

Bootstrapping in General*

Jonathan Weisberg
University of Toronto

I Introduction

The following procedure seems epistemically defective. Suppose I have no reason to think the gas gauge in my car is reliable, and I attempt to establish its reliability as follows. I read the gauge on many occasions, concluding each time that the tank is as the gauge says; when the gauge reads ‘full’, I conclude that the tank is full, similarly for ‘empty’, etc. Eventually I conclude by induction that the gauge is reliable, since it was correct each time.

Even if my beliefs in this chain of reasoning are all true, I have done nothing to establish that the gauge is reliable: I do not know that it is reliable, nor am I justified in believing that it is. Call this sort of defective procedure *bootstrapping*.¹ Our topic here is: what is defective about this sort of reasoning, and what epistemological lessons can we learn from its defectiveness?

Vogel (2000) argues that bootstrapping presents a problem for reliabilist theories of knowledge. According to Vogel, reliabilism says that I can use the reasoning just outlined to come to know that my gauge is reliable. Cohen (2002; 2005), Van Cleve (2003), and others argue that bootstrapping actually poses a more general problem, afflicting any view that allows for *basic knowledge*, i.e. any view that allows one to gain knowledge from a source without prior knowledge that the source is reliable.² Thus bootstrapping is also a problem for foundationalists who allow one to gain knowledge from perception without prior knowledge that perception is reliable.

* I am grateful to David Christensen, Stewart Cohen, Sinan Dogramci, Igor Douven, Christoph Kelp, Sari Kisilevsky, Phil Kremer, Jennifer Nagel, Jim Pryor, Sherri Roush, Adam Sennet, Jonathan Vogel, Roger White, and audience members of the 2009 Rutgers Epistemology Conference for many helpful suggestions and criticism.

¹ This ostensive characterization of bootstrapping is not precise, but will serve for now. A clearer characterization will emerge alongside the diagnosis defended below (§3.3).

² Many authors prefer to put things in terms of justification, allowing for *immediate justification* instead of basic knowledge. These authors allow that some sources can justify beliefs even absent any justification for believing that the source is reliable. (For a recent example of such a view, see (Pryor 2000, 2005).) In the interest of simplicity and continuity with recent discussions of bootstrapping, I will focus on basic knowledge, though much of our discussion will bear on views about immediate justification as well. (See, for example, fns. 5 and 7.)

I will argue that bootstrapping poses an even more general challenge. Versions of the bootstrapping problem can be constructed even on strongly internalist theories of knowledge; even if one must always know that one's source is reliable to gain knowledge from it, bootstrapping is still possible. I will then consider some solutions the internalist might offer for her bootstrapping problem, and defend the one I find most plausible: that bootstrapping involves an abuse of inductive reasoning akin to generalizing from a small or biased sample. Finally, I will argue that this solution is equally available to the reliabilist. The moral will be that the issues raised by bootstrapping are orthogonal to questions about internalism and basic knowledge, having more to do with the nature of good inductive reasoning.

2 Generalizing the Bootstrapping Problem

Bootstrapping is problematic even for strongly internalist views that reject reliabilism and basic knowledge. To see why, let's first see how the problem arises for reliabilism, and consider why internalism is supposed to be immune.

According to reliabilism, a belief is knowledge just in case it is true and was formed by a reliable process, even if one does not know that the process is reliable.³ According to Vogel, reliabilism says that knowledge of reliability is gained in the following case:

The Gas Gauge. The gas gauge in Roxanne's car is reliable, though she has no evidence about its reliability. On one occasion the gauge reads *F*, leading her to believe that the tank is full, which it is. She notes that on this occasion the tank reads *F* and is full. She then repeats this procedure many times on other occasions, eventually coming to believe that the gauge reliably indicates when the tank is full.

On each occasion Roxanne knows that the tank is full, according to the reliabilist, because her belief is true and formed by a reliable process, namely trusting the gauge. Her true belief that the gauge reads *F* is also reliably formed, being based on perception. Thus she knows in each instance that the gauge reads *F*, and that the tank is full. She then uses induction to conclude that the gauge reading *F* reliably indicates that the tank is full. Since induction is also reliable, she knows that the gauge is reliable. Intuitively, however, she cannot gain this knowledge in this way.

³ This is a crude formulation of only one kind of reliabilism. Readers who are concerned that salient details are being overlooked are referred to Vogel's (2000) more careful discussion, especially pp. 602–609 and 611–15.

Why doesn't the internalist face the same problematic result? Because she can place requirements on knowledge that prevent Roxanne from knowing in each instance that the tank is full. The internalist can say that knowledge requires justification, and for a belief based on a source to be justified one must know that the source is reliable.^{4,5} Since Roxanne does not know that her gauge is reliable by hypothesis, she is not justified in believing that the tank is full, and thus does not know it is full.

The relevant difference between the reliabilist and the internalist appears to be the requirement of antecedent knowledge that one's source is reliable. If antecedent knowledge of a source's reliability is not needed to gain knowledge from it, bootstrapping threatens to yield knowledge of reliability in an illicit way. Thus Cohen (2002; 2005) and others argue that basic knowledge is the source of the bootstrapping problem. This makes bootstrapping a more general problem, afflicting even some internalist views, e.g. foundationalist theories that allow one to gain knowledge by perception without antecedently knowing that perception is reliable.^{6,7}

Of course, one can reason bootstrappishly even when one antecedently knows that one's source is reliable. Suppose Roxanne had already known her gauge to be reliable. According to the internalist she could then know on each occasion that the tank is full and that it reads *F*. Couldn't she then conclude by induction that her gauge is reliable, and isn't that a problem for the internalist? Even if it is, the internalist still does not face the problematic result the reliabilist faces, since in this variant, Roxanne does not *gain* knowledge. She merely concludes what she already knows: that her gauge is reliable. Thus reliabilism appears to face a serious problem that internalism does not, since reliabilism is committed to an intuitively repugnant gain in knowledge, while internalism is not.

So the internalist blocks bootstrapping by driving a wedge between cases. In cases where there is no antecedent knowledge of reliability, there is no knowledge in each instance from which to gain knowledge of reliability by induction, so knowledge of reliability cannot be gained. In cases where there is antecedent

⁴ Not all internalists can impose this requirement, of course, as we'll note in a moment.

⁵ If we frame our discussion in terms of immediate justification rather than basic knowledge (see fns. 2 and 7), the relevant difference is the requirement that one antecedently possess justification for believing that one's source is reliable.

⁶ Suppose one can know that *P* on the grounds that it looks as though *P*, whether or not one knows that the way things look is a reliable indicator of the way they are. Then one can come to know that *P* based on visual perception, know that it looked as if *P* on that occasion by introspection, and by repeating this procedure on many occasions come to know that the way things look is a reliable indicator of the way things are.

⁷ See (White 2006) for the use of bootstrapping as an objection to views that allow immediate perceptual justification.

knowledge of reliability, there is knowledge in each instance, but induction on those instances yields a conclusion that is already known, so knowledge of reliability still cannot be gained.

Still, one can threaten to gain knowledge via bootstrapping even when one knows antecedently that one's source is reliable. The trick is to note that, even if we antecedently know that the gauge is reliable, there is still room to improve our epistemic standing with respect to the gauge's reliability. For example, one thing we may not know antecedently but might conclude by bootstrapping is that the gauge is *super*-reliable. That is, we may conclude by bootstrapping that the gauge is not only reliable enough to grant knowledge, but even more reliable than it needs to be. Consider:

The Super-Reliable Gas Gauge. Charlie knows that the gauge in his car is reliable, and it is in fact super-reliable. On one occasion the gauge reads *F*, leading him to believe that the tank is full, which it is. He notes that on this occasion the tank reads *F* and is full. He then repeats this procedure many times on other occasions, coming to believe that the gauge is not only reliable, but super-reliable.

Charlie has antecedent knowledge of reliability, so the internalist cannot appeal to a lack of justification in each instance to block his bootstrapping procedure. Moreover, Charlie concludes something stronger than what he knew antecedently, so the internalist cannot dismiss his case as one where no knowledge is gained. Thus the wedge the internalist drove between cases is not actually sufficient to defuse the threat of bootstrapping. Even in cases where antecedent knowledge of reliability is present, there is room to bootstrap.⁸

There are ways the internalist can try to block Charlie's bootstrapping. She might say that Charlie's reasoning is circular, so it does not justify his conclusion. We will consider such responses in the next section. What I want to note now is just that the internalist has to say *something* to block Charlie's bootstrapping. This is important because, whatever she does say about Charlie, the reliabilist may be able to say the same thing about Roxanne. For example, another response the internalist might try is that Charlie's use of inductive reasoning is flawed. And if his inductive reasoning is flawed, Roxanne's inductive reasoning may be flawed too, letting the reliabilist off the hook. This is roughly the view I will defend below.

⁸ White (2006) discusses a similar point, attributing it to Greg Epstein, Matt Kotzen, and Nico Silins.

3 Internalist Solutions

How might the internalist block Charlie's bootstrapping? I will survey three options and argue in favor of the third.

3.1 Appealing to Circularity

Vogel (2008) offers the following restriction to handle bootstrapping: one cannot justify the belief that a rule is reliable using reasoning in which that same rule is applied. Strictly speaking, Charlie does not violate this restriction, since what he concludes is not that the rule "trust the gauge" is reliable, but that it is super-reliable. So a slight amendment to Vogel's proposal is in order:

No Rule Circularity (NRC) A belief about the reliability of rule R cannot be justified by the application of R . That is, neither the conclusion itself nor any belief which supports it may be justified in virtue of the application of R .⁹

Charlie does violate NRC, so it seems a promising way to handle bootstrapping. The problem is that NRC is both too strong and too weak. There are cases where NRC incorrectly says the subject is not justified, and cases where the subject bootstraps but NRC does not block his justification.

Here is a case where NRC should not block justification but does:

The Times Studies. Eliza knows that the *The New York Times* is reliable. She reads an article on the front page reporting three independent studies of the *Times*'s reliability, all of which found that the *Times* is significantly more reliable than most readers think. She concludes that the *Times* is more reliable than she thought.¹⁰

Eliza is right to draw this conclusion,¹¹ but NRC says she is not justified. Her rule is to trust the *Times*, and she applies that rule in the reasoning that leads her to believe the *Times* is more reliable than she thought.

To see that NRC is too weak, notice that bootstrapping can be used to draw conclusions about things besides reliability, but these applications of bootstrapping

⁹ Vogel uses 'NRC' to label the unamended restriction.

¹⁰ White (2006: 530) offers similar arguments.

¹¹ We can support this claim probabilistically: when Eliza reads the *Times*' report, the probability that the *Times* is more reliable than she previously thought goes up. This follows from the fact that $p(E|H) > p(E|\neg H)$, where E is the proposition that the *Times* makes such a report and H is the proposition that the *Times* is more reliable than Eliza thought. (Thanks to Roger White for suggesting this argument, correcting an erroneous one that appeared in an earlier draft.)

are not blocked by NRC. Here is a case:

The Sunday Times. Matt knows *The New York Times* to be reliable every day of the week, though less reliable on Sundays than other days. Not knowing the day of the week, he reads the paper front to back and comes to believe of each sentence that the *Times* says it and it is true. He concludes that today is not Sunday, since the *Times* is normally not perfectly accurate on Sundays.¹²

Matt's rule is to trust the *Times*, but he does not apply the rule to justify a belief about its reliability. Instead his conclusion is about the day of the week, so NRC does not apply.

There is a separate diagnosis of Matt's case that the internalist might appeal to, consistent with her appeal to NRC in other cases. Matt's beliefs are reminiscent of the preface-paradox: when Matt reads today's paper front to back, he is justified in believing each claim individually, but the conjunction of all those claims might be so improbable that he should not believe it. (In Kyburg's (1970) terminology, Matt suffers from conjunctivitis.) But, the story goes, he must believe that conjunction in order to conclude that today is not Sunday, so his conclusion is not justified.

There are some wrinkles here. We might challenge the assumption that Matt must believe this grand conjunction in order to arrive at his conclusion that today is not Sunday. But if this assumption does no harm, the internalist may help herself to it if it allows her to get the right results. Instead we might question whether the conjunction is too improbable to believe. If we specify the case right, this is not obviously so. Suppose the *Times* is infallible Monday through Saturday, and averages one mistake every other Sunday. The probability of the conjunction is then $13/14$ (approximately 93%),¹³ arguably high enough for justified belief.¹⁴

These considerations do not decisively show that the preface-paradox-style diagnosis cannot handle Matt's case. One might insist that a higher probability is necessary for justification in Matt's case. And to make the conjunction more probable we have to either decrease the prior probability that today is Sunday or increase the reliability of the Sunday paper. If we do the former, Matt may already be justified in believing that today is not Sunday, in which case there may be no harm in his bootstrapping reasoning. If we do the latter, the fact that today's paper is perfectly accurate does less to discriminate between Sunday and

¹² Thanks to Adam Sennet for the example.

¹³ This calculation assumes that each day of the week is equally probable at first. We could tinker with the probabilities so that this won't be so, but the considerations to follow will still apply.

¹⁴ We do not, by talking about probabilities being "high enough for justified belief", commit to the threshold view about probabilities and belief that yields the lottery paradox. We merely suppose that, in some cases, passing a (possibly variable) threshold is necessary for justification.

not-Sunday, so Matt's justification for believing that today is not Sunday may be lost. If the internalist insists that the numbers cannot be balanced so as to escape this dilemma, the debate is likely to stalemate.

But notice, if a preface-paradox-style diagnosis is adequate here, it should also be adequate for the cases that were supposed to be handled by NRC. Charlie's case has the same character, after all. To conclude that his gauge is super-reliable, he must believe that it read accurately in many instances, more instances than it would if it were merely reliable. Given that Charlie does not know antecedently whether the gauge is reliable or super-reliable, the number of cases needed to support super-reliability will be such that a conjunction across all of them is not probable enough for acceptance. At least, if the probabilities balance out this way in Matt's case, they should balance the same in Charlie's case. So if the preface-style treatment is adequate for Matt's case, it should be adequate for Charlie's. Why treat them differently then, shouldn't a unified diagnosis be preferred?

We have seen that NRC blocks justification where intuitively it should not (Eliza's case), and that there are cases of a bootstrappish character that NRC does not cover (Matt's case). For the cases not covered, the preface-style diagnosis was offered as a supplement. But given that NRC is too strong anyway, and that the supplement offered was not clearly adequate, we have good reason to abandon the circularity approach. Given also that the preface-style diagnosis would, if adequate, make NRC otiose, the reasons are even weightier.

3.2 Appealing to Conjunctivitis

Why not just appeal to the preface-style diagnosis alone then? Because there are bootstrapping cases that do not share the preface-paradox-like quality of our previous cases, which suggests that we have yet to get to the heart of the bootstrapping problem. Consider:

Slight Bootstrapping. Starla knows that the *Times* is reliable. She opens today's paper, reads the first sentence, P , and comes to believe that the *Times* says P and P is true. She then ever so slightly increases her epistemic probability¹⁵ that the next sentence, Q , is true.

¹⁵ I intend 'increases her epistemic probability' to be neutral between competing interpretations of 'epistemic probability'. For example, we might understand it to mean that she increases her degree of belief, or that she revises her outright belief about what her evidential probability is.

The preface-style diagnosis will not apply to Starla, since she uses only two premises in her reasoning: that the *Times* says P and that P is true.¹⁶ Her premises are weaker than before, so the gain in her epistemic position is weaker too: only a slight increase in her estimate of an epistemic probability. Nevertheless, she is not entitled to even this slight gain.

The beginning idea behind Starla's case is that bootstrapping can be used to make illicit gains on epistemic fronts besides justified belief and knowledge, like epistemic probability. So if the internalist handles Matt's case by appealing to its preface-paradox-like character, we can avoid preface-type complications by setting all-or-nothing belief aside and focusing directly on the probabilities. The driving idea then is that the interesting error in Matt's case is not his disregard for the significant probability of error that accrues as his sample grows large. Rather, it his tacit presumption that the elements of a sample gathered in his way could support his conclusion at all. Starla's case brings this point out by looking at a single-element sample, and asking whether it provides even minimal support for her conclusion. Since the preface-style diagnosis does not apply to Starla, it seems not to get at the heart of our problem. I'll now turn to the diagnosis that I think does.

3.3 Appealing to Feedback

All our examples of bootstrapping have two distinctive features which together suggest a diagnosis. First, they are all instances of reasoning via lemmas: preliminary conclusions are drawn from premises, and those preliminary conclusions are then used as lemmas from which an ultimate conclusion is drawn.¹⁷ In the case of the super-reliable gas gauge, for example, Charlie uses the fact that the gauge reads F in one instance to infer that the tank is full in that instance, and he collects many such instances to further infer that the gauge is super-reliable. Second, they are all cases where the premises, lemmas, and conclusions exhibit a distinctive probabilistic pattern: the premises alone do not increase the probability of the conclusion, but with the help of the lemmas they do. The fact that the gauge reads F in a series of instances does not make it any more likely that the gauge is

¹⁶ NRC will not apply either, since the epistemic boost Starla gets from her "check" does not bear on a belief about general reliability. It bears instead on the epistemic probability of a single instance.

¹⁷ We will have to interpret this picture somewhat liberally if we want to include cases where the "lemmas" are derived from non-doxastic sources, e.g. from perception. I intend my diagnosis to apply to such cases, but I will interpret all our cases as ones where the premise-lemma-conclusion format applies. This allows us to handle all our cases within a traditional probabilistic framework, dodging the thorny problems that arise when considering perception's influence on epistemic probabilities. See (Jeffrey 1965; Field 1978; Garber 1980; Christensen 1992; Weisberg 2009) for some background on the complications that arise here.

super-reliable. But combined with the lemma that the tank was actually full in those instances, the probability of super-reliability goes up.

These observations suggest that bootstrapping is illegitimate because it violates a probabilistic restriction on the use of lemmas. So let's conjecture that the lemmistic reasoning in a bootstrapping case is defeated because the premises of the bootstrapping argument do not by themselves provide the probabilistic support the conclusion requires. The lemmas together with the premises may make the conclusion sufficiently probable, but the lemmas came from the premises, so they may not be used to amplify the probabilistic support of the premises.

To capture this thought, I'll postulate the following defeater for inductive reasoning:

No Feedback If (i) $L_1 - L_n$ are inferred from $P_1 - P_m$, and (ii) C is inferred from $L_1 - L_n$ (and possibly some of $P_1 - P_m$) by an argument whose justificatory power depends on making C at least x probable, and (iii) $P_1 - P_m$ do not make C at least x probable without the help of $L_1 - L_n$, then the argument for C is defeated.¹⁸

The idea is that the amplification of an already amplified signal distorts the original signal, resulting in feedback, and bootstrapping is just “epistemic feedback”. Bootstrapping is an undesirable result of amplifying the output of ampliative inference without restriction. No Feedback prevents such distortion by blocking the feedback loop when distortion will result. No Feedback restricts when the outputs of ampliative inference may be fed back in to ampliative inference: if the second amplification would yield a conclusion not sufficiently probabilistically supported by the premises of the first amplification, the second amplification is defeated.

No Feedback clearly applies to all our bootstrapping cases. The fact that Roxanne and Charlie's gauges say F does not make it any more probable that their gauges are reliable. The fact that Matt and Starla's newspapers say P does not make it any more probable that it is not Sunday, or that the next claim in the paper is true. Only together with the lemmas drawn from these premises — that the tank is full, that P is true — do the conclusions of the respective bootstrapping arguments become more probable. So, if justifying these conclusions requires making them

¹⁸ An importantly different version of No Feedback appeared in a previous draft delivered at the 2009 Rutgers Epistemology Conference. Clause (ii) of the earlier version read “(and possibly other premises)” rather than “(and possibly some of $P_1 - P_m$)”. Douven & Kelp ([forthcoming](#)) point out that the earlier formulation is too strong; it fails to specify that the “other possible premises” are only to include P_i 's. As a result, it defeats perfectly fine reasoning like, “Antoinette said she's going to the party tonight (P), so she'll go (L). And since she usually dances at parties (other premise), she'll dance tonight (C).”

sufficiently probable, No Feedback prevents them from being justified. The premises alone do not make the conclusions sufficiently probable. The premises and the lemmas together would make these conclusions sufficiently probable, but this use of lemmas is precisely what No Feedback forbids.

No Feedback also gives the right verdict in the case that showed NRC was too strong. The fact that Eliza's *Times* reports those studies of its own reliability does, by itself, increase the probability that the *Times* is more reliable than she thought (see fn. 11).

4 No Feedback: A Closer Look

I have been arguing that the internalist should appeal to No Feedback to handle cases like Charlie's. Before considering whether the reliabilist can use the same diagnosis to handle cases like Roxanne's, I want to take a closer look at No Feedback. First I want to clarify its content and explain certain aspects of its formulation. Second I will anticipate some objections. Third I will consider the possibility of extending No Feedback to cover deductive reasoning.

4.1 Clarifying No Feedback

No Feedback's formulation may puzzle some readers; I will briefly consider four points of clarification.

1. *Cumulative and Plain Transitivity.* Why does clause (ii) allow, but not assume, that some of $P_1 - P_m$ will be used in the argument for C ? All our bootstrapping cases are ones where $P_1 - P_m$ are used in the argument for C , and their use appears to be essential. Charlie, for example, uses both the fact that the tank was full (L_1) and that the gauge read 'full' (P_1) in his argument that the gauge is super-reliable, and he could not do without the latter. So why allow that $P_1 - P_m$ might not be used in the argument for C ?

Because it makes No Feedback more general, allowing it to handle cases of a slightly different logical structure. Take a standard preface case, where a historian is writing a grade school textbook, and suppose he reasons as follows: "My first piece of evidence says that the first claim in my book is true, so it is true; my second piece of evidence says that the second claim in my book is true, so it is true; etc. So all the claims in my book are true."¹⁹ So I am an unusually reliable

¹⁹ One might object that the author's reasoning already goes off the rails at this deductive step, so a diagnosis that rejects the next step by appeal to No Feedback misses the mark. We can block this step by extending No Feedback to cover "inductive-deductive" chains of reasoning, a possibility considered below (§4.3).

researcher, since I've written the first error-free textbook in history."²⁰ This reasoning might be construed as having the following structure. P_1 therefore L_1 ; P_2 therefore L_2 ; etc. L_1, \dots, L_n , therefore $L_1 \wedge \dots \wedge L_n$. Finally, $L_1 \wedge \dots \wedge L_n$, therefore C . And on that construal, No Feedback would not apply if the use of $P_1 - P_m$ were required in the argument for C , since they are not used in this case.

Put in terms of principles of transitivity, the point is that our bootstrapping cases all presuppose that inductive reasoning obeys *cumulative* transitivity, while this case presupposes that it obeys *plain* transitivity. Since either presupposition yields feedback, No Feedback is formulated to curtail both.

2. *Evidentialism*. It's helpful to think of No Feedback as saying that one cannot amplify the power of evidence by taking the conclusions one draws from one's evidence and adding them to one's initial stock of evidence.²¹ But No Feedback does not presuppose an evidentialist picture. All it requires is a distinction between what is inferred and what it is inferred from.²² The core insight we are trying to capture is just that if the outputs of ampliative inference are always allowed to serve as further inputs to the ampliative process, feedback results, with bootstrapping emerging as a symptom. There is no need to divide up the inputs and outputs into evidence and non-evidence.

3. *Inferences vs. Bases*. Because No Feedback is framed in terms of inferences, rather than the bases of beliefs more generally, it may not cover all the cases some might want it to. To illustrate, suppose Charlie initially infers his beliefs about the states of the tank from readings of a dipstick, but later forgets those readings and comes to base those beliefs on the readings of the gauge instead. Many will feel that his belief that the gauge is super-reliable then acquires a bootstrappish basis and is no longer justified. No Feedback could be extended to cover this case, but we would have to reformulate it in terms of the basing relation rather than inferences. In the interest of simplicity, and to sidestep controversial questions about the basing relation, I have opted for the more conservative version here.

4. *Measuring Probabilistic Support*. No Feedback presupposes a simplistic view of probabilistic support, but it can be reformulated using more sophisticated views. The idea behind No Feedback is that $P_1 - P_m$ must provide whatever

²⁰ Pollock (1986; 1995) seems to be committed to this reasoning. See Christensen (2004) for some other "downstream" consequences of believing that one's book is error-free.

²¹ Compare Williamson's (1997; 2000) view that all knowledge is evidence, or Levi's (1980) view that we ought to conditionalize on our full beliefs. Slight bootstrapping cases like Starla's present a problem for such views (without even appealing to deductive closure principles).

²² No Feedback can even do work within holistic theories of justification, like Bonjour's coherentism (Bonjour 1985). On such views, a belief's justification comes from membership in a justified corpus, and a corpus's justification depends (at least in part) on the inferential connections within it. If those inferential connections are representable as step-by-step arguments, then No Feedback can constrain what inferential connections there are, and thus what corpuses are justified.

probabilistic support for C the argument from $L_1 - L_n$ (and possibly some of $P_1 - P_m$) requires in order to propositionally justify C . No Feedback takes this probabilistic support to amount to putting C 's probability above some threshold, x . But perhaps, in some cases, the kind of probabilistic support $L_1 - L_n$ (and possibly some of $P_1 - P_m$) must lend C will be of a different kind. For example, it might not matter whether C attains a certain probability threshold, so long as C 's likelihood ratio is sufficiently high; that is, provided

$$\frac{p(L_1, \dots, L_n, P_1, \dots, P_m | C)}{p(L_1, \dots, L_n, P_1, \dots, P_m | \neg C)}$$

is sufficiently high (assuming, for the moment, that all of $P_1 - P_m$ are being used in the argument for C). Or, the argument for C might require that both the likelihood ratio and the final probability of C be above certain thresholds. Or some measure of support besides the likelihood ratio might be appropriate.²³

The current formulation provides a simple and concrete statement capable of handling all the cases on the table. But its simplicity is bought at the price of ignoring the subtleties and lessons of the literature on measures of probabilistic support. This is especially important since Douven ([manuscript](#)) shows that all popular measures of probabilistic support are intransitive. Thus there may be cases the current formulation does not handle, but which would be handled by a more sophisticated formulation that does not presuppose a simple, single conception of probabilistic support. So the current formulation should not be taken as the final word in this respect. Instead it should be seen as a proof of concept, and a template for more refined proposals in the same vein.

4.2 Objections to No Feedback

Here I will briefly consider three objections to No Feedback.

1. *No Feedback is ad hoc.* Even supposing No Feedback gets all of our cases right, does it have any independent plausibility? Or is it just an ad hoc postulate, cooked up to solve our problems here?

Reply. No Feedback does have independent plausibility. Traditional patterns of inductive reasoning are widely acknowledged to have multiple defeaters, and those defeaters are sensitive to probabilistic factors. Take simple induction to a generalization: we observe a sample of F s that are all G s and conclude that all F s are G s. This inference pattern can be defeated if the sample is not large enough, if it is not diverse enough, or if the sample is biased. These are all factors that rob

²³ See ([Fitelson 1999](#); [Christensen 1999](#)) for some background on various ways of measuring probabilistic support.

the sample of its ability to make the generalization probable. So there is precedent for defeaters that prevent inductive inference patterns from drawing improbable conclusions.

Moreover, we can expect there to be defeaters along the lines of No Feedback. If we are going to chain together inductive inferences, as we seem to do (e.g., *Times says peace talks broke down* → *peace talks broke down* → *there will be renewed violence*), then there are bound to be cases where the end point is not made probable by the starting point. This follows from the fact that probabilistic support is intransitive.²⁴ So, if defeaters are generally in the business of keeping step-by-step inductive reasoning on track with probability, and probability is intransitive, there should be defeaters for the general practice of inductive reasoning via lemmas, since the use of lemmas effectively assumes the transitivity of probabilistic support.

2. *No Feedback places too much emphasis on probability.* No Feedback might appear to grant probability sovereignty over justification. If the probability of the conclusion C isn't made sufficiently high by the premises P_1 – P_m , then the argument is defeated. Doesn't that make probability the sole arbiter of justification? Come to think of it, what is the point of step-by-step inductive reasoning if the probability of C given P_1 – P_m is all that matters in the end?

Reply. First, No Feedback does not entail that probability is the sole arbiter of justification. Clause (i) of No Feedback restricts it to cases where probabilistic support is *necessary* for justification. We thus leave open whether other factors might contribute to justification, and whether probability is sometimes unnecessary for justification.

Second, even if probability were the sole arbiter of justification, step-by-step inductive reasoning might still be indispensable. It could serve to discover the hypotheses that we evaluate the probabilities of, it could be a way of evaluating probabilities that are not epistemically transparent, it could be a shortcut heuristic to be used when probabilistic reasoning is too difficult or time-consuming, etc.²⁵

3. *No Feedback is too strong.* Won't we need to violate No Feedback at least sometimes, for how else could we ever come to know elementary facts about our own reliability? If No Feedback is right, then we could never use the deliverances

²⁴ More precisely, the relation *A increases the probability of B* is intransitive, as is the relation *A makes B at least x probable*. These are well-known truisms about the probability calculus, and similar results can be proven for a wide range of probabilistic notions of confirmation/support (Douven manuscript).

²⁵ It may help to compare the role of lemmas in deductive reasoning. Even if logical entailment were the sole arbiter of justification for deductive reasoning, we would still need step-by-step deductive argumentation. We cannot always recognize directly that P entails C , sometimes we have to go via lemma L . The difference between the inductive and deductive cases is just that deductive entailment is transitive, whereas probabilistic support is not. So step-by-step deductive reasoning needs no analogue of No Feedback.

of our faculties to establish facts about their reliability. How then do we come to know how reliable our vision, introspection, memory, etc. are?

Reply. No Feedback does not endanger such knowledge; it does not entail, for example, that we cannot use the deliverances of our senses to evaluate their reliability. As long as the facts about what our various senses are saying do, by themselves, make it probable that our senses are reliable, No Feedback will not prevent us from finding out. We are still allowed to reason in ways like, “perceptual faculties 1–4 testify that P , so P . But faculty 5 testifies that $\neg P$, so faculty 5 is not functioning accurately.” No Feedback allows such reasoning, provided the probability of $\neg P$, given that faculties 1–4 testify P and faculty 5 testifies $\neg P$, is high. Similarly, No Feedback allows reasoning like, “perceptual faculties 1–5 all testify that P , and P is consonant with other things we already know. So the reliability of faculties 1–5 gains support.”

4.3 No Feedback and Deduction

All our bootstrapping cases were “inductive-inductive” cases: the subjects began with some premises, inferred some lemmas by non-deductive means, and inferred a conclusion from the lemmas by non-deductive means again. But some well-known “inductive-deductive” cases are problematic too. To adapt an example from Cohen (2002), consider someone who knows her vision to be reliable and, spying a red table, reasons as follows: “the table appears red, so it is red, so it is not a white table trickily lit to appear red.” Even though this subject knows that her vision is reliable, and this may support the proposition that the table’s appearance is not deceptive, there is something suspicious about her reasoning here. Even if her background knowledge about reliability supports her conclusion, the appearance of the table does not add to that support.²⁶ So it seems the subject should not form, base, or in any way support a belief that the table is not a trickily lit white one using this chain of reasoning.

We can extend the idea behind No Feedback to diagnose the error in such inductive-deductive cases. It is well known that incremental probabilistic support can fail to transmit across entailments: sometimes P increases the probability of L , and L entails C , but P does not increase the probability of C . The red table case above is one example of this phenomenon. More sophisticated ways of measuring probabilistic support, like the likelihood measure, also fail to transmit across entailments. And the “cumulative” transmission property fails too: sometimes P increases the probability of L , and L and P entail C , but P does not increase the probability of C . (We’ll see an example in a moment.)

²⁶ Cf. (Markie 2005) and (Cohen 2005).

So it's natural to ask whether some suspicious deductive inferences can be treated in the style of No Feedback. I will briefly describe three cases where such a diagnosis might be tried before remarking on the promise of such treatments.

1. *Easy Knowledge*. Some worry that, in a bootstrapping case like Charlie's, the subject's reasoning goes off the rails before the final inductive step that leads to the conclusion about reliability. It is natural to describe Charlie as reasoning something like, "First the gauge said 'full', so the tank was full; so the gauge was right on that occasion. Later the gauge said 'empty', so the tank was empty; so the gauge was right on that occasion too. Etc. So the gauge is super-reliable." So described, Charlie's reasoning includes suspicious inference-chains of the sort, *The gauge said the tank was in state X, so the tank was in state X; so the gauge was right on that occasion*, where the conclusion is deduced from the beginning premise and the subsequent lemma.²⁷

Charlie's suspicious deductions are cases where probabilistic support fails to transmit across the entailments in question.²⁸ More specifically, they are failures of *cumulative* transmission. The fact that the gauge read 'full' supports the hypothesis that the tank was full, and these two facts entail that the gauge read correctly; yet the fact that the gauge read 'full' does not support the hypothesis that the gauge read correctly. If, in the spirit of No Feedback, we think that beliefs' justificatory powers are limited by the probabilistic support their inferential ancestors can provide, then it is natural to think that Charlie would be unjustified in deducing that the gauge read correctly from the facts that the gauge read 'full' and that the tank was full.

2. *Harman-Vogel Cases*. A similar probabilistic pattern shows up in a puzzling class of cases identified by Harman (1973) and expanded by Vogel (1990). To adapt one of Vogel's examples, suppose you parked your car on Avenue A an hour ago and, noting that you remember parking it there, you infer that your car is currently parked on Avenue A. If I then I ask you whether your car has been stolen and moved somewhere else in the last hour, it would be inappropriate for you to respond that it has not, since it is on Avenue A. Apparently, the chain of reasoning, *I remember parking my car on Avenue A → It is on Avenue A → It has not been stolen and moved somewhere else* is illegitimate, despite the last inference being deductive.

Again, this is a case where probabilistic support fails to transmit across the relevant entailment, though this case is plain rather than cumulative. The fact that

²⁷ Charlie does have some independent reason for believing that the gauge was right on this occasion, since he knows that the gauge is reliable. But the fact that the gauge read 'full' adds nothing to his justification. So it seems inappropriate for Charlie to appeal to the fact that the gauge read 'full' in deducing that it read correctly. (Cohen 2005: 421)

²⁸ White (2006: 551) makes this point using a different example.

you remember parking your car on Avenue A an hour ago does probabilistically support the proposition that it is still parked on Avenue A, but it does not increase the probability that it has not been stolen and moved. So here too, a treatment in the spirit of No Feedback seems promising.

3. *Moorean Refutations*. Finally, consider a bit of modern Moorean reasoning: “It appears as if I have hands, so I have hands, so I am not a brain in a vat” (*Appears* \rightarrow *Hands* \rightarrow \neg *BIV*, for short). Some authors find this reasoning unproblematic, but many are suspicious (even if the subject executing the reasons knows her vision to be reliable). Indeed, White (2006) and Silins (2008) suggest that it is defective precisely because it is a case where probabilistic support fails to transmit across the relevant entailment.

If your having hands is probabilistically supported by its appearing as if you have hands, that appearance also supports the standard brain-in-a-vat alternative. By stipulation, the brain-in-a-vat hypothesis predicts the same observations as its common-sense counterpart, so it is confirmed by the same observation reports. So the probabilistic support for *Hands* does not transmit across the entailment to \neg *BIV*; *Appears* increases the probability of *Hands*, and *Hands* entails \neg *BIV*, but *Appears* does not increase the probability of \neg *BIV*. In fact, *Appears* must decrease the probability of \neg *BIV*, since it increases the probability of *BIV*. (Silins 2008: §3.3)

These observations, together with the success of No Feedback, support the conjecture that these inductive-deductive cases are defective because of the failure of probabilistic support to transmit across entailments. But there are obstacles to such a treatment.

First, there is an important disanalogy between inductive-inductive cases and inductive-deductive cases. In an inductive-inductive case, it can be that *L* is highly probable given *P*, and *C* is highly probable given *L*, yet *C* is not highly probable given *P*. So there is a sort of consequentialist motive for blocking the inference from *L* to *C*: if such inferences aren’t avoided, we will be able to draw wildly improbable conclusions from true starting points, and will frequently come to false conclusions in the long run. But the same cannot be said of inductive-deductive case. However probable *L* is given *P*, if *L* entails *C* then *C* must be at least that probable given *P*. So the same danger is not present.

Still, the fact that when *L* is at least *x* probable given *P*, *C* must also be at least *x* probable given *C*, does not change the fact that *P* may nevertheless fail to support *C*. This leaves open the possibility of motivating a No Feedback-like treatment of inductive-deductive cases by appeal to other sorts of undesirable consequences. For example, suppose you reason $P \rightarrow L \rightarrow C$ though, in fact, *P* is not what makes *C* probable; it is probable on independent grounds. Then, if you later abandon your belief that *P*, you might also abandon your belief in

C, thinking it was based on a falsehood. So a No Feedback-like restriction on inductive-deductive chains might serve to prevent the loss of highly probable beliefs, and thus the loss of many true beliefs in the long run.

A second obstacle to a No Feedback-like treatment of inductive-deductive cases is the *prima facie* tension with closure principles. If you cannot deduce that your car has not been stolen and moved from the fact that it is on Avenue A, then it seems you either do not know that it is on Avenue A, or else popular knowledge-closure principles (Williamson 2000; Hawthorne 2005) are false.

There is room to manoeuvre here. For example, Nagel (forthcoming) suggests that you might know that your car is parked on Avenue A, yet be unable to deduce that it has not been stolen and moved *from your knowledge that it is on Avenue A*, preserving the heart of closure principles. Knowing that your car is on Avenue A, if you actively judge that it is on Avenue A in trying to settle the question whether it has been stolen and moved, that judgment may not manifest your knowledge. If that's right, then you cannot deduce that your car hasn't been stolen and moved *from your knowledge* that it is on Avenue A; you can only deduce it from a non-knowledge-manifesting judgment that it is on Avenue A.

These worries and others need more careful consideration. So properly defending a No Feedback-style treatment of inductive-deductive cases would take us beyond the scope of the current discussion. Here I only want to note that many problematic deductive inferences discussed in the literature have the probabilistic structure noted above, and that the similarity to No Feedback suggests attempting a similar treatment. The parallel with our inductive-inductive cases and the possibility of such an extension will inform my defence of the view that the reliabilist can, like the internalist, use No Feedback to handle bootstrapping.

5 Reliabilism Revisited: Cribbing

We've seen that bootstrapping is a problem even for the internalist who requires antecedent knowledge of reliability, and that the internalist's response should be an appeal to No Feedback. I'll now argue that the reliabilist can crib the internalist's solution.

The reliabilist can solve her bootstrapping problem if she can argue that Roxanne's conclusion is not the result of a reliable process. This appears impossible at first, since trusting a reliable gauge is reliable and induction is presumably reliable as well. But one moral of our earlier discussion is that induction is *not* reliable, unless one respects the various defeaters for one's inductive steps. Naive applications of enumerative induction can lead one badly astray if one is not on guard against small samples, biased samples, etc. The obvious strategy for the reliabilist, then, is to argue that she is only committed to the reliable use

of induction, which the internalist has acknowledged requires respect for No Feedback. Since Roxanne does not respect No Feedback, her inductive reasoning in the second step of the bootstrapping argument is not reliable.²⁹

This defense of reliabilism invites a number of objections. I will very briefly consider four.

First Objection. This solution only partially blocks bootstrapping. If Roxanne is disposed to respect No Feedback almost all the time, then her use of induction will be reliable, but she might still bootstrap on occasion because she occasionally ignores No Feedback. On the occasions where she does disregard No Feedback, her process might still be called reliable: she deploys the same cognitive process she always does, even if it behaves unusually in this instance. So she knows the conclusion of her bootstrapping argument.

Reply. Even if this is an embarrassing result for the reliabilist, it is not a new embarrassment. The reliabilist already faces the problem that one can “get lucky”, abusing an ordinarily reliable process to gain knowledge. For example, suppose I am ordinarily careful in my use of inductive reasoning but, on one occasion, I carelessly conclude that all ravens are black from a sample of one black raven. Then I seem to know that all ravens are black according to reliabilism. Whatever the reliabilist says about this case, she will presumably say the same about the rare occasions where Roxanne carelessly overlooks No Feedback. She may bite the bullet, or she may argue that occasional misuses of a reliable method do not count as instances of the method. But whatever her response, bootstrapping adds nothing to her difficulties here.

Second Objection. The reliabilist is hiding behind the ambiguity that creates the generality problem for reliabilism.³⁰ Whether a token process is reliable depends on what general type it is subsumed under, and we can subsume Roxanne’s final, inductive step in the gas gauge case under different general types, some reliable and some not. If we regard Roxanne’s conclusion that her gauge is reliable as resulting from a process of the general type *an application of enumerative induction to true premises*, then her conclusion results from a reliable process, even if she is not disposed to respect No Feedback. To deem Roxanne’s conclusion the result of an unreliable process, we must view it as the result of a different kind of process, something like *an application of enumerative induction to believed premises*.

²⁹ Vogel (2000: p. 615) anticipates a superficially similar defense of reliabilism. He anticipates the reliabilist responding that Roxanne’s conclusion is the result of a bootstrapping process, but bootstrapping is unreliable, so she does not know her conclusion. The difference between our reliabilist and Vogel’s is this: Vogel’s reliabilist complains that bootstrapping as a whole is an unreliable process, whereas our reliabilist complains more narrowly that the second stage in Roxanne’s reasoning — an inductive inference without sensitivity to No Feedback — is unreliable.

³⁰ See (Goldman 1979; Feldman 1985; Conee & Feldman 1998) for background on the generality problem.

For applying enumerative induction to believed premises without regard for No Feedback is not reliable. So the reliabilist's defense depends on her classification of the final, inductive step in Roxanne's bootstrapping argument.

Reply. There is an obvious reason to prefer the latter classification. Roxanne tracks truth by belief, so her tendency will be to apply enumerative induction to believed premises, and only to true premises insofar as they are believed. Given two natural classifications of a process, it is more natural to subsume it under the type that better fits the system's dispositions to undergo similar processes. So it is more natural to describe what Roxanne does as applying enumerative induction to believed premises, in which case her conclusion is not the result of a reliable process (unless she is sensitive to No Feedback).

This response appears to hang on a stipulation about how Roxanne uses induction. Couldn't we just stipulate that Roxanne is disposed to use enumerative induction on true premises, whether or not she believes them? We could, but then we would have a case that does not pose any novel problem for reliabilism. If Roxanne uses enumerative induction on premises of the form *The gauge reads X and the tank is X*, not because she believes them but because they are true, then she does not bootstrap. She does not use the gauge's readings, or any conclusions drawn from them, to arrive at the conclusion that the gauge is reliable. What she does is admittedly mysterious without further elaboration; how she manages to dispose herself to use enumerative induction on true premises without using belief as a mark of truth is puzzling. But we have yet to hear any reason to suspect that an agent who is so disposed does anything more distasteful than what Roxanne already does just by trusting the unauthenticated gauge in her car.

Third Objection. If the reliabilist endorses No Feedback to block Roxanne from bootstrapping, she hamstring herself, rendering basic knowledge generally impotent.³¹ Consider a subject who does not know whether appearances are a reliable indicator of reality. Seeing what appears to be a hand, she can acquire the basic knowledge that there appears to be a hand, since introspection is reliable. But she cannot infer from there that there really is a hand, since (for her) the appearance as of a hand does not increase the probability that there really is a hand. In general, subjects who do not know what sources are reliable will not be able to use their basic knowledge to gain further knowledge the reliabilist wants.

Reply. The reliabilist is not as hamstrung as the objection makes out. In the example given, the subject may not be able to *infer* that there is a hand from her knowledge that it appears that way. But she can still come to know that there is a hand by forming that belief in the usual way. If she forms the belief that there is a hand in response to her vision, rather than via inference from introspective

³¹ I am grateful to Roger White for pointing out this concern.

knowledge about appearances, she knows that she has a hand, since vision is reliable.

No Feedback might defeat some reasoning that some reliabilists would miss, like van Cleve's (1984) inductive justification of induction. But the details here are contentious. First, it is contentious whether the inductive justification of induction is good reasoning, so it is not a clear strike against No Feedback if it rules it out. Second, it is not clear whether No Feedback does rule it out. Whether it does depends on facts about the correct epistemic probabilities for a subject who does not know that induction is reliable. One might think that the probability that future *F*s will be *G*s is unaffected by the fact that past *F*s have typically been *G*s, unless one knows that induction is reliable. In that case, No Feedback will block the inductive justification of induction. But one could insist that, even if a subject does not know that induction is reliable, the probability that future *F*s will be *G*s given that past *F*s have typically been *G*s is high. In that case, a track-record of successful inductive projections makes it highly probable that future inductions will be successful, and No Feedback lets the inductive justification of induction go through.

Fourth Objection. Even given No Feedback, the reliabilist is committed to an unsavoury consequence the internalist avoids. According to the reliabilist, Roxanne knows in each instance what the gas gauge says and thus what the state of the tank is. So she can deduce and come to know that the gauge made no errors. Even if No Feedback prevents her from inferring from there that the gauge is reliable, she has already gone astray: how could she know that the gauge made no errors just by reading what the gauge says? The internalist avoids this unseemly result, since she says that Roxanne does not know in each instance what the state of the tank is. (Roger White (2006) raises this point as a problem for dogmatism (Pryor 2000), and I am grateful to him for pressing me on it in conversation.)

Reply. I see two problems with this objection. First, even if the reliabilist does face a problem here, it is not the problem she set out to solve with No Feedback. The problem here differs from the original in important ways: it arises from deductive reasoning rather than inductive reasoning, it does not involve verifying one's method by illicit use of that method, and it does not involve an illicit double-amplification of one's evidence. It does have one thing in common with the original bootstrapping problem: the conclusion reached is not probabilistically supported by the starting premises. But to the extent that we see this as unifying this problem with the original one, we should be sympathetic to the idea of extending No Feedback to inductive-deductive cases, as discussed in §4.3. And if we do so extend it, Roxanne cannot deduce that her gauge made no errors.

This brings us to the second point, that there is no special problem for the reliabilist here; the internalist faces similarly unsavoury commitments. As we

saw in §4.3, even internalists face suspicious inductive-deductive chains, like those identified by Harman and Vogel. To emphasize the similarity between the reliabilist's and internalist's positions in this respect, consider what happens if Charlie, who knows that his gauge is reliable, performs a deduction similar to Roxanne's. After many repeated readings of the gauge, he could deduce that it had made no errors in an arbitrarily long sequence of readings. This is an unsavoury result in two respects. First, Charlie's reasoning uses premises of the sort *The gauge read X*, which do not probabilistically support the conclusion (even given his background knowledge that the gauge is reliable). And second, if the sequence is long enough, the conclusion is improbable even given the truth of his premises. The internalist can block this result by extending No Feedback to inductive-deductive chains, of course, but then the reliabilist can avail herself of the same solution.

6 Conclusion

I have argued for three claims. First I argued that bootstrapping is a problem even for internalists who reject basic knowledge. My argument for this claim proceeded by presenting cases. Second I argued that, of the solutions we considered on the internalist's behalf, appealing to the No Feedback defeater is the most promising one. There were essentially two threads of argument here: only No Feedback handles all of our cases correctly, and No Feedback is independently motivated by broad considerations to do with induction, justification, and probability. Third and finally, I argued that the No Feedback solution to the bootstrapping problem looks to be available to the reliabilist as well. The moral, I suggest, is that bootstrapping has less to do with internalism and basic knowledge, and more to do with the relationship between probabilistic support and step-by-step inductive reasoning.

References

- Bonjour, Laurence. 1985. *The Structure of Empirical Knowledge*. Harvard University Press.
- Christensen, David. 1992. Confirmational holism and bayesian epistemology. *Philosophy of Science* 59(4): 540–557.
- Christensen, David. 1999. Measuring confirmation. *Journal of Philosophy* 96: 437–461.
- Christensen, David. 2004. *Putting Logic in its Place: Formal Constraints on Rational Belief*. Oxford University Press.

- Cohen, Stewart. 2002. Basic knowledge and the problem of easy knowledge. *Philosophy and Phenomenological Research* 65(2): 309–329.
- Cohen, Stewart. 2005. Why basic knowledge is easy knowledge. *Philosophy and Phenomenological Research* 70(2): 417–430.
- Conee, Earl & Richard Feldman. 1998. The generality problem for reliabilism. *Philosophical Studies* 89(1): 1–29.
- Douven, Igor. manuscript. Further results on the intransitivity of evidential support .
- Douven, Igor & Christoph Kelp. forthcoming. Proper bootstrapping. *Synthese* .
- Feldman, Richard. 1985. Reliability and justification. *The Monist* 68(2): 159–174.
- Field, Hartry. 1978. A note on jeffrey conditionalization. *Philosophy of Science* 45(3): 361–367.
- Fitelson, Branden. 1999. The plurality of bayesian measures of confirmation and the problem of measure sensitivity. *Philosophy of Science* 66 (Proceedings): S362–S378.
- Garber, Daniel. 1980. Field and jeffrey conditionalization. *Philosophy of Science* 47(1): 142–145.
- Goldman, Alvin. 1979. What is justified belief? In George Pappas (ed.) *Justification and Knowledge*, 1–23. D. Reidel.
- Harman, Gilbert. 1973. *Thought*. Princeton University Press.
- Hawthorne, John. 2005. *Knowledge and Lotteries*. Oxford University Press.
- Jeffrey, Richard C. 1965. *The Logic of Decision*. University of Chicago Press.
- Kyburg, Henry E. 1970. Conjunctivitis. In Marshall Swain (ed.) *Induction, Acceptance, and Rational Belief*. D. Reidel.
- Levi, Isaac. 1980. *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability, and Chance*. The MIT Press.
- Markie, Peter. 2005. Easy knowledge. *Philosophy and Phenomenological Research* 70.
- Nagel, Jennifer. forthcoming. The psychological basis of the harman-vogel paradox. *Philosophers' Imprint* .
- Pollock, John L. 1986. The paradox of the preface. *Philosophy of Science* 53(2): 246–258.
- Pollock, John L. 1995. *Cognitive Carpentry: A Blueprint for How to Build a Person*. MIT Press.
- Pryor, James. 2000. The skeptic and the dogmatist. *Noûs* 34(4): 517–549.
- Pryor, James. 2005. There is immediate justification. In Matthias Steup & Ernest Sosa (eds.) *Contemporary Debates in Epistemology*. Blackwell.
- Silins, Nicholas. 2008. Basic justification and the moorean response to the skeptic. In *Oxford Studies in Epistemology*, vol. 2. Oxford University Press.
- Van Cleve, James. 1984. Reliability, justification, and the problem of induction.

- Midwest Studies in Philosophy* 9(1): 555–567.
- Van Cleve, James. 2003. Is knowledge easy — or impossible? externalism as the only alternative to skepticism. In *The Sceptics: Contemporary Essays*. Aldershot, UK: Ashgate Publishing.
- Vogel, Jonathan. 1990. Are there counterexamples to the closure principle? In M. Roth & G. Ross (eds.) *Doubting: Contemporary Perspectives on Skepticism*. Kluwer.
- Vogel, Jonathan. 2000. Reliabilism leveled. *Journal of Philosophy* XCVII(11): 602–623.
- Vogel, Jonathan. 2008. Epistemic bootstrapping. *Journal of Philosophy* CV(9): 518–539.
- Weisberg, Jonathan. 2009. Commutativity or holism? a dilemma for conditionalizers. *British Journal for the Philosophy of Science* 60(4): 793–812.
- White, Roger. 2006. Problems for dogmatism. *Philosophical Studies* 131(3): 525–557.
- Williamson, Timothy. 1997. Knowledge as evidence. *Mind* 106(424): 717–741.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford University Press.