

BLACK'S CLOSURE PATCH, AND WHY IT'S NO HELP FOR THE RELIABILIST

JOHN MILANESE

ABSTRACT. This paper argues that Tim Black's [1] closure-based solution to the problem of easy knowledge fails to save the reliabilist from the bootstrapping objection.

1. BOOTSTRAPPING AND THE PROBLEM OF EASY KNOWLEDGE

According to BonJour, reliabilism (of roughly the 1979 Goldman variety [9]) is the thesis that “what makes a belief justified is the cognitive reliability of the causal process via which it was produced, that is, the fact that the process in question leads to a high proportion of true beliefs” [7, p. 244]). This characterization of what it is for a belief to be justified, or count as knowledge, is *externalist*, because it holds that at least one element of justification need not be cognitively available to the knower forming the belief, *viz.* the *fact* about the reliability of the process that produced the belief.¹

Vogel [11] and Cohen [3] have argued that reliabilism makes knowledge *too easy*, and more specifically, that it makes second-order knowledge too easy. For the reliabilist, second-order knowledge is knowing *that one knows* that *P because* one has a true belief (formed by a reliable process) that *P* was brought about by a reliable process. More specifically, the problem for the reliabilist is that one can, via a *single* process such as color vision, arrive at

¹Following BonJour, let *internalism* be the thesis that “all of the elements needed for a belief to satisfy [the condition for justification] must be cognitively available to the person in question,” and let *externalism* be the negation of this thesis, that at least one element is not required to be cognitively available for a belief to be justified [7, p. 234].

knowledge about the reliability of color vision; using color vision alone, one can establish that color vision is reliable. This seems to contradict a common intuition, formulated by Heil [8] that one cannot use an instrument to check that instrument's reliability: one needs evidence from an independent source about the instrument's reliability. Heil provides an appropriate example from Wittgenstein:² one doesn't check the accuracy of a claim made in a newspaper by purchasing another copy of that paper and checking again.

For the problem of easy knowledge to arise, the instrument one gains easy reliability knowledge about need not be a cognitive process. Consider the following example of bootstrapping, analogous to the sort of scenario Vogel³ and Cohen describe. John hasn't a clue about how to speak ancient Aramaic, though he watched a one hour PBS special on the language a few months ago. As he walks to school, John stumbles across a huge tome. Opening it and glancing briefly at the pages, he has the occurrent belief that it is a manual for translating English sentences into ancient Aramaic (in fact, *it is*, and a rather good one as well).

- 1a. John notes (occurrent belief) that the manual says to translate English sentence S_1 into aramaic sentence A_1 .
- 1b. John knows that the manual says S_1 translates to A_1 .
- 1c. John thinks, "Huh, I guess this (A_1) is how you would say this English expression (S_1) in ancient Aramaic."
- 1d. Unbenounced to John, S_1 actually does translate as A_1 .
- 1e. John knows that S_1 translates as A_1 .

²Wigggenstein 1953/1961, §265, as cited by Heil [8, p. 316].

³For a reconstruction of Vogel's gasoline gauge example [11, p. 612-20], where the reliabilist gets to establish the reliability of a gauge by looking *only* at the gauge (and never checking the level of gas in the tank, etc.), see *Appendix 1*.

- ∴1f. Since John knows that S_1 translates to A_1 , and he knows *that the manual says* S_1 translates to A_1 , he further knows that, on *this occasion*, the manual is accurate.
- 2a. John notes that the manual says to translate S_2 as A_2 .
- ∴
- ∴2f John knows that, on this occasion, the manual is accurate.
- ∴
- ∴
- 3 John knows that the manual gives accurate translations on many occasions.
- ∴4 So, he knows that the process ‘believe what the manual says’ is a reliable belief-forming process.
- 5 He notes the manual says ‘ S_i translates to A_i ,’ believes that S_i translates to A_i , (and, unbenounced to him, it does, and so he knows it), and he further notes that his belief is formed by the reliable process ‘believe what the manual says,’ and so, John *knows that he knows* that the English sentence (S_i) “For if a man considers that he is something, when he is not, his soul is deceived,” translates into ancient Aramaic as (A_i) “*An gir anash sebar d’aytohi medem kad la aytohi naphsheh metaia*” [10].

In case you were wondering, the source I used for the last translation was a result of an arduous expedition that lasted all of 20 seconds on Google search engine. I know nothing about the translator, or for that matter about the website which linked to this PDF file (except that it was of a religious bent). Supposing that this translation I’ve found is fairly accurate, all I must do *to know that I know* that “To Whom the glory be, forever and ever, amen,” is “*D’leh shubkha lai-lam al-meen aw-mayn*” in ancient Aramaic, is just copy

down a few more of these sentences and have a few more occurrent beliefs. This is ridiculous. *Maybe* I know (first-order) what the Aramaic sentence means (though one might also reasonably say that I don't really know unless I have at least *some* reason for suspecting the source is reliable — this is a debate about the whether testimony is a basic source of knowledge). But surely second-order knowledge requires a little more work than I've done in the last five minutes. This is the problem of easy knowledge for reliabilism.

2. IS RELIABILISM THE SOURCE OF THE EASY KNOWLEDGE PROBLEM?

The reliabilist ought to be worried at this point. If the problem of easy knowledge has its source in reliabilism, and is not a direct result of a more general problem in epistemology, this is not an unreasonable ground for rejecting reliabilism. If the problem of easy knowledge has its source in externalism, the internalist has new grist to add to their mill. Just as one does not want their epistemology to be committed to skepticism, so to do we not want to make second-order knowledge effortlessly attainable. Lucky for the reliabilist and the externalist, neither position seems to be the *source* of the problem. In 2002, Cohen argued that any epistemology that allows for *basic knowledge* will face a similar problem of easy knowledge. Cohen defines an epistemology that allows for basic knowledge as a theory that *rejects* the following principle [3, p. 309]:

KR “A potential knowledge source K can yield knowledge for S , only if S knows K is reliable.”⁴

⁴KR entails skepticism, because KR entails an infinite regress. Assume there is a first piece of knowledge, K . It is necessary for K that the knower *already knew* the process that produced K is reliable. So, on KR, knowledge is impossible for finite beings.

But Cohen does not need to define basic knowledge this narrowly for his purposes. For his examples of easy knowledge, basic knowledge can be understood more broadly as any knowledge that is non-inferentially-justified true belief. Cohen argues that for “evidentialist foundationalism,” according to which, for example, I can know that the table “is red on the basis of its looking red,” it follows that I can know that the table is not white but illuminated with red lights such that it is *visually indistinguishable* from a red table [3, p. 312-4]. This should not be able to be known through vision alone. About this implausible consequence of evidentialist foundationalism, Cohen has the following in mind [3, p. 314]:

But the trivial entailment is not what’s in question here. And neither is the closure principle which allows you to know that the table is not white with red lights shining on it on the basis of your knowing it’s red. Both sides of the dispute agree that *if* we know the table is red, *then* we know it’s not white but illuminated by red lights. The point of the objection is that it seems implausible that we could come to know such a thing in this way.

Cohen seems to think that because the innocuous closure principle deductively *entails* something false, *viz.* that we know the skeptical hypothesis in question is false on the basis of visual evidence alone, it would follow by *modus tollens* that we cannot know it is red simply on the basis of its looking red.⁵ Thus, Cohen concludes that evidentialist foundationalism falls prey to a problem of easy knowledge; as he puts it, “if you allow for basic knowledge, there is

⁵But Cohen is not selling a *reductio* of basic knowledge. Instead, Cohen embraces a distinction from Sosa’s virtue epistemology, according to which one distinguishes between *animal* knowledge (basic, non-inferentially justified knowledge) and *reflective* knowledge. Part of Cohen’s solution to the easy knowledge problem requires denying that a closure principle holds over animal knowledge.

nothing to stop us from acquiring, by trivial inferences, all sorts of knowledge about how we are not deceived or misled by our belief sources” [3, p. 305].

This is good news for one who originally suspected that the problem of easy knowledge was a function of relaxing the internalist constraint on knowledge, for it is now argued (for example, by Bonjour [7]) that the internalist *must be* a foundationalist of the sort that Cohen describes, and thus would face the problem of easy knowledge. The reason the internalist must be foundationalist is that the internalist cannot be a coherentist (setting aside the options of infinitism and skepticism). See *Appendix 2* for this proof. At least, then, it seems, the internalist couldn’t use the problem of easy knowledge as an argument against externalism or reliabilism (if the internalist is also to avoid skepticism).

However, it has been argued more recently that the source of the easy knowledge problem is not basic knowledge. In 2005, Peter Markie [6] provides an example in response to Cohen: let a person have inductive evidence from previous experience that their color vision is reliable. Since they see that the table is red, and know by induction that seeing red things is a good indication of their color, this person knows (reflectively, not basically) that the table is red, and then by simple deductive inference they also know that the table is “not black but made to appear red to us by a deceptive god, and so on” [6, p. 409 *f.*]. But exorcizing demons in this way, by reflectively having a true belief about a table, seems a bit too easy. It appears then, *prima facie*, that basic knowledge is not the source of the easy knowledge problem.

Further, Markie points out that it is not clear that bootstrapping is always possible, and moreover that bootstrapping into knowledge about the reliability of *basic* cognitive processes may be *less* easy than doing so for non-basic processes. For example, bootstrapping into knowledge of the reliability of *color*

vision seems impossible, since the necessary *instances* of accuracy required for the bootstrapping *themselves require* a previously learned ability to identify colored objects reliably, which presumably relies on other faculties and processes “e.g. reports from others that what we are calling ‘red’ is red,” after which it is hard to say that this person lacks evidence that their color vision is reliable [6, p. 410].

Markie proposes that bootstrapping violates a so-called Independence Principle (IP): “Where a belief gains its *prima facie* justification for *S* just from the fact that it was produced by a particular faculty...the belief is not supporting evidence for *S* for beliefs concerning the reliability of that very faculty,” and since this principle can easily be independently motivated, the reliabilist could appeal to it against the charge of being able to bootstrap [6, p. 412]. However, Markie ends up arguing that a solution solely based on (IP), among other problems, will run afoul of the following “intuitively attractive principle: if it is reasonable for us to believe *P*, and the truth of *P* increases the likelihood that another prop, *Q*, is true, then *P* is a reason (perhaps a defeasible) for us to believe *Q*” [6, p.413].

Cohen replies [4] to the charge that theories which allow for non-basic knowledge will also exhibit the easy knowledge problem, by appealing to his contextualism in these case,⁶ that “the standards [that] govern how strong one’s evidence must be in order for one to know” are context-sensitive [4, p. 421].

⁶Cohen admits that this move is slightly *ad hoc*. “Why can I appeal to Contextualism to handle Markie’s case of easy knowledge that arises even for [non-basic structure] theories,” when “I rejected a contextualist treatment of the [basic-knowledge structure] case on the grounds that, by my lights, the BKS account of how we know that the table is not deceptively illuminated is implausible, even when the context is held fixed”? [4, p. 422].

In this debate about easy knowledge, one starts to get the feeling that there are all sorts of ways one *could* solve the problem, as it crops up here and there in different varieties. But we ought to be concerned with finding a principled, single solution to the problem. I will now shift gears and consider a current attempt at just this. Instead of making extra distinctions (how closure behaves over *animal* vs. *reflective* knowledge), or appealing to contextualism about knowledge in some cases, Tim Black [1] proposes to fix what he takes to be the cause of the problem, namely the *closure principle*: that “if S knows that P , and S knows that P entails Q , then S knows (or at least is in a position to know) Q ” [1, p.].⁷

3. BLACK’S CLOSURE PATCH AND BOOTSTRAPPING

In motivating a strategy such as Black’s, consider how closure allows for easy knowledge in each of the cases worried over. In the easy knowledge case Cohen describes, about knowing that the table is red by looking at it, Cohen recognizes that it is closure that commits the foundationalist to saying that one can in fact know there is no evil demon, if one knows the table is red, since the former follows easily from the truth of the latter.⁸ And Black argues that closure is behind both the bootstrapping case, as well the non-basic knowledge case of knowing the table is red by induction on past perceptual success.

In seeing this, consider again how the bootstrapping case gets off the ground. In the Aramaic manual case, by what rule is it the case that (1) one’s realizing that the manual says “ S translates as A ,” and (2) one’s knowing that S does

⁷Call this version *traditional* closure.

⁸As Cohen says, “But the trivial entailment is not what’s in question here. And neither is the closure principle which allows you to know that the table is not white with red lights shining on it on the basis of your knowing it’s red” [3, p. 314].

translate as A , make (3) follows: one knows that the manual is *accurate* in this case? From the perview of what is cognitively accessible, all one gets to ‘see’ is that the manual makes a claim that one believes to be true, and therefore it seems accurate in this case. But the reliabilist can’t push this line far, for the salient point they are committed to is the following. S translates as A , and one knows it. And the manual translates S as A , and one knows that too. What *makes* one have knowledge that the manual is *accurate* on this occasion is the following simple and obvious fact: it is blatantly analytic that if the manual provides a true claim about a translation *on one occasion*, then it *must* be accurate *on that occasion*; if I know the one, then I know the other (or am in a position to know). By what rule would the external fact that I know something *automatically* put me in a position to know anything else? So long as simple logical connections are part of my cognitive repertoire, I know the manual is accurate *because of* closure.

And closure is also sufficient for easy knowledge *via* a non-basic knowledge source (like induction on past information) that Markie describes: “closure guarantees that since I know the table is red, I know some skeptical hypothesis is false...Closure is essential, then, to each version of the problem of easy knowledge” [1, p. 2]. Black’s strategy about how to solve the problem of easy knowledge is utterly reasonable: first isolate the variable that is associated with every kind of easy knowledge and then *just change* (or improve upon) *that variable*. Why treat the symptoms by focusing on the level of theories (as Cohen does, distinguishing between animal and reflective knowledge for basic knowledge theories, or adopting contextualism in the case of non-basic knowledge), when one can treat the cause, which seems to be the closure principle?

After trying out a few patches on the closure principle that he ends up rejecting, Black settles on the following principle, which he maintains solves the problem of easy knowledge for both basic and non-basic theories of knowledge. The principle, *single source closure*, is adapted and revised from Markie's Independence Principle [1, p. 13]:

SSC “If S knows via K that P , and if S knows via K^+ that P entails Q , and if K or K^+ will allow S reasonably to believe that Q , then S knows that Q .”

Consider how subbing in this closure principle for *traditional* closure blocks the first step of a bootstrapping argument from knowing *via* color vision that (1) a table is red and one knows it and (2) the table looks red, to the intermediate conclusion of the bootstrapping procedure that (3): one knows that color vision is reliable on this occasion. Black notes that (1) and (2) set up (3) to follow by SSC according to the following closure formula (4), subbing in ‘color vision’ for K , (1) for P , and so on [1, p. 16]:

(4) “If I know *via* vision that the table is red and *via* introspection that the table looks red, and if I know *via* reason that its both looking red and being red entails that the table’s looking red is an accurate indication of its being red, and if vision, introspection or reason will allow me reasonably to believe that the table’s looking red is an accurate indication of its being red, then I know that the table’s looking red is an accurate indication of its being red.”

Notice, as Black points out, that in order for SSC to allow (4) to deliver the conclusion that I know that color vision is accurate in this case, I would need the following additional premise (5): “Vision, introspection or reason will allow

me reasonably to believe that the table's looking red is an accurate indication of its being red" [1, p. 16]. But *in this case*, where the competing hypothesis is that the table looks *exactly as if* it is red but is white and illuminated by red lights (or that it's black but the demon makes it look red), (5) is clearly false. By hypothesis, vision alone won't allow me to reasonably believe it's red and not white but looks exactly as if it was red — and so vision combined with any amount of introspection about the content of my visual experience won't do the trick via SSC *to establish* accuracy. Further, not only does reason not fail to help, but reason alone is capable only of "[reminding] us that the table can look red even when it *is not* red" [1, p. 16]. (5) is false, and as a result, SSC does not allow bootstrapping to get off the ground.

Black's general strategy here is utterly reasonable, and I think it can work in particular cases for the reliabilist. In epistemology, avoiding skepticism can result in letting in too much as knowledge, and when that happens, as it does in the problem of easy knowledge, one must let less in. Black is right in that the closure principle seems to be the leak that's letting too much in as knowledge, and so it needs to be patched up. Further, a natural way to do this is to build in something like the clauses in SSC into closure. But I doubt that this solution will work as a general solution *for the reliabilist* against bootstrapping. My doubt stems from how much the *reasonability clause* of SSC can demand of the reliabilist. To what extent could such a clause be acceptable for a reliabilist as an addenda to something so basic as closure, of veritable axiomatic status?

On the one hand, I take it that this reasonability constraint can't be too demanding, if it is to be independently motivated for the reliabilist. If my color vision is reliable, then I do get to know *via* vision alone that the table is red, for example, so long as the table is red. But the reliabilist has a right to be suspicious as to why this knowledge cannot be transmitted farther *via*

SSC and vision alone in general. Simply adding (as Black says), that *in this case* a relevant competing hypothesis is *this or that* skeptical scenario, like there being red lights that are themselves out of view (or demons), doesn't change the fact the reports of reliable vision are reasons to believe the table isn't white but actually red — *unless*, that is, it's the *context* that *makes* it unreasonable, such as in this skeptical context. But making context make all the difference in this way as to what counts as reasonable is just to sneak contextualism into the closure principle. If this is the suggestion, then Black's solution is just Cohen's 'theory-based' contextualist solution applied to closure. Moreover, and more importantly, I take it that the reliabilist may be unwilling to concede such general points, such as a contextualist amendment to closure: whether you are in epistemology class or not, the reliabilist does not want the *mere possibility* of demons being enough to taking away their knowledge that there aren't evil demons *if there aren't*. It is one of the worthwhile properties of reliabilism, and not one of the costs, that if there is no evil demon, the mere fact that it *seems possible to me* that there there could be shouldn't rule out that I could know that I'm not being deceived. At least, it's plausible to think a thoroughgoing externalist would agree with this intuition. If I do *in fact* know the table is red (and that I have two hands, etc.), then I do also know *in fact* that I'm not being deceived about this, plain and simple for the thoroughgoing externalist. The mere conceivability of a skeptical hypothesis shouldn't be enough, all by itself and in general, to keep me from knowing that it's false if it is, and knowing this *reasonably*. Again, it is at least plausible that the externalist reliabilist does not wish to abandon this intuition. So, if the reasonability clause of SSC is asking the reliabilist to set this intuition aside on contextualist grounds, it is asking too much.

On the other hand, if the reasonability clause is *weak enough* that it will not be too demanding for the reliabilist, or unacceptable, or not independently motivated, then bootstrapping may still be possible. Consider the following variant of BonJour's clairvoyance.⁹ scenario, which seems to satisfy SSC in a way that allows for bootstrapping to go through. While Norman slept, he developed an odd power of reliable clairvoyance as to the current location of the President. The way this clairvoyance works is that Norman will see (*via* vision) glowing letters appearing in the mirror¹⁰ on the opposing wall spelling out this location of the President, and he will believe it. Norman was also bestowed with the separate power to be able to reliably identify which beliefs result from his newfound clairvoyant ability.¹¹ Norman awakens the next morning and in awe says, "Woah! I guess the President is in New York City," as he stares at the glowing letters spelling this out. A small but relevant detail: earlier that morning, Norman's wife, like she usually does, brought the paper into the bedroom to read (the conservative press, mind you, which usually covers the President's main agenda for the day).

Here's a question. Can Norman, without checking the President's location, bootstrap his way into second-order knowledge about where the president is? Yes. Here's the first iteration:

- (1) Norman's belief forming process '**believe whatever you see the glowing letters say**' is in fact reliable.

⁹BonJour [2] As cited by Alvin Goldman's entry "Reliabilism" on the Stanford Encyclopedia of Philosophy website [5].

¹⁰Thanks to the imagination of P.D. Magnus.

¹¹But this isn't too hard: they're the beliefs that result directly from seeing the glowing letters!

(2) The President is in New York City, Norman believes this *via* the clairvoyance, and so he knows it.

∴(3) Since Norman knows that the President is in New York, and he knows that his clairvoyance formed this belief, he further knows that this means his clairvoyance was reliable on this occasion.

A closure principle is what would allow (3) to follow. Will restricting closure to SSC help keep out this problematic inference to (3)? We can find out by replacing the appropriate clauses (to get 4) and then inspect (4) to see if the missing premise needed, (5), is false. If it is, then (3) is blocked:

(4) “If I know *via* clairvoyance and vision that the President is in New York, and if I know that this seems right to me *via* clairvoyance and vision and thus I know *via* reason that my clairvoyance was accurate on this occasion, and if clairvoyance, vision or reason will allow me reasonably to believe that the glowing letters are an accurate indication of the President’s location, then I know that the glowing letters are an accurate indication of the President’s location.”

Now the missing premise which is needed for SSC to allow (3) to follow is (5): “Clairvoyance, vision or reason will allow me to reasonably believe that the glowing letter’s giving the President’s location is an accurate indication of the President’s location.” But (5) is true in the given case. It may be true, as Black says, that vision cannot be a reasonable ground for checking between normal-world and demon-worlds. But vision *will allow* for checking the newspaper, in which Norman would easily see reliable information as to the whereabouts of the President. Of course, if Norman actually picks up the paper, he’ll find

out his clairvoyance is indeed reliable, but then the ‘bootstrapping’ would no longer be illicit — he’d have independent evidence.

But here is the crucial point: it is not clear that Norman is required to actually check the newspaper for SSC to be satisfied; all that needs to be satisfied for SSC is that, in this case, vision *will allow* for reasonable belief, and it will, since if Norman checks the paper he’ll have all the evidence he needs, unlike in the table case, where vision will never be able to corroborate the ‘red table’ theory over the ‘black but seems red because of the demon’ theory.

In this case, vision will allow for reasonable belief in the accuracy of the clairvoyance process on this occasion, for looking at the newspaper will do. Black will be unsuccessful in persuading the reliabilist that Norman must now *actually check the paper for closure to work*. A core tenant of reliabilism is that one need not have cognitive access to a process K being reliable to for K to produce knowledge, and, if K is a very reliable, *the following process is necessarily no less reliable*: ‘believing deductive consequences of beliefs formed via reliable processes K .’ If K is clairvoyance, the process ‘believing the deductive consequences of one’s clairvoyant beliefs’ is *just as* reliable. It is unclear why the reliabilist would allow something tantamount to the denial of externalism at the level of closure. Why would the externalist conceded that a knower must cognitively possess the reasons (and not merely have them available) for justification to be transmitted over truth-preserving deductive inference? Again, if the reasonability clause of SSC is too strong, it will not be independently motivated for the reliabilist. And even if a reliabilist embraces a strong reading of the reasonability clause of SSC to block easy knowledge about the absence of demons and the like, it would look *ad hoc*.

If the reasonability clause of SSC is even read just slightly weaker, it's not strong enough to block bootstrapping. For example, an acceptable interpretation for the externalist of the 'will allow *S* to reasonably believe' clause is: 'would allow *S* to reasonably believe *had S* checked under the exact same circumstances.' But, if this is how the reliabilist would like to read SSC, then Norman gets to complete the first iteration of the bootstrapping procedure, and many others, for it would be enough that, had he checked the paper, he could reasonably believe *via* vision that the clairvoyant message is reliable on this occasion. I cannot identify a reading of the reasonability clause of SSC that (1) would be independently acceptable for the reliabilist, yet (2) require either that Norman actually check corroborating evidence (and so stop bootstrapping) or deny that step (3) of the procedure follows altogether for any normal cognitive faculty, as in the red table *vs.* demon table case, where one cannot distinguish between the two *by hypothesis*.

4. CONCLUSION

I think Black is right, that the closure principle is the culprit in all the varieties of the problem of easy knowledge. And so his strategy is a good one: fix the source of the problem. Nevertheless, in the case of bootstrapping as a problem for the reliabilist, it is hard to see how the reliabilist could have independently motivated reasons for accepting Black's proposed closure solution. Either the reasonability clause of Black's closure principle is to be read with some force, in which case it may be unacceptable or unmotivated for the reliabilist, or else it is read slightly weaker, in which case bootstrapping seems possible.

5. APPENDIX

1. Here's a reconstruction of Vogel's gasoline gauge example of bootstrapping [11, p. 612-20], where the reliabilist gets to establish the reliability of a gauge by looking *only* at the gauge (and never checking the level of gas in the tank, etc.):

- 1a. Roxanne notes (occurrent belief) on occasion 1 that her gauge reads 'F' (full).
- 1b. Roxanne knows the gas gauge reads 'F.'
- 1c. Unbenounced to her, the tank is full.
- 1d. Roxanne knows that her tank is full.
- ∴1e. Roxanne knows that, one this occasion, her gauge is accurate.
- 2a. Roxanne notes on occasion 2 that her gauge reads x .
- ∴
- ∴2e Roxanne knows that, on occasion 2, her gas gauge is reading accurately.
- ∴
- 3 Roxanne has knows that the gauge has been reading accurately on many occasions.
- ∴4 Roxanne knows the gauge is reliable.
- 5 She notes the gauge reads 'F,' the tank is full, she knows the gauge is reliable, and so she knows that she knows the tank is full.

If only we could know a gauge is working properly in this manner (and *know* that we know). We ought to tell the auto mechanics that they never have to dirty a dipstick or a rag over a properly working gauge again!

2. Here's a reconstruction of BonJour's argument that the internalism cannot hold a coherence theory of justification, and so, setting aside skepticism

and infinitism, the internalist must be a foundationalist [7, p. 241]. Suppose, for *reductio*, that some internalist epistemology is coherentist. This position would have to hold both the claims that (1) all elements of justification must be cognitively available, and (2) that all beliefs depend on other beliefs for their justification. It follows from (2) that this system of beliefs must be self-supporting in a way that is virtuously circular, otherwise justification faces either an infinite regress or vicious circularity, neither of which is any sort of justification at all. So, it follows from (2) that justification depends upon coherence (thus, a coherentist account). but consider some knower s and an arbitrary belief of s : b_i . By (2), whether b_i is justified is a function whether s 's system of belief is sufficiently coherent and whether b_i coheres strongly with the other beliefs. As required by (2), the coherence of this system is an element of justification. By (1), the fact about the coherence of this system (call it C) must be cognitively available to s . To say that this fact C must be available to s is to say that s must be able to have a belief about it. Let b_j be this belief about C . Since b_j is a belief of s , b_j is a member of s 's belief system. But, since b_j being justified *depends* upon C , yet b_j *asserts* this fact, b_j is immediately circular, since it's content amounts to 'my beliefs, *including this one*, are warranted.' In order to avoid this sort of immediate circularity, a belief such as b_j must somehow be about their belief system such it is not the fact that C that makes b_j justified: that is, b_j must enjoy warrant from the outside (or *independently*) of other beliefs. From here, the only way out for the internalist is to have certain beliefs *not* depend upon others for their justification. But this is to embrace foundationalism and abandon the possibility for a internalist coherentist epistemology.

REFERENCES

- [1] Black, Tim. "Solving the Problem of Easy Knowledge". *The Philosophical Quarterly*, forthcoming (2008).
- [2] Bonjour, Laurence. Externalist Theories of Empirical Knowledge. *Midwest Studies in Philosophy*, Vol. 5 (1980):53–73.
- [3] Cohen, Stewart. Basic Knowledge and the Problem of Easy Knowledge. *Philosophy and Phenomenological Research*, Vol. 65, No. 2 (2002):309–29.
- [4] Cohen, Stewart. Why Basic Knowledge is Easy Knowledge. *Philosophy and Phenomenological Research*, Vol. 70, No. 2 (2005):415–30.
- [5] Goldman, Alvin. Reliabilism. *Stanford Encyclopedia of Philosophy*, page plato.stanford.edu.
- [6] Markie, Peter. Easy Knowledge. *Philosophy and Phenomenological Research*, Vol. 70, No. 2 (2005):404–14.
- [7] Moser, Paul K., editor. *The Oxford Handbook of Epistemology.*, chapter "Internalism and Externalism", pages 234–65. OUP (2002).
- [8] Moser, Paul K., editor. *The Oxford Handbook of Epistemology.*, chapter "Mind and Knowledge", pages 316–35. OUP (2002).
- [9] Pappas, George, editor. *Justification and Knowledge: New Studies in Epistemology.*, chapter "What is Justified Belief?", pages 1–19. Reidel (1979).
- [10] Roth, Andrew Gabriel. "Word for Word Translation of the Epistle to the Galatians from the Ancient Aramaic Peshitta Text".
- [11] Vogel, Jonathan. Reliabilism Leveled. *The Journal of Philosophy*, Vol. 97, No. 11 (2000):602–23.