent memories have been coherent in the past without employing one's memories of their past coherence. And one cannot infer from the coherence of past reasoning-based judgments that reasoning is reliable without employing one's faculty of reasoning, which I take to include the capacity for making inferences to the best explanation. It is not obvious that the IBE strategy's failure to generalize to memory and reasoning undermines its plausibility as a response to perceptual skepticism, for it is not obvious that a response to skepticism ought to be uniform across different kinds of skeptical doubts. I hope to address these issues in greater depth in future work.

whether your gauge is reliable, and so it should be no surprise that without independent evidence about the truth or falsity of this testimony, you will have no reason to think a reliable gauge more likely than an unreliable one to read 'full'. Any procedure for evaluating your gauge's reliability based on its readings, then, is bound to be epistemically worthless—it will treat any reading the gauge might give equivalently, as further confirmation of the gauge's reliability. Given that the procedure in some sense is guaranteed in advance to give a favorable assessment of your source's reliability, there seems to be no value in carrying it out.

In the course of defending a reductionist account of epistemic circularity, however, we have contrasted bootstrapping with another epistemically circular procedure that we have called 'self-verification'. In contrast to bootstrapping, the self-verification procedure employed by Raymond is not guaranteed in advance to lead to a favorable outcome. Consequently, a favorable outcome has the potential to be informative in a way that a favorable outcome in a bootstrapping procedure does not. Indeed, we have seen that in some cases one is justified in increasing one's confidence in the reliability of a source based on that source's self-verifying testimony, whereas one never is justified in increasing one's credence in a source's reliability on the basis of bootstrapping.

It might therefore be wondered whether these differences between bootstrapping and self-verification might be sufficient to underwrite a satisfactory response to the skeptic. A suggestion along these lines has been proposed by DeRose (1992) in his defense of Descartes against the charge of vicious circularity in his attempt to verify the reliability of reason using reason. The charge, which goes back to Descartes' earliest critics, was succinctly expressed by Thomas Reid as follows:

If a man's honesty were called into question, it would be ridiculous to refer to the own man's word, whether he be honest or not. The same absurdity there is in attempting to prove, by any kind of reasoning, probable or demonstrative, that our reasoning is not fallacious, since the very point in question is, whether reasoning may be trusted... Every kind of reasoning for the veracity of our faculties amounts to no more than taking their own testimony for their veracity.³⁷

In defense of Descartes, DeRose replies that

we should remember that the use of a faculty could result in one's coming to learn that the faculty is unreliable... In light of this, if Descartes were right that his faculty of clear and distinct perception is self-verifying (rather than self-undermining), this would by no means be an obviously worthless result. Such an epistemically circular verification may well be of some significant value.³⁸

Unfortunately, I do not think this strategy succeeds as a defense of Descartes or as an independently motivated response to the skeptic. One problem the strategy faces

³⁷ Reid (1983, p. 276), also quoted in Plantinga (2002, p. 242).

³⁸ DeRose (1992).

is that self-verifying testimony, even when it provides some evidence in favor of a source's reliability, often provides evidence that is too weak to justify belief (see, e.g., the example of Carol in Sect. 6). But this strategy faces an even more serious problem in the context of responding to Cartesian skeptical doubts. For Descartes himself was concerned not only with the skeptical hypothesis that his faculty of reason is merely unreliable, but moreover the hypothesis that it is anti-reliable, on account of the faculty's having been designed by a malevolent demon. The result that the faculty is self-verifying *is* worthless in response to this hypothesis, since an erroneous self-verification is precisely what one should expect from an anti-reliable faculty. We surely do have good reasons to believe our cognitive faculties are not anti-reliable, but the mere fact that these faculties are self-verifying is not among them.

For further illustration, consider a simple case in which you know a source is either 100 or 50 or 0 % reliable, and you do not consider one possibility any likelier than the others. Suppose you ask the source whether it is 100 % reliable. If the source says Yes, this will to some extent disconfirm that it is 50 % reliable, and support that it is 100 % reliable, since this is the answer you would expect a reliable source to give. However, a Yes answer also is what you would expect an anti-reliable source to give, so the hypothesis that the source is 0 % reliable will also be supported to the same degree. So, the source's claim that it is reliable will support the more extreme hypotheses about its reliability at the expense of the moderate hypothesis that it is 50 % reliable. Since it will remain just as likely that the source is 0 % reliable as it is that the source is 100 % reliable, its expected reliability will remain intact at 1/2.

The point generalizes. Call a probability distribution over reliability hypotheses **symmetrical** just in case for all x ,

$$\Pr(\text{Rel} = 50\% + x) = \Pr(\text{Rel} = 50\% - x).$$

Note that symmetrical distributions yield an expected reliability of 1/2 for the source, although a source can have an expected reliability of 1/2 even if its probability distribution is not symmetrical. Thus, a source's having an expected reliability of 1/2 is a necessary but not sufficient condition for one's probability distribution over reliability hypotheses to qualify as symmetrical

Call a proposition **favorable evidence** if that proposition, if verified independently of the source, would increase the source's expected reliability—e.g., the proposition that the source committed no errors on a number of recent trials. Finally, say that a source's claim that p gives **symmetrical support** for your reliability hypotheses if for all x ,

$$\begin{aligned} & \Pr(\text{Rel} = 50\% + x | S(p)) - \Pr(\text{Rel} = 50\% + x) \\ & = \Pr(\text{Rel} = 50\% - x | S(p)) - \Pr(\text{Rel} = 50\% - x). \end{aligned}$$

Trivially, if your priors are symmetrical, then a claim that offers symmetrical support will leave you with symmetrical posteriors, and so the source's expected reliability will remain intact at 1/2. When p is independently more likely than not to be true, the source's claiming that p will increase the source's expected reliability

even when p is not favorable evidence. And this can be true when p is favorable evidence as well—for example, when p is the disjunction of some highly probable proposition and the proposition that the source is reliable. But when p is favorable evidence with prior probability of $1/2$, the source's claiming that p will provide symmetrical support in any realistic case.³⁹ The upshot is that if your priors for a source are symmetrical, then almost any favorable claim the source could make about its own reliability will not increase its expected reliability.

8 Internal faculties versus external sources

We encountered above Reid's charge that "every kind of reasoning for the veracity of our faculties amounts to no more than taking their own testimony for their veracity." The charge, aimed at Descartes' procedure for verifying the reliability of his faculty of reason, holds that there is no essential difference between trusting the self-verifying testimony of an external source and employing one's own internal cognitive faculties to verify their reliability. Many anti-skeptical philosophers have followed Reid in thinking that internal faculties and external sources are epistemically on a par in this way, but I think we should not take Reid's charge for granted. Elsewhere I argue for the historical claim that Descartes himself distinguished between the epistemic roles of internal faculties and external sources, and that this distinction was central to his strategy for responding to skepticism.⁴⁰ I also argue elsewhere that we, too, should distinguish between the epistemic roles of internal faculties and external sources, and accept credulism about the former but not the latter.⁴¹ Here I will limit myself to sketching some potential anti-skeptical consequences of doing so.

In the uncontroversial examples of illegitimate circularity that we have discussed, an agent trusts the testimony of an external source, such as the claims of another person or the readings of a fuel gauge. As incredulists about testimony, we have affirmed the principle that **CREDIBILITY REQUIRES APPARENT RELIABILITY**, which says that one is justified in trusting such a source's testimony only if one has prior justification to believe that that source is reliable. We have seen how the incredulist is able to explain the intuitive illegitimacy of Roxanne and her cohort's bootstrapping procedures merely by appealing to the principle that **CREDIBILITY REQUIRES APPARENT RELIABILITY** (and the stronger principle that **CREDIBILITY EQUALS EXPECTED RELIABILITY**). And we have said that this lends support to **reductionism**, the claim that it is illegitimate to verify a source's reliability by trusting the source's

³⁹ For the source's claim that p to increase its expected reliability, p will have to selectively confirm an ad hoc collection of favorable and unfavorable hypotheses about its reliability. For example, suppose you know that the source is either 100 or 80 or 60 or 40 or 20 or 0 % reliable, with none more likely than another. If the source claims that it is either 100 or 80 or 0 % reliable, this will rule out that it is 0 % reliable, nudging its expected reliability up from 50 % to roughly 63 %.

⁴⁰ See Barnett (MS a).

⁴¹ See Barnett (MS b, MS c).

own testimony simply because it is illegitimate to trust an unverified source's testimony in the first place.

We are now in a position to consider a potential anti-skeptical upshot of this sort of reductionism: If we have reasons to accept credulism about an agent's own cognitive faculties, even while we deny credulism about testimony from external sources, then we also will have reasons to accept the potential legitimacy of using one's cognitive faculties to verify their own reliability. For if one can be justified in one's beliefs achieved through the use of one's own cognitive faculties even in the absence of justification to believe that those faculties are reliable, then as reductionists about epistemic circularity we should expect there to be no further problem with using those faculties in order verify their own reliability.

Acknowledgment The author is grateful for helpful comments from Paul Boghossian, Justin Clarke-Doane, Andrew Cling, Jonathan Cottrell, Sinan Dogramaci, Hartry Field, Don Garrett, Jesper Kallestrup, Colin Marshall, Tom Nagel, Jim Pryor, Stephen Schiffer, Sharon Street, Peter Unger, David Velleman, and the participants of the NYU dissertation seminar.

References

- Adler, J. (2012). Epistemological problems of testimony. In E. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Available at plato.stanford.edu/archives/fall2012/entries/testimony-episprobl/.
- Alston, W. P. (1986). Epistemic circularity. *Philosophical Studies*, 47, 1–28.
- Avnur, Y. (2012). Closure reconsidered. *Philosophers' Imprint*, 12(9), 1–16.
- Barnett, D. J. (MS a). Knowledge, luck, and the Cartesian circle. Available at www.davidjamesbar.net.
- Barnett, D. J. (MS b). Is memory merely testimony from one's former self? Available at www.davidjamesbar.net.
- Barnett, D. J. (MS c). The agony of higher-order defeat. Available at www.davidjamesbar.net.
- Bergmann, M. (2004). Epistemic circularity: Malignant and benign. *Philosophy and Phenomenological Research*, 69(3), 709–727.
- Boghossian, P. (2001). How are objective epistemic reasons possible. *Philosophical Studies*, 106(1/2), 1–40.
- Brueckner, A. (forthcoming). Bootstrapping, evidentialist internalism, and rule circularity. *Philosophical Studies*.
- Christensen, D. (2004). *Putting logic in its place: Formal constraints on rational belief*. New York: Oxford University Press.
- Cohen, S. (2002). Basic knowledge and the problem of easy knowledge. *Philosophy and Phenomenological Research*, 65(2), 309–329.
- Cohen, S. (2005). Why basic knowledge is easy knowledge. *Philosophy and Phenomenological Research*, 70(2), 417–430.
- Cohen, S. (2010). Bootstrapping defeasible reasoning, and a priori justification. *Philosophical Perspectives*, 24, 141–159.
- DeRose, K. (1992). Descartes, epistemic principles, epistemic circularity, and scientia. *Pacific Philosophical Quarterly*, 73, 220–238.
- Dogramaci, S. (2010). Knowledge of validity. *Noûs*, 44(3), 403–432.
- Elga, A. (2005). On overrating oneself...and knowing it. *Philosophical Studies*, 123(1–2), 115–124.
- Elga, A., & Egan, A. (2005). I can't believe I'm stupid. *Philosophical Perspectives*, 19(1), 77–93.
- Fumerton, R. (1995). *Metaepistemology and skepticism*. Lanham, MD: Rowman and Littlefield.
- Kallestrup, J. (forthcoming). Bootstrapping and rollback: Generalizing epistemic circularity. *Synthese*.
- Kornblith, H. (2009). A reliabilist solution to the problem of promiscuous bootstrapping. *Analysis*, 69, 263–267.
- Lewis, D. (1980). A subjectivist's guide to objective chance. In R. Jeffrey (Ed.), *Studies in inductive logic and probability* (Vol. 2). Berkeley: University of California Press.

- Plantinga, A. (2002). Reply to Bielby's cohorts. In J. Bielby (Ed.), *Naturalism defeated?*. Ithaca: Cornell University Press.
- Pryor, J. (2000). The skeptic and the dogmatist. *Noûs*, 34, 517–549.
- Pryor, J. (forthcoming). Problems for credulism. In C. Tucker (Ed.), *Seemings and justification: New essays on dogmatism and phenomenal conservatism*. Oxford: Oxford University Press.
- Pryor, J. (MS). Is Moore's argument an example of transmission-failure. Available at <http://www.jimpryor.net/research>.
- Reid, T. (1983). In R. E. Beanblossom & K. Lehrer (Eds.), *Essays on the intellectual powers of man in inquiry and essays*. Indianapolis: Hackett Publishing Company.
- Roush, S. (2005). *Tracking truth: Knowledge, evidence, and science*. Oxford University Press
- Schechter, J. (2013). Rational self-doubt and the failure of closure. *Philosophical Studies*, 163(2), 428–452.
- Smullyan, R. (1978). *What is the name of this book?—The riddle of dracula and other logical puzzles*. Englewood Cliffs, NJ: Prentice Hall International.
- Titelbaum, M. (2010). Tell me you love me: Bootstrapping, externalism, and no-lose epistemology. *Philosophical Studies*, 149(1), 119–134.
- Van Cleve, J. (1979). Foundationalism, epistemic principles, and the Cartesian circle. *Philosophical Review*, 88, 55–91.
- Van Cleve, J. (2003). Is knowledge easy—or impossible? Externalism as the only alternative to skepticism. In S. Luper (Ed.), *The skeptics: Contemporary essays*. Aldershot: Ashgate Publishing.
- Vogel, J. (2000). Reliabilism leveled. *Journal of Philosophy*, xcvi(11), 602–623.
- Vogel, J. (2005). The refutation of skepticism. In M. Steup & E. Sosa (Eds.), *Contemporary debates in epistemology*. Malden, MA: Blackwell Publishing.
- Vogel, J. (2008). Epistemic bootstrapping. *Journal of Philosophy*, 105(9), 518–539.
- Weisberg, J. (2010). Bootstrapping in general. *Philosophy and Phenomenological Research*, 81(3), 525–548.
- Weisberg, J. (2012). The bootstrapping problem. *Philosophy Compass*, 7(9), 597–610.
- White, R. (2006). Problems for dogmatism. *Philosophical Studies*, 131, 525–557.
- White, R. (2009). On treating oneself and others as thermometers. *Episteme*, 6(3), 233–250.
- Wright, C. (2004). Warrant for nothing (and foundations for free)? *Supplement to the Proceedings of the Aristotelian Society*, 78(1), 167–212.
- Zalabardo, J. L. (2005). Externalism, skepticism, and the problem of easy knowledge. *Philosophical Review*, 114(1), 33–61.