

# IS JUSTIFICATION EASY OR IMPOSSIBLE? GETTING ACQUAINTED WITH A MIDDLE ROAD<sup>1</sup>

BY: SAMUEL A. TAYLOR

FORTHCOMING IN *SYNTHESE*

Can a belief source confer justification absent antecedent justification for believing it's reliable? A negative answer quickly leads to vicious regress and skepticism. A positive answer, on the other hand, appears to commit one to allowing a pernicious form of reasoning known as "epistemic bootstrapping." Most discussions of bootstrapping do so within the context of foundational justification. By starting with a consideration of the justificatory power of deductively valid inferences, however, I will pave the way for a middle road that resolves our puzzle.

I will show that puzzles surrounding bootstrapping arise because we illicitly assume that justification either does not require any awareness of a source's epistemic credentials or justification requires doxastic awareness of said credentials. We can resolve our puzzle by splitting the horns and requiring a *non-conceptual* awareness of, or direct acquaintance with, a source's legitimacy.

Our problem concerns theories of justification that allow **basic sources** that can confer justification absent prior reason to believe they're reliable. Consider the following principle:

**JR** Necessarily, a subject S's belief source  $\Phi$  can confer justification only if S has antecedent justification for believing that  $\Phi$  is reliable.

Most contemporary epistemologists reject JR and allow basic sources. Process reliabilism, for instance, holds that a belief B is justified if and only if B is reliably produced.<sup>2</sup> The *de facto* reliability of my color vision is (absent defeaters) sufficient for the justification of the resulting color beliefs. Similar points apply *mutatis mutandis* to

---

<sup>1</sup> I would like to thank Richard Fumerton, Ali Hasan, Evan Fales, Gregory Landini, David Cuning for all of their feedback and conversations on these topics during my dissertation during which the ideas of this paper started to take form. I would also like to thank Ian MacMillan, Greg Stoutenburg, the anonymous reviewers, the audience at the 38<sup>th</sup> Midsouth Philosophy Conference, the audience at the 2014 Minnesota Philosophical Society Conference, and the audience at the 2014 Central States Philosophical Association Meeting for helpful feedback at various stages of this project.

<sup>2</sup> This is a sketch of process reliabilism. A fully developed version would distinguish foundational and inferential justification, explain defeaters, address the generality problem, etc. See Goldman (1979, 1986, 1993, 2008) and Lyons (2009) for good developments and defenses of the view.

## IS JUSTIFICATION EASY OR IMPOSSIBLE?

other externalist theories' rejection of JR. However, it's important to note that even most internalists allow basic sources. Consider the mentalist thesis that justification supervenes on one's mental states. A mentalist could deny basic sources and accept JR but most explicitly reject any such requirement. Phenomenal Conservatism (PC), for instance, is the view that *P*'s seeming true provides defeasible justification for believing *P*. PC requires (for foundational justification) that a subject *S* has a non-doxastic seeming representing *P* as being true. PC requires neither that these seemings are reliable nor that we have any reason to believe they are. Thus, PC holds that justification supervenes on one's mental states (i.e. seemings) but denies JR.<sup>3</sup>

Why not follow suit, deny JR, and allow basic sources? Assume color vision is a basic source. I can now *generate* justification for believing my color vision is reliable in ways that are too easy. I look to my left, see a basketball that looks orange, and thereby justifiably believe that the basketball *is* orange.<sup>4</sup> Next, I introspect and notice that the basketball *looks* orange. My color experience is accurate. Finally, I iterate this reasoning so as to produce inductive data supporting the reliability of my color vision. This is commonly referred to as bootstrapping.

Something has gone wrong. When bootstrapping a source confirms its own accuracy or reliability absent any independent checks. I can know at the outset that this procedure is guaranteed to reach the favorable conclusion.<sup>5</sup> Consider Cohen's

---

<sup>3</sup> Views closely resembling my sketch of PC have been defended by Foley (1983); Huemer (2001), (2006), (2007); McCain (2008), (2012); Pryor (2000); Skene (2013); and Tucker (2010). See Tucker (2013) for an anthology of papers both defending and criticizing these kinds of theories.

<sup>4</sup> Different epistemological theories deny JR in different ways. As such, these different theories will give different explanations of *why* the basketball's looking orange provide me with justification for believing that it's orange. On a process reliabilist view it's because the color experience plays a crucial causal role in a belief forming mechanism that is in fact reliable. On a reliable indicator theory the color experience provides justification for this belief because the existence of this color experience reliably indicates that an object of that color is present. According to PC, my color experience provides this justification because the propositional content of my experience and the resulting belief stand in a determinate-determinable relation (since the experience represents the basketball to be a specific shade of orange and my belief represents the basketball as merely having the more determinable property of being orange).

<sup>5</sup> This point is a bit tricky. Isn't there at least *some* sense in which perception could turn on itself when you use the iterated reasoning to reach a conclusion about a source's reliability? Couldn't you look at an object and it perceptually seems to be square but when you look at the object again it perceptually seems to be a different shape? *If* you have background beliefs that physical objects are not constantly changing shape randomly *then* there is a sense in which the bootstrapping reasoning could turn on itself. In which case, we simply have to explain that someone *without* these background *perceptual* beliefs about physical objects couldn't go through the bootstrapping reasoning and fail to reach the favorable conclusion.

## IS JUSTIFICATION EASY OR IMPOSSIBLE?

(2002) case where one employs this reasoning while looking at a white table illuminated by red lights. The table would deceptively appear red, trusting the perceptual appearance leads to the belief that the table is red, you introspectively note that the table *looks* red, and you thereby infer the (false) conclusion that your perception is accurate. Bootstrapping fails to include anything that helps discriminate accurate/reliable from inaccurate/unreliable sources in any interesting sense. None of the materials providing justification for the premises can plausibly be construed as good bases for believing the conclusion.<sup>6</sup>

One might respond by rejecting the possibility of basic sources and adopting JR. Why trust a thermometer, newspaper, or another source if we lack any reason to believe it usually gets things right? Unfortunately, this leads to vicious regress and skepticism. In order for a belief source  $\Phi$  to confer justification we need antecedent

---

<sup>6</sup> I need to be clear about the problematic nature of bootstrapping. Saying the justification is “too easy” has rhetorical pull but is too vague. There are two ways to state the worry. First, we can know prior to the bootstrapping reasoning that it couldn’t fail to reach the favorable conclusion. Thus, that the reasoning reaches the favorable conclusion tells us nothing *new* about the accuracy or reliability of the source. But, given the way we’ve allowed basic sources, a subject can nonetheless *gain* justification for believing the bootstrapping premises without being aware of anything that could provide justification for the reliability of that source. A different way to state the worry is that reasoning *transmits* justification and so the conclusion of an argument isn’t *ultimately* based on the premises but, rather, on the bases of those premises. Premises serve to *link* their bases to the conclusion. Now consider visual perception. The basis for the bootstrapping premises is the visual perception itself, nothing more! A visual perception of X isn’t evidence for that perception’s accuracy. This is ultimately the only basis for the conclusion of the bootstrapping reasoning, so when one goes through such reasoning, one doesn’t seem to base the conclusion on anything that actually supports the conclusion!

I’m *not* arguing that the problem with bootstrapping is circularity. Some forms of circularity might: (a) allow for the possibility of failure and (b) be consistent with the *ultimate* basis of the conclusion actually providing epistemic support to such a conclusion. Consider an argument for memory’s reliability that relies on the fact that one has both a current apparent memory as of a *previous* seeming to remember X and a current apparent memory that X. Such an argument for memory’s reliability is circular in some sense. Nonetheless, such a procedure could fail to reach the favorable conclusion. One could easily have a current apparent memory of previously seeming to remember parking one’s car in the west parking lot accompanied by a current apparent memory of the car having been in the east parking lot. And an apparent memory of X isn’t evidence of its own accuracy or reliability. However, the combination of an apparent memory of X with an apparent memory of a *previous* apparent memory of X is more plausibly construed as evidence of the former’s accuracy. I’m not endorsing these claims here. I’m merely rejecting the claim that the fundamental worry with bootstrapping is one of circularity. Finally, I should make one last clarification that the problem isn’t that bootstrapping generates justification for the conclusion! What is problematic is that, *given the way we’ve rejected JR*, the reasoning appears to generate justification for the conclusion even though the ultimate basis for this conclusion fails to include anything that can plausibly be said to support it. One way to resolve our puzzle—my way—is to provide a way of allowing basic sources that still guarantees that when one has justification for the bootstrapping premises will require that a subject uses materials that can act as adequate bases for believing the source is accurate/reliable. The tricky part is showing how to do this without accepting JR.

justification for believing it's reliable. Perhaps we can appeal to an independent source  $\Psi$  to confirm  $\Phi$ 's reliability. But, assuming JR,  $\Psi$  only confers justification if we have prior justification for believing it's reliable. The strategy of appealing to independent sources to provide *prior* confirmation of a source's reliability will be frustrated by the fact that we have a finite number of belief sources at our disposal.<sup>7,8</sup>

### 1. INFERENCE, REGRESS, AND BOOTSTRAPPING

It's trivial that sources of foundational justification are basic sources. However, basic sources *in our current sense* needn't be foundational sources. Consider induction. Induction only confers justification if one has justification for input beliefs about past observations. This is consistent with denying that induction's justificatory power is dependent on one's having antecedent justification for a meta-belief about its reliability. Denying the need for this meta-belief about the inferential rule is sufficient for it being a basic source in the sense relevant to bootstrapping. Let's call any such inference an epistemically **basic inference**.<sup>9</sup>

---

<sup>7</sup> The regress is more vicious than suggested above. Even *if* we have an infinite supply of sources, the regress is vicious. If JR is true and we must appeal to an independent source in order to provide prior justification for a source's reliability, there could never be an initial source to get the justificatory juices flowing.

<sup>8</sup> I don't have the space for detailed discussions of alternative solutions, but there are some very interesting and important attempts to solve the bootstrapping worry that I should at least mention. First, Vogel (2008) attempts to avoid bootstrapping by adopting a principle that essentially bans rule-circularity. I admit that bootstrapping is rule-circular; the problem with Vogel's proposal is that we can contrast bootstrapping cases with other rule-circular arguments where it's obvious that the bootstrapping reasoning is worse than other rule-circular arguments. Thus, rule-circularity doesn't adequately capture what is problematic about bootstrapping. Second, there is Titelbaum's (2010) diagnosis of bootstrapping as "no-lose investigations." Fn. 5 shows my sympathies with this diagnosis—though I'd make some qualifications. However, this only tells us why bootstrapping is worrisome, it doesn't provide a theory of justification that avoids the problem. Third, there is Weisberg's (2010) appeal to a principle concerning defeaters that he calls the "no feedback principle." Again, fn. 5 hints that I think Weisberg correctly identifies a problematic feature of bootstrapping, i.e. the bases of the premises of bootstrapping reasoning don't by themselves constitute an adequate basis for (e.g. don't make probable) believing the conclusion. Nevertheless, adopting Weisberg's no-feedback principle doesn't adequately resolve our puzzle. I reject Weisberg's no-feedback principle for the reasons given in Cohen (2010).

<sup>9</sup> Notice that the relevant notion of a basic inference is an epistemic and not a psychological concept. A basic inference is one that can *confer justification* absent a meta-belief about its reliability. What we might call a "psychologically basic inference" would just be an inference that we *make* absent any meta-belief about its reliability. An inference might be psychologically basic without being epistemically basic.

## IS JUSTIFICATION EASY OR IMPOSSIBLE?

What is required for an application of a particular inference rule to provide one with inferential justification for a particular belief? Consider the following principle of justification:

**IR** An application of inference rule R to justified premises provides S with justification for believing Q only if S has prior justification for believing R is (conditionally) reliable.<sup>10</sup>

IR has some initial plausibility. Why trust an inference to get things right if we lack any reasons to believe the inference reliably preserves truth?

Examples can also motivate IR. Consider Greco's (2010) case of a careless math student:

Suppose S is taking a math test and adopts a correct algorithm for solving a problem. But suppose that S has no understanding that the algorithm is the correct one to use for this problem. Rather, S chooses it on a whim... By hypothesis, the algorithm is the right one, and so using it to solve the problem constitutes a reliable [inferential rule]. It seems wrong to say that S thereby [has justification for believing] the answer to the problem (p. 149).

S is completely oblivious to the algorithm's reliability. It's a matter of luck from S's perspective that he followed a truth-preserving inference rule. We thereby judge that S lacks justification for the belief that results from this inference. This example and others like it provide some motivation for IR.<sup>11</sup> Nevertheless, there are good reasons to deny IR and allow basic inferences.

IR over-intellectualizes inferential justification. It's plausible that young children can transmit justification via applications of simple deductively valid rules such as *modus ponens* (MP). My young nephew Jack knows that his mom is working tonight and that *if* his mom is working *then* his dad will be watching him tonight. Jack can justifiably infer that his dad will be staying home and watching him. It's doubtful, however, that Jack can think about the structure of this inference let alone think about its conditional reliability.

An even more definitive reason to reject IR is that it leads to a Carroll style regress (See Carroll 1885). If MP isn't a basic inference then S's justified beliefs **that**

---

<sup>10</sup> IR is an inferential analogue of JR and very similar to Fumerton's (1995, 2004, 2006) inferential internalism.

<sup>11</sup> See Fumerton (1995, 2004, 2006) for additional cases motivating principles like IR.

## IS JUSTIFICATION EASY OR IMPOSSIBLE?

**P** and **that if P then Q** cannot, by themselves, produce in S a justified belief **that Q**. S must also have the justified meta-belief **that this inference rule is reliable**. This initiates a regress *even if S can know the inference rule is reliable a priori!*<sup>12</sup>

Requiring the justified meta-belief has the result that, in an apparent application of MP, S's forming an inferentially justified belief *Q* actually requires basing this on three distinct beliefs:

- B1. *P*
  - B2. If *P* then *Q*
  - B3. An application of MP is reliable
- 
- ∴ *Q*

Let's name the inference from B<sub>1</sub>-B<sub>3</sub> to *Q* *modus ponens\** or MP\*.

MP\* falls under IR's domain and transmits justification to S's belief that *Q* only if S has prior justification for believing it's reliable; an inferentially justified belief *Q* now requires four beliefs:

- B1. *P*
  - B2. If *P* then *Q*
  - B3. An application of MP is reliable
  - B4. An application of MP\* is reliable
- 
- ∴ *Q*

Name the inference from B<sub>1</sub>-B<sub>4</sub> to *Q* *modus ponens\*\** or MP\*\*. This too falls under IR's domain, and so the same reasoning repeats ad infinitum.

Accepting IR commits one to the view that an inference transmits justification only if we have an infinite number of meta-beliefs of ever increasing complexity as premises. This leads to skepticism about the possibility of ever transmitting justification via inference. The possibility of inferential justification requires the possibility of basic inferences.<sup>13</sup>

---

<sup>12</sup> This fact is often overlooked in the literature on the Lewis Carroll regress.

<sup>13</sup> Two anonymous reviewers rightly suggested that the regress for IR doesn't arise as straightforwardly as I've suggested. The Carroll style regress assumes justification requires a subject to *base* her belief in the conclusion on the belief that the inference rule is reliable. IR says following an inference rule transmits justification only if one has justification for the meta-belief about the rule's reliability, but it doesn't say that the subject must *base* the inferred belief on the meta-belief. IR is a principle of a belief's being *propositionally* justified. The Carroll style regress occurs at the level of doxastic justification (or well-foundedness). Justification for the meta-belief about the inference rule's reliability is part of what provides justification for believing an output belief *Q*. But in order for one's belief to be *justified* (i.e. have doxastic justification) rather than merely to be *justifiable* (i.e. have propositional justification) one must base the belief on whatever it is that provides propositional justification. So, if IR is true, following an inference will transmit justification only if the subject *bases*

## IS JUSTIFICATION EASY OR IMPOSSIBLE?

So let's assume MP is a basic inference and applying MP to justified premises is sufficient to transmit justification to a belief  $Q$ . Now we encounter the bootstrapping worry. I need only figure out that my belief  $Q$  is the result of MP in order to justifiably conclude that MP got things right. And by iterating such reasoning I generate justification for MP's reliability:

### BOOTSTRAPPING ON MP

1.  $Q_1$  (Justified by an application of MP to justified premises.)
  2. I formed my belief that  $Q_1$  on the basis of an application of MP to true premises.
  3. Therefore, applying MP to true premises got things right regarding  $Q_1$ .
  4.  $Q_2$  (Justified by an application of MP to justified premises.)
  5. I formed my belief that  $Q_2$  on the basis of an application of MP to true premises.
  6. Therefore, applying MP to true premises got things right regarding  $Q_2$ .
- 
- $\therefore$  MP is (conditionally) reliable.

Bootstrapping to a belief about MP's reliability without independent checks on the resulting beliefs cannot generate new justification for this conclusion.

Think back to Greco's careless math student. The student could confirm algorithm A's reliability by using A to produce a number of answers and using an independent check (e.g. a calculator or a solution manual) to confirm the answers. However, surely the student couldn't *discover* A's reliability by bootstrapping without any of the independent checks on A's results. But it's this reasoning without

---

the output belief on the relevant meta-belief. One possible maneuver here is to claim that the meta-belief isn't providing propositional justification but, rather, that it's part of the basing relation itself. Such a view avoids the regress because the meta-belief would be required for doxastic justification in virtue of being a component of the basing relation but avoids the need to *base* the inferred belief on the meta-belief by denying that it's playing a role in providing propositional justification. This response fails. If the purpose of the meta-belief is merely to fulfill a basing requirement it isn't clear what motivates IR's claim that one needs *justification* for the meta-belief. Presumably IR construes the justification for the inference's reliability as part of what provides the subject with assurance that the inferred belief is true. But if it's providing this assurance it looks like the meta-belief is part of the propositional justification, which leads to the basing requirement that gets the Carroll style regress started. We can further motivate the basing requirement by imagining I have a justified belief  $P$  and that I have justification for believing that inference rule R is reliable. Now imagine that I form belief  $Q$  by applying rule R to my belief  $P$ . However, if I don't follow rule R *because* of my justification for the meta-belief (I follow the rule because I recently learned the name of the inference pattern and liked the sound of it) then, even though I *have* justification for the inferred belief available to me, it doesn't seem that the resulting belief is actually justified. Surely there is something epistemically problematic with my belief. This suggests that, *if* IR is true, doxastic justification would require basing an inferred belief on the justified meta-belief that the inferential rule is reliable.

independent checks that is structurally isomorphic to the bootstrapping on MP that would generate justification if we assume that MP is a basic inference.<sup>14</sup>

## 2. ACQUAINTANCE AND A MIDDLE ROAD

Our puzzle arises because we overlooked a middle road regarding how an inference can transmit justification. We shouldn't leap from the fact that requiring *doxastic* awareness of the connection between one's premises and conclusion is too strong to the conclusion that *no* awareness is required. We should take the motivations for IR's claim that one must "see" the connection between the premises and conclusion seriously. However, these motivations by no means force us to construe the required "seeing" in terms of a justified *belief*.

Here the traditional notion of acquaintance is helpful. Acquaintance is a real relation of direct awareness that holds between a subject and a fact, state, property, or event. Acquaintance is not a representational state; it's the relation we stand in to things when they are directly given to us in consciousness. It's likely that children and even unsophisticated animals can be acquainted with a red patch in their visual fields or with the fact that they are in searing pain even if they lack the conceptual resources to represent such things in thought. Acquaintance is not representational and so it isn't the kind of thing that can be true/false or justified/unjustified. To be acquainted with a fact is not to know or to have a justified belief that the fact obtains.<sup>15</sup>

---

<sup>14</sup> Bootstrapping on MP may strike *us* as silly and irrelevant since *we* (i.e. ordinary humans) know a priori that MP is reliable. *We* don't need the bootstrapping reasoning. Noting that MP's reliability is knowable a priori, however, doesn't resolve the worry. That we can know MP is reliable a priori only allows us to *avoid using* bootstrapping to generate the relevant justification; we must nevertheless *permit* such reasoning to generate justification. Just because one can have a priori justification for a belief doesn't mean that one actually has it. In the case of the careless math student, algorithm A's reliability is a necessary truth and know-able a priori. Bootstrapping is still problematic when we assume that the student actually lacks this a priori justification. MP's reliability may be obvious to us (normal humans) while the reliability of other more complicated deductively valid inference rules is not. Nevertheless, there are possible creatures where things are reversed. MP's reliability is obscure and non-obvious to these creatures. If MP is a basic inference, we must permit these creatures to use bootstrapping to *generate* justification for beliefs about MP's reliability. You might respond that MP isn't justification conferring for us because its reliability is know-able a priori but, rather, because we actually have a priori justification regarding it's reliable; MP isn't justification conferring for the possible creatures under consideration because they lack the relevant a priori justification. Such a response is simply an appeal to IR and brings back the problems of over-intellectualization and Carroll style regress.

<sup>15</sup> Russell was the first philosopher to use term "acquaintance." See Russell (1912). Recently, a variety of epistemologists have attempted to resurrect a kind of classical foundationalism by appealing to this notion of acquaintance. See Bonjour (2000, 2003, 2010), Fales (1996), Fumerton (1985, 1995,



## IS JUSTIFICATION EASY OR IMPOSSIBLE?

The notion of non-conceptual awareness, or acquaintance, is controversial. Many find acquaintance mysterious and are skeptical that such a relation exists. So I should say something in defense of this critical notion. However, I want to note that an adequate defense of the existence of acquaintance would require an entire journal article dedicated to the topic. One could probably dedicate an entire book to the topic. Unfortunately, given my focus on showing how acquaintance resolves our bootstrapping puzzle, my comments must be brief and await further development.

It's often claimed that acquaintance is *sui generis* and unanalyzable. Richard Fumerton (1995, 2006) claims as much but suggests that he knows acquaintance exists because he's acquainted with it! Such a maneuver is *epistemically* legitimate.<sup>16</sup> Given his acquaintance theory of justification, this is exactly what Fumerton should say. Acquaintance with acquaintance would provide the best possible justification one could have for believing that this relation exists. Nevertheless, it's clear that this maneuver is *dialectically* ineffective. Claiming acquaintance with acquaintance does nothing to ease the concerns of someone who isn't sure what acquaintance is or whether it exists. If we're going to provide a *convincing* defense of the existence of acquaintance we should look elsewhere.

Getting a clearer grasp on what the relation is supposed to be and noticing various examples can help one to realize that such a relation exists. I've tried to say some informative things about acquaintance above. It's non-conceptual, it's the relation children and animals stand in to mental states like searing pain. However, further elucidation and examples are helpful. Consider an attempt from Hasan and Fumerton (2014) to single out and elucidate the relation of acquaintance:

[I]magine the following sort of familiar case. One is in pain but as one engages in an interesting conversation, one doesn't notice the pain for a while. After the conversation ends, one is again aware of the pain... [One possibility] is that the pain continued, but that the subject was temporarily unaware of it. On the supposition that [this possibility] makes sense, acquaintance is that relation the subject had to pain before

---

2006), Hasan (2012), McGrew & McGrew (2007), and Moser (1989) for attempts to resurrect an epistemology founded on the notion of acquaintance.

<sup>16</sup> As I mentioned in fn. 5, on my view bootstrapping is problematic but *not* because it's circular. As such, it is perfectly consistent for me to say that bootstrapping is problematic while at the same time admitting that acquaintance with acquaintance could justify one in believing acquaintance exists.

## IS JUSTIFICATION EASY OR IMPOSSIBLE?

the conversation, the relation that ceased during the conversation, and that began again shortly thereafter.

Despite the fact that Hasan and Fumerton take acquaintance to be a simple concept and indefinable, here they help draw our attention to the relevant phenomenon and thereby motivate its existence.

An anonymous reviewer wondered whether the non-conceptual awareness of acquaintance could just be identified with the relation of being “conscious of.” I admit that such a view of acquaintance is attractive. It’s plausible that we (along with children and animals) can be conscious of a pain or a red patch in our visual field without needing to form a higher-order conceptual representation of the fact that we’re in pain or that there is a red patch. Such a construal would make acquaintance no more mysterious than consciousness—pretty mysterious still but surely consciousness exists! Moreover, this construal might seem plausible when considering the case from Hasan and Fumerton. Unfortunately, identifying acquaintance with the “conscious of” relation isn’t going to allow acquaintance to play the heavy epistemological role that I want it to play. The reasons stem from the problem of the speckled hen. It’s likely that there are features of my mental life of which I’m conscious, e.g. a visual image of a 47-speckled hen, but to which it’s doubtful that I stand in any kind of epistemologically relevant relation, e.g. I don’t have justification for believing that I have a visual image as of a hen with exactly 47-speckles.<sup>17</sup>

I think all of this, however, actually leads to and suggests an account of acquaintance that makes it an incredibly familiar phenomenon and that gives us additional resources for motivating its existence. The key is to construe acquaintance as attention.<sup>18</sup>

First, there is phenomenological motivation for positing the existence of this kind of attentive awareness. We are all familiar with the phenomenon of mentally

---

<sup>17</sup> For discussions of the speckled hen see Feldman (2004), Fumerton (2004), Poston (2007), and Sosa (2009). Towards the end of his article, Feldman (2004) very briefly mentions the solution via an appeal to a kind of attentive awareness that I prefer.

<sup>18</sup> I actually think it’s best to understand acquaintance not as attention itself but rather a states being available for uptake by attention mechanisms. However, the reasons for this preference are tangential to my purposes in this paper. Nothing crucial will turn on whether we identify acquaintance with actually attending to a feature of my experience or that feature being *available* to attention.

“pointing” or attending to various features of our mental states. In fact, in Hasan and Fumerton’s example its plausible to explain the case as one where when we are engrossed in conversation our attention is diverted from the pain and then afterwards our attention was drawn back to the pain.

Second, identifying acquaintance with a kind of attention avoids worries associated with the speckled hen. *Perhaps* we can be conscious of a visual image with 47-speckles. Unfortunately, normal humans don’t attend to the feature of 47-speckledness. In fact, given normal cognitive limitations (i.e. there is a bottleneck regarding the amount of information one can attend to and take up for cognitive processing) we simply *cannot* attend to the fact that the image has exactly 47-speckles. Thus, attention seems better suited for doing epistemological work than the “conscious of” relation.

Third, while many advocates of acquaintance have advanced the claim that it’s sui generis and undefinable, one needn’t adopt this view. There has been a lot of very good work within psychology and philosophy of mind on the nature of attention. If we identify acquaintance with attention then it’s open to appeal to this literature in an attempt to elucidate the relation of acquaintance. In fact, one gets the feeling that some philosophers who express concerns about the mysteriousness of acquaintance are really just worried that acquaintance cannot fit into a naturalistic framework. However, the physical underpinnings of attention have been very well researched. If one identifies acquaintance with attention to a mental state or feature then giving a naturalistic account of acquaintance poses no special difficulties—i.e. it poses no more difficulties than giving a naturalistic account of mental states in general.<sup>19</sup>

Finally, we can treat acquaintance as a theoretical posit. I’ve suggested that acquaintance is a non-conceptual and non-representational form of attentive awareness. Why should we believe that a relation of awareness with said features exists? One way to motivate the claim that something exists is by showing the theoretical advantages that result from positing such an entity—think of David

---

<sup>19</sup> I’m not here endorsing the view that acquaintance is a physical or natural relation. Rather, I’m merely pointing out that when we identify acquaintance with attention then, *if* one can give a naturalistic/physicalist explanation of mental states *then* there is no reason to think one can’t do the same for acquaintance.

## IS JUSTIFICATION EASY OR IMPOSSIBLE?

Lewis' defense of modal realism. If positing the existence of acquaintance can solve various theoretical problems in epistemology and elsewhere, then this could justify positing its existence. This is probably the most dialectically effective strategy for arguing that acquaintance exists.

There is quite a bit of other theoretical work that acquaintance, understood as attentive awareness, can do. Attentive awareness can be used to make sense of how we acquire phenomenal concepts. Presumably an experience of redness is necessary for acquiring a phenomenal concept of redness. However, simply having an experience of redness isn't *sufficient* for acquiring the phenomenal concept. Presumably one must actually attend to that experience in order to uptake the experience into one's cognitive processing and use it to produce a phenomenal concept. In order to perform this task the attention mechanism must be non-conceptual since it's used as an explanation of how one *acquires* the concept in the first place. Attention can also be put to work in explaining how both demonstrative and *de re* thought is possible. When I have a certain experience I can form the demonstrative thought: "What is *this*?" But how can I think about a particular experiential feature without using some conceptual description of the experiential feature? Presumably I individuate this experiential feature via attention mechanisms that uptakes it into my cognitive processing. Yet again, in order to perform this task the attention mechanisms must be non-conceptual. I admit that an entire article dedicated to developing these and other theoretical motivations for position acquaintance is desirable. Unfortunately, if we are going to get to the main task of showing how an appeal to acquaintance can solve our bootstrapping puzzle we must rest content with this modest development.

Actually, given my stress on thinking of acquaintance as a theoretical posit, one can think of the remainder of this paper as developing in detail *one* of these theoretical benefits to positing an acquaintance relation. You can see the dialectic of the paper as follows: Acquaintance is a controversial notion but positing such a relation has the theoretical advantage of resolving regress and bootstrapping worries. Thus, we have *some* motivation for acquaintance as a theoretical posit.

So how does the relation of acquaintance help us regarding our current puzzle? Appealing to acquaintance opens up the possibility of a more moderate theory of

inferential justification that only requires a *non-conceptual* or direct awareness of the relevant connection between the input and output beliefs. This leads to the following principle of inferential justification:

**IA:** S has justification for believing  $Q$  on the basis of her justified belief(s)  $P$  only if S is *non-conceptually* aware of, or acquainted with, the fact that  $P$  epistemically supports  $Q$ .<sup>20</sup>

An inference will be epistemically basic for a subject S only if that subject is acquainted with an appropriate relationship holding between the premises and conclusion. When S is acquainted with this relationship, S's inference can transmit justification without S's needing to form any meta-belief.

However, before detailing how IA's appeal to acquaintance solves our puzzle, we need to consider a difficulty that arises if we were to identify epistemic support with any species of reliability.

### 3. AN INTERLUDE ON NON-REDUCTIVE EPISTEMIC SUPPORT RELATIONS

When introducing the fully general version of our puzzle I mentioned that a variety of both externalist and internalist theories of justification reject JR and thereby avoid the vicious regress and resulting skepticism. Similarly, a variety of both externalist and internalist views would attempt to avoid the skeptical worry about inferential justification by rejecting IR. So resolving this puzzle requires more finesse than simply rejecting externalism. I'm suggesting that the additional finesse amounts to requiring that a subject be non-conceptually aware of, or acquainted with, an

---

<sup>20</sup> An anonymous reviewer has suggested that the "only if" in IA is too strong. Couldn't a subject have a justified conceptual awareness of the fact that  $P$  supports  $Q$  without being non-conceptually aware of this fact (i.e. maybe a logician tells me that the inference is valid but I can't "see" it directly myself)? In which case, wouldn't this be sufficient for me to be justified in inferring  $Q$ ? I do think something like this is possible. However, since one cannot actually directly see the connection between  $P$  and  $Q$  and one might even become justified in believing the connection is present when it isn't—maybe the logician that gave you the testimony is generally reliable but made an error in this particular case—the *justified belief* that the connection holds between  $P$  and  $Q$  must be part of one's basis for believing  $Q$ . See fn. 13. However, in such a case I think part of what explains our justification for believing  $Q$  is the fact that we're acquainted with the fact that the beliefs *that P* and *that there is an appropriate connection (entailment, reliability, etc.) between P and Q* together support believing  $Q$ . In this case I think the support relation between these premises and conclusion is so obvious to us (I'm not sure I can even imagine what it would be like to be aware of this combination of beliefs and fail to see that they jointly support believing  $Q$ ) that we tend to ignore our awareness of it and fail to notice that it's playing a crucial role in providing our justification.

appropriate connection between one's premises (or—perhaps—evidence) and one's conclusion.

However, my suggested maneuver only provides a satisfactory solution to our puzzle if we also reject externalist accounts of what that *appropriate* connection is. Facts about reliability are best understood in terms of contingent facts about actual or counterfactual frequencies. Such facts don't seem to be the kind of thing with which a subject could be acquainted. Therefore, combining IA with reliabilism about epistemic support quickly leads to inferential skepticism.

We can avoid this problem only by rejecting an identification of epistemic support relations with a reliability relation in favor of a more internalist friendly account. In the end I think the right metaphysics of epistemic reasons identifies epistemic support with Keynesian probability relations.<sup>21</sup> Keynes (1921) suggested a logical theory of probability relations where these relations are modeled on entailment. When a proposition  $P$  makes epistemically probable (without entailing)  $Q$ , we can metaphorically describe this as *partial* entailment. Entailment is simply the upper limit of the epistemic probability relation.

If we adopt this approach, there are four key features of entailment that should be carried over to our conception of the epistemic support relation. Epistemic support (like entailment) is: (i) objective and independent of any minds conceptualizing the relation as existing; (ii) non-reducible to contingent facts about relative frequency or nomological relations<sup>22</sup>; (iii) an internal relation (i.e. the existence and natures of the relata necessitates the existence of the relational fact);

---

<sup>21</sup> See Keynes (1921). See Fumerton (2004) for a detailed defense of identifying epistemic probability with something like the Keynesian notion of probability.

<sup>22</sup> The best one could do in an attempt to reduce facts about entailment to facts about relative frequencies is to reduce them to frequencies *across all possible worlds*. But even this proposed reduction will fail. The fact that  $P$  entails  $Q$  doesn't obtain *because* of the fact that across all possible worlds  $Q$  is invariably true whenever  $P$  is true. We must remember that talk about possible worlds is metaphorical; the actual world is the one and only world and we must find truth-makers for claims about possible worlds in this world. What fact about the actual world acts as the truth-maker for the claim that across all possible worlds  $Q$  is invariably true whenever  $P$  is true? The best candidate is just the fact that  $P$  entails  $Q$ . Thus, if anything, the relevant fact about relative frequencies across possible worlds holds *because* of certain facts about entailment relations. Facts about relative frequencies across possible worlds are reduced to facts about entailment not the other way around! Here I'm being somewhat cavalier in my dismissal of modal realism. See Lewis (1986) for a defense of modal realism. I'll simply note that I hold the view that modal realism is quite implausible. Discussing the issue, however, would take us too far afield.

and (iv) knowable a priori. The essential idea is that epistemic support relations are fundamental and irreducible parts of the world.

Needless to say, this view of epistemic support is controversial. Nevertheless, this view can be motivated by well-known thought experiments; the most important in this regard is the “New Evil Demon.” Consider a modified version of the thought experiment. We think that our beliefs about past observations of the sun rising each day collectively constitute a sufficient epistemic reason to believe that the sun will rise tomorrow. Just thinking about the nature of our beliefs about the past and this belief about the future, we come to the belief that the relevant rational relation holds between these beliefs. But what if (unbeknownst to us) we actually live in a world where the laws of nature randomly change even though (by chance) the world has every appearance of regularity up to this point? Despite appearances, the past is not actually a reliable guide to the future. Nevertheless, since the world would still have every appearance of regularity, we intuitively think that my past observations would still provide epistemic support for my belief about the future. If this is correct then epistemic support cannot be identified with any external relation such as reliability or proper function.<sup>23</sup>

If I’m correct that the Keynesian notion of probability (with entailment as the upper-limit of making probable) best captures our notion of epistemic probability, then it isn’t implausible that we could be acquainted with epistemic support relations. A subject’s experiences and thoughts are paradigms of the kinds of things with which he or she could be acquainted. If epistemic support is an internal relation (like entailment) then we needn’t look beyond the nature of these experiences and thoughts themselves (items with which we are often acquainted) in determining the existence of an epistemic support relation.<sup>24</sup>

---

<sup>23</sup> See Cohen (1984) for the canonical presentation of the new evil-demon problem. See Fumerton (2006) for a dedicated defense of identifying epistemic probability with Keynesian probability relations. Recently the new evil demon thought experiment is often used to motivate the idea that only current mental states (or perhaps only current *conscious* mental states) can provide epistemic support. However, the other lesson to learn from the new evil demon isn’t one about the *states* that can provide epistemic support but rather about the nature of the epistemic support *relation* itself. This is the lesson that Cohen originally takes away from the new evil demon thought experiment.

<sup>24</sup> None of this commits me to saying that this is an easy task. Sometimes getting in a position where we “see” (i.e. are acquainted with) an epistemic support relation is incredibly difficult. In fact, in

Here I've appealed to the new evil demon thought experiment in order to offer some initial support for the idea that epistemic support (like entailment) is a necessary and irreducible relation. Such a view of epistemic support is admittedly contentious. For the remainder of this paper I'll assume this non-reductivism about epistemic support. If I'm right that the account helps resolve bootstrapping worries, then this provides additional theoretical reason to adopt the account.

#### 4. DEVELOPING OUR ACQUAINTANCE BASED SOLUTION

In order to see the solution that emerges, let us again restrict our focus to the conditions under which following a deductively valid inference rule such as modus ponens generates inferential justification. The two worries for a strong theory of inferential justification that accepts IR were: (a) over-intellectualization and (b) the Carroll style regress. A failure to require *any* awareness of the inferential connection, however, led to worries of bootstrapping.

My appeal to the more moderate acquaintance theory of inferential justification (i.e. IA) avoids the worry of over-intellectualizing inferential justification. Acquaintance is a non-conceptual form of awareness and it's likely that even children can stand in this kind of relation to simple entailment relations.

Children are probably not acquainted with *general* facts about the legitimacy of *general* forms of inference. However, it's plausible that children can be acquainted with *particular* instances of entailment relations that hold between *particular* propositions (even if they cannot categorize this relation as one of entailment). Just think back to your experiences teaching introduction to logic. It's often much easier for students to "see" the validity of a particular instance of an inference rule (i.e. a case where you use specific propositions instead of the variables "P" and "Q") than it is for students to "see" the validity of the inference considered as a *general* form.

Or, consider again the case of my nephew Jackson. If I were to ask Jackson, "Do  $P$  and *if P then Q* jointly entail  $Q$ ?" he would surely just stare at me in confusion. However, when I ask Jackson why he believes that his Dad will be watching his brother and him tonight, he will often say something like, "Because Mom has to

---

some cases it might be psychologically impossible for humans to become acquainted with an existing epistemic support relation.



## IS JUSTIFICATION EASY OR IMPOSSIBLE?

work and Dad stays home when Mom works.” Jackson is acquainted with the relevance his beliefs have to his conclusion that his Dad will be watching him tonight. Nevertheless, he isn’t acquainted with the relevance when we abstract away from the particular instance to the more general inference form. Acquaintance with the particular instance of an entailment is all that is required by IA for the instantiation of an inference rule to be justification conferring. Thus, we avoid worries with over-intellectualization.

My appeal to acquaintance also avoids Carroll style regress and the resulting inferential skepticism. Acquaintance with the fact that ***Q* is entailed by the conjunction of *P* and *If P then Q*** doesn’t require forming any belief about this entailment. Unlike theories of inferential justification that adopt IR, we needn’t treat the awareness of the evidential connection between premises and conclusion required by IA as anything like an additional premise in an argument. Thus, IA avoids the worries associated with accepting the stronger principle IR.

What about the bootstrapping worries associated with rejecting IR? Consider the belief (B) *that P* and the belief (B\*) *that if P then Q*. B and B\* are not evidence for the conditional reliability of modus ponens (MP). Thus, if following modus ponens can provide me with justification for believing *Q* absent *any* awareness of the entailment, then bootstrapping on MP would generate justification for believing that MP is reliable *even though I’m not aware of any evidence for this conclusion!* This is the problem with bootstrapping that we need to avoid.

The key to our solution is that conceptual relations and relations between propositions can constitute evidence for various beliefs.<sup>25</sup> We refer to this kind of evidence as *a priori* evidence. Consider my concepts “unmarried” and “bachelor.” These concepts stand in a part-whole relation and the existence of this relation is

---

<sup>25</sup> I use the term “evidence” in a way that some philosophers might find idiosyncratic. Evidence is in some sense a relational concept in that evidence is always evidence *for* some claim. Nevertheless, something’s being evidence doesn’t depend on its standing in any relation to an agent. On my usage, something can *be* evidence for believing a claim even if no one *possesses* that evidence. Having evidence is like having an apple. Something can be an apple even though no one *possesses* that apple. One consideration in favor of this conception is the idea that we can gather or *find* evidence to support various views or theories. We think of the evidence as being out there ready for us to find. But this presupposes that something can be evidence for a claim independently of our currently having that evidence. Some philosophers, however, might think that having evidence is more like having a golf partner. George cannot be a golf partner independently of someone *having* George as a golf partner (see Schroeder 2008).

## IS JUSTIFICATION EASY OR IMPOSSIBLE?

evidence for the claim that all bachelors are unmarried. Similarly, consider any propositions  $P$  and  $Q$  where the former entails the latter. The existence of the entailment is a priori evidence for believing that following an inference rule that moves from  $P$  to  $Q$  is conditionally reliable. *That  $X$  entails  $Y$*  itself entails (is conclusive evidence) that  *$X$  is a reliable indicator of  $Y$* . I'll show below that this fact is crucial for understanding why requiring acquaintance with the entailment relation is capable of resolving issues surrounding bootstrapping on deductively valid inference rules.

According to IA, an application of MP to justified premises confers justification only if I'm acquainted with the fact that the premises entail the conclusion. The existence of this entailment is itself evidence for the conditional reliability of my inference. Therefore, if I use MP to form a justified belief that  $Q$  as a premise in bootstrapping, such reasoning doesn't reveal any new evidence for its conclusion. All bootstrapping does is reveal the fact that something of which I was already aware (i.e. the entailment) is itself a priori evidence for MP's reliability. But this isn't worrisome. It's perfectly reasonable to maintain that an awareness of an entailment relation can justify a belief in the reliability of that inference. Bootstrapping cannot reveal anything that I wasn't already in a position to know independently of that reasoning.<sup>26</sup>

The point I'm making is subtle and could easily be misunderstood. I'm not claiming that nothing *ever* changes when you go through the bootstrapping reasoning.

---

<sup>26</sup> There is a difference between the way one's awareness of the entailment relation provides justification for believing in the inference's reliability *absent* the bootstrapping and *with* the bootstrapping. If one were to get justification for believing the conclusion *absent* the bootstrapping one would have a kind of foundational justification that depends on being acquainted with the entailment and acquainted with the fact that this entails the inference's reliability (see 5's acquaintance theory of foundational justification). However, when one goes through the bootstrapping one doesn't (or at least *might not*) *directly* grasp the fact that the entailment entails reliability. Rather, one's grasp of this support relation between the entailment and the reliability is had indirectly via one's acquaintance with a *chain* of support relations. This is analogous to the following idea. Consider a necessarily true material conditional *If  $P$  then  $Q$* . One might arrive at justification for believing this conditional in two different ways. First, one might be able to immediately see a priori that the entailment holds. Second, a person might be such that they cannot immediately see the connection between  $P$  and  $Q$ . However, she might be able to see the connection between  $P$  and  $R$ ,  $R$  and  $T$ ,  $T$  and  $V$ , and  $V$  and  $Q$  in such a way that she can go through a conditional proof to arrive at justification for believing the conditional. Each is in a sense aware of the connection between  $P$  and  $Q$  but the former is directly aware of this connection and the latter is only indirectly aware of the connection via a chain of connections leading from  $P$  to  $Q$ .

## IS JUSTIFICATION EASY OR IMPOSSIBLE?

Prior to having justification for believing the premises of bootstrapping I'm aware of a priori evidence for the reliability of MP. Nevertheless, I might not have justification for believing that MP is reliable. IA has the result that S's following MP is epistemically efficacious only if S is aware of a priori *evidence* for MP's reliability but this isn't the same as requiring that S have a priori *justification* for believing that MP is reliable. The latter requirement leads to the Carroll style regress.

Let me explain more. We can distinguish between a weak and a strong notion of S's being aware of evidence for an inference's reliability. On the weak reading, S is aware of something that *is* evidence for an inference's reliability but needn't be aware of its status as evidence for the reliability. On the strong reading, S is aware of the evidence and aware of the fact that it is evidence for the inference's reliability. IA has the implication that an inference is only justification conferring if one is aware of a priori evidence for the inference's reliability, *when this is given the weak reading*. This is why bootstrapping doesn't produce new evidence for a source's reliability. Bootstrapping just reveals the epistemic support relation that holds between the entailment relation and the claim that there is a reliability relation. Bootstrapping can move you from a mere weak awareness of a priori evidence for MP's reliability that doesn't include an awareness of this evidential status to the strong awareness that does. However, the bootstrapping reasoning isn't necessary for achieving this result. When you arrive at the conclusion of the bootstrapping reasoning you realize that you could have known this reasoning would reach the favorable conclusion prior to actually engaging in it; you didn't need to go through the bootstrapping in order to justify the conclusion.<sup>27</sup>

---

<sup>27</sup> As an anonymous reviewer correctly points out, given my distinction between evidence and justification, my view does have the result that while the bootstrapping doesn't reveal new evidence for the reliability of MP it does produce new justification. However, I do not see this as problematic since, given how I've made room for basic sources, justified beliefs in the premises of the bootstrapping reasoning must be based on an acquaintance with the entailment. Thus, part of the ultimate basis for the conclusion of the bootstrapping reasoning is one's acquaintance with the entailment relation. The fact that X entails Y, however, is in fact a good basis (or evidence) for the inference from X to Y being reliable. As such, I don't see this result as problematic. The problem with bootstrapping is that the views considered earlier allowed an application of MP to transmit justification to *Q* when a subject had justified beliefs that *P* and *if P then Q* but lacked any awareness of the entailment. Thus, the ultimate basis for the conclusion that MP is reliable (i.e. the basis of the premises) didn't have to include anything (e.g. the entailment) that actually supported this conclusion. This is the problem that I've avoided. These views allowed one to use bootstrapping to justify the reliability of source using bases that didn't provide any support for this conclusion. On my view,

## IS JUSTIFICATION EASY OR IMPOSSIBLE?

At this point someone might worry: how do we know that acquaintance is reliable?<sup>28</sup> I want to say in response that acquaintance is a *factive* relation of awareness. One cannot be acquainted with a pain if a pain does not exist. Your attention mechanisms cannot uptake your pain state for cognitive processing if that pain state does not exist. The question of the reliability of acquaintance is, by definition, a non-starter. Some might say that acquaintance is infallible. However, I think that a more accurate response to this worry is that applying these concepts of infallibility and reliability to acquaintance is a category mistake! Acquaintance is a non-judgmental and non-conceptual awareness. To be acquainted with pain is not to be in a state the *represents* you as being in pain. Insofar as acquaintance doesn't involve entertaining a representation of the objects of awareness the question of misrepresentation simply doesn't make sense.<sup>29, 30</sup>

---

bootstrapping can only provide new justification for the conclusion when one bases her belief in the premises on one's awareness of the entailment, which really is a good basis for believing the conclusion!

<sup>28</sup> Thanks to an anonymous reviewer for mentioning this concern and pushing me to respond.

<sup>29</sup> It only makes sense to apply the concepts of infallibility or reliability to the beliefs resulting from one's acquaintance with certain facts. Clearly one's belief that one is in pain is fallible. People might base this belief on something other than their acquaintance with the fact that they are in pain. In fact, a belief even when based on acquaintance might be fallible since, given the principle of foundational justification I present in section 5, one might be acquainted with: (i) a non-doxastic state that merely makes probable (maybe a mere seeming that one is in pain) the fact that one is in pain and (ii) the epistemic support relation. In such a case one has defeasible foundational justification for believing one is in pain. At this point you might ask what evidence we have that beliefs formed on the basis of acquaintance are reliable. I think this is to ask the wrong question. My acquaintance isn't what reliably indicates or stands in an epistemic support relation to my belief that I'm in pain. It's the object of my acquaintance that acts as evidence and epistemically supports my belief. My acquaintance is just part of what enables this evidence to be epistemically efficacious. Thus, the relevant question is what evidence I have for believing that the *object* of my acquaintance is a reliable indicator of the belief's truth. We'll see in section 5 that the answer is that the epistemic support relation between my non-doxastic foundational evidence and my foundational belief is evidence with which I'm acquainted for the reliability of the non-doxastic foundational evidence. See section 5 for more details.

<sup>30</sup> The same anonymous reviewer raised another nearby concern. Even if acquaintance is *factive* there seem to be cases where a subject *believes* she is acquainted with something when she isn't. So why don't we need evidence that we're reliable at distinguishing cases of genuine acquaintance from cases where we falsely believe ourselves to be acquainted with an object. This is an incredibly interesting question and relevant to assessing the prospects of an acquaintance theory. This issue revolves around skeptical arguments from the possibility of deception, but it's somewhat tangential to the *structural* issue that arises when we consider bootstrapping. Moreover, one should (and many have) dedicate an entire paper to evaluating the implications of this skeptical possibility. For nice discussions of the issue see Ballantyne (2012), Fumerton (1995, 2006b, 2010), Hasan (2013) [Hasan doesn't explicitly deal with the issue but some of his remarks are indirectly relevant], and Poston (2010). I think the best response to this worry involves arguing that it rests on a level-confusion. Finding out that I'm acquainted with pain is relevant to justifying my belief *that I'm justified in believing I'm in pain* but not to justifying my belief *that I'm in pain* (an awareness of my awareness of pain doesn't seem to provide any

## IS JUSTIFICATION EASY OR IMPOSSIBLE?

I've only explained how adopting IA allows us solve our puzzle as it applies to the justificatory power of following deductively valid inference rules. It remains to be seen how to extend this solution to both non-deductively strong but deductively invalid inference rules and to supposed sources of foundational justification. Before doing this, however, it will be helpful to summarize the upshot of my discussion. My discussion suggests reversing an intuitive order of explanation. Following a deductively valid inference rule doesn't make beliefs rational *because* we are rational in believing that following the inference rule is reliable. Nor does following a deductively valid inference rule make a belief rational because it is in fact reliable. Rather, in some sense, we are rational in believing that following a basic deductively valid inference rule is reliable *because* we are aware of the fact that such an inference makes the output beliefs rational (i.e. the input and output beliefs of the inferential rule stand in the relevant entailment relation).

### 5. GENERALIZING OUR SOLUTION TO A FULL THEORY OF JUSTIFICATION

My remarks in the last section apply *mutatis mutandis* to any supposed basic inference that is deductively valid. But can we generalize the solution in order to allow for basic inferences, such as induction, that are not deductively valid?

Assume that some past observations  $O_1$ - $O_N$  epistemically support but fail to entail a claim  $Q$  about the future (or the unobserved past). I've suggested understanding this as (metaphorically speaking) partial entailment. This is an irreducible epistemic support relation between propositions. According to my acquaintance theory of inferential justification (i.e. principle IA), inferring  $Q$  from my justified beliefs that  $O_1$ - $O_N$  will only produce a justified belief that  $Q$  for me if I'm acquainted with the partial entailment relation.

---

assurance in addition to that provided by my being aware of the pain). In order to be justified in believing I'm in pain I need only *be* acquainted with the fact that I'm in pain and the fact that this entails the truth of my belief that I'm in pain. I needn't be justified in *believing that I'm acquainted with my pain*. The question of whether we can reliably distinguish genuine cases of acquaintance from merely believed cases is relevant to giving an epistemic assessment of the higher-order belief *that my belief is justified*. Many acquaintance theorists are adamant that one can appeal to acquaintance while rejecting access requirements on justification such as the JJ principle. As I said, there are incredibly complicated and interesting issues surrounding the possibility of believing you're acquainted with something when you aren't. I wish I had time to say more about this issue. However, this skeptical possibility is best tackled separately from that of bootstrapping since this possible deception is to be construed as a skeptical worry rather than a worry about the *structure* of justification.

Such a theory avoids worries about over-intellectualization and about Carroll style regress in the same way it avoids these worries within the context of deductively valid inference rules. Acquaintance with the partial entailment doesn't require a belief or even the conceptual resources to represent the inference and its reliability.

However, applying our solution to the bootstrapping problem in the case of basic non-deductive inferences might appear more difficult. Partial entailment is to be construed as a necessary relation that holds in all possible worlds, including worlds where the past isn't actually a reliable guide to the future. This is one of the upshots of the new evil demon thought experiment discussed earlier. As I noted, epistemic support relations are not analyzable in terms of contingent nomological relations or frequencies across possible worlds. Reliability, however, is best understood in these terms. Therefore, the existence of epistemic support relations and reliability relations can theoretically come apart. The fact **that  $O_1-O_N$  epistemically supports  $Q$**  does *NOT* entail **that  $O_1-O_N$  are a (conditionally) reliable indicator for  $Q$** .<sup>31</sup>

Adopting IA doesn't guarantee that one could only bootstrap on a non-deductive inference rule if one was antecedently aware of a *conclusive* reason for believing that following the rule is reliable. One might thereby conclude that bootstrapping on non-deductive inference rules ought to be banned. Perhaps we should move to a compromise view that accepts a modified IA with its domain restricted to deductively valid inferences but also accepts a modified IR with its domain restricted to non-deductive inferences. Unfortunately, this quickly leads to the implausible view often associated with classical foundationalism that only deduction can transmit justification.

Remember, however, that entailment is supposed to be the upper-limit of the epistemic support relation. The fact that *P's entailing Q entails P's being a reliable indicator of Q* should be seen as the upper limit of a more general feature of epistemic support; *P's epistemically supporting Q* itself epistemically supports

---

<sup>31</sup> This fact is the basis of Greco's (1999) criticism that Fumerton's principle of inferential justification fails to make a subject "sensitive" to a source's reliability. What I say below can be used to respond to this kind of criticism. Also note that, as I mentioned earlier, things are different when we restrict our focus to the upper limit of a case where one's evidence *entails* one's belief. The fact that *X entails Y* itself entails that *X is a reliable indicator of Y*.

the claim that *E is a reliable indicator of Q*. In a case where our evidence entails a certain claim, this fact entails the relevant reliability. In a case where our evidence only provides defeasible reason to believe a claim, this fact provides defeasible reason to believe that the relevant reliability relation holds. Bootstrapping is an illustration of this very fact!

It might help if I also give some intuitive motivation for the idea that the existence of an epistemic probability relation itself makes epistemically probable the existence of a reliability relation. Consider a fair 6-sided die. When I roll the die, each side is equally epistemically likely to land face up. Thus, for any individual roll, it's epistemically likely that the die will not land with the 6-side up. Now I ask, "When I roll this die 1,000 times, is it likely that it will not usually land with the 6-side up?" Of course! The fact that it's epistemically likely that, on any individual role, the die will not land with the 6-side up makes it epistemically likely that the die will not tend to land with the 6-side up. None of this is threatened by the fact that we are all well aware that it is *possible* for a fair die to be rolled 1,000 times and *usually* land with the 6-side up. It's even possible to roll a fair die 1,000 times and it *always* land with the 6-side up. We still think that there are "more" possible worlds<sup>32</sup> where a fair die is rolled 1,000 times and does not tend to land with the 6-side up than there are worlds where a fair die does tend to land with the 6-side up. Here we have an illustration of my claim that the existence of an epistemic probability relation itself provides (defeasible) epistemic support for a claim about reliability.<sup>33</sup>

With this in mind, let us reconsider the possibility of a basic inference that is not deductively valid. Induction—here I have in mind enumerative induction—appears

---

<sup>32</sup> One might be worried about how such a claim could be true given that there are an infinite number of possible worlds where each scenario holds. This arises from a skepticism that probabilities are well defined when the relevant events or fact belong to an infinite class. I do think, however, that we have an intuitive grasp on this idea. What we need is a measure on the worlds in order to determine the relevant ratio. What is the truth-maker that determines the proper measure and the resulting ratio? In my view the relevant truth-maker in the actual world is just the existence of the epistemic probability relation. In the same way that the existence of an entailment relation in the actual world is the truth-maker for certain claims about *perfect* reliability across all possible worlds. I actually take the fact that the non-reductive account of epistemic support can make sense of the claim that there are "more" possible worlds (when quantifying over *all* possible worlds) where the die lands without the 6-side face up than possible worlds where the die lands with the 6-side face up to be a mark in its favor.

<sup>33</sup> The examples from Hawthorne's (2002) and Turri's (2011) defense of the contingent a priori can also be modified so as to help illustrate this idea.

## IS JUSTIFICATION EASY OR IMPOSSIBLE?

to be a prime candidate. So let us assume that an application of induction can confer justification absent a justified higher-order belief that induction is reliable. Now let us consider what happens when we attempt to use the bootstrapping reasoning:

### BOOTSTRAPPING ON INDUCTION

1.  $P_1$  (Inferred from inductive data supporting  $P_1$ )
  2. Following the rule of enumerative induction is what tells me that  $P_1$ .
  3. Therefore, following enumerative induction got things right regarding  $P_1$ .
  4.  $P_2$  (Inferred from inductive data supporting  $P_2$ )
  5. Following the rule of enumerative induction is what tells me that  $P_2$ .
  6. Therefore, following my inductive rule got things right regarding  $P_2$ .
  - [Repeat for  $P_3$ - $P_N$ ]
- 
7. Therefore, enumerative induction is reliable.

Such bootstrapping is worrisome if we can be justified in moving from 1 to 2 without any awareness of this inference's legitimacy. If induction can confer justification without *any* awareness of its legitimacy then we can use bootstrapping to gain justification for induction's reliability even though we aren't aware of any evidence for this conclusion. Therefore, *some* awareness of the legitimacy of the application of the inductive rule should be required in order for it to transmit justification.

However, we shouldn't attempt to avoid the bootstrapping problem by requiring *doxastic* awareness of induction's legitimacy. The same bootstrapping worries arise for *any* basic non-deductive inference. Thus, such a response puts pressure on us to reject the very possibility of basic defeasible inferences. This leads to a far-reaching skepticism. If there are no basic defeasible inferences then a defeasible inference can only provide justification if we have antecedent justification for believing that it's reliable. This leads to two worries. First, it brings back the Lewis Carroll regress for defeasible inferences. Second, any justification for the reliability of a defeasible inference would itself need to appeal to defeasible evidence. *If* there are no basic defeasible sources then we couldn't ever get the necessary antecedent justification (this is precisely what gives rise to the Humean problem of induction). Thankfully, we are now in a position to see the middle road offered by my acquaintance theory of inferential justification.

According to IA, a defeasible inference is only basic *for me* if I'm acquainted with the partial entailment that holds between my premises and conclusion. Having



## IS JUSTIFICATION EASY OR IMPOSSIBLE?

accepted this moderate theory of inferential justification, if induction is a basic defeasible inference for S then the justification for believing induction is reliable conferred by bootstrapping cannot exceed the degree to which this conclusion is epistemically supported by the existence of the partial entailment... *of which the subject was already aware in the move from 1 to 2 above!* This vindicates the feeling that such reasoning doesn't generate anything new. Just as important, however, is that this result is achieved merely by requiring *acquaintance* with the partial entailment. Requiring acquaintance is not equivalent to requiring a justified *belief* that there is any connection between the beliefs about past observations and the belief about the future. Thus, we avoid the Carroll regress. Induction doesn't make beliefs rational because it's reliable or even because we are rational in believing that it's reliable. We are (defeasibly) rational in believing that induction is reliable because we are aware of induction making beliefs (defeasibly) rational.

It's important to stress, however, that the evidence for induction's reliability with which one must be acquainted is *defeasible*. If you continue to use induction to make predictions that tend to conflict with independent sources (such as perception, testimony, etc.) then this provides additional empirical evidence that defeats your justification for believing induction is reliable. When a set of evidence E makes epistemically probable *P* but doesn't entail *P*, it's possible to encounter additional evidence E\* such that E+E\* doesn't make epistemically probable *P*.

At this point it should be clear how to extend this solution to sources of foundational justification, and so I'll be brief. In cases of inferential justification, the reason or basis that provides S with justification consists of some of S's further beliefs. In cases of foundational justification, however, the reason or basis that provides S with justification is a non-doxastic state such as a perceptual experience, memorial experience, intuition, conceptual relation, or fact.

I prefer to identify facts as the non-doxastic states that provide foundational justification. Facts are general enough so as to incorporate all plausible candidates for evidence. There are facts about perceptual experiences, facts about memorial experiences, facts about intuitive experiences, facts about conceptual relations, etc. Intuitively, facts about my having perceptual experiences as of *P* (defeasibly) epistemically support believing *P* and can thereby provide foundational justification.

## IS JUSTIFICATION EASY OR IMPOSSIBLE?

However, if we insist that these perceptual experiences can provide justification without any awareness of the fact that they epistemically support the resulting belief then we invite problems of bootstrapping for foundational sources such as perception and memory. On the other hand, if we require a *doxastic* awareness of the connection between our perceptual or memorial evidence and our belief, then it turns out that the relevant belief isn't *foundationally* justified after all and we're led down our vicious infinite regress.

Yet again, our puzzle arises by overlooking a middle road. We shouldn't leap from the fact that requiring doxastic awareness of the epistemic credentials of our perceptual or memorial evidence is too strong to the claim that we shouldn't require *any* awareness. Requiring non-conceptual awareness avoids the problematic regress associated with adopting a higher-order doxastic requirement such as JR. Moreover, the epistemic support relation between our foundational evidence and our foundationally justified belief is itself (defeasible) evidence for the reliability of that non-doxastic evidence. Thus, requiring acquaintance with the epistemic support relation in order for our non-doxastic evidence to be epistemically efficacious ensures that we cannot use bootstrapping to produce justification for a claim for which we possess no evidence. Acquaintance with the fact, F, that our perceptual experiences as of *P* provide (defeasible) epistemic support for believing that *P* is in fact an awareness of something (F) that provides (defeasible) epistemic support for believing our perceptual experiences are reliable.<sup>34</sup>

---

<sup>34</sup> The bootstrapping problem is a *structural* problem. My solution is to suggest that in order for a basic foundational source to provide justification one must: (a) be acquainted with non-doxastic states that *do* stand in an epistemic support relation to the foundational belief; and (b) be acquainted with this epistemic support relation. This allows me to solve the structural problem by saying that *if* a non-doxastic state E supports *P* and one is acquainted with that support relation, *then* one gets justification for the premises of the bootstrapping reasoning but the resulting reasoning isn't problematic since we're actually aware of evidence for the conclusion. My solution to the structural problem only requires me to say that *if* a perceptual experience or seeming that *P* is a basic source of justification for believing *P*, then we must be acquainted with the experience and the support relation. Resolving the structural problem doesn't require defending the claim that a single perceptual experience as of *P* is in fact a basic source of justification that stands in epistemic support relation to *P*. One might suggest that it's only when one has a variety of experience as of *P* in multiple modalities (perhaps accompanied by memory experiences) that one's experience will epistemically support *P*. In which case, the basic source of foundational justification would be this larger set of non-doxastic experiences rather than an individual experience. If you adopt this latter view then my view claims that, in order for that set of experiences to provide foundational justification, the subject must be acquainted with the set of experiences and with the fact that this set supports believing *P*.

## IS JUSTIFICATION EASY OR IMPOSSIBLE?

Bootstrapping in order to move up a level to beliefs about the accuracy or reliability of your evidence cannot provide you with any more assurance concerning the likely truth of your belief than was present prior to moving up a level. When a fact  $f$  epistemically supports  $P$ ,  $P$  has the property of being probably true. When I'm acquainted with the support relation I have everything constitutive of  $P$ 's likely truth directly before my mind. My perspective on  $P$  provides me with an assurance regarding  $P$ 's truth. Bootstrapping to the further conclusion that forming my belief that  $P$  on the basis of the fact  $f$  was accurate (or even reliable) doesn't bring any information before my mind that improves my perspective on  $P$ 's truth.<sup>35</sup> All of this is achieved without accepting the principles that generated a vicious regress and lead to skepticism (i.e. JR and IR).

We are thus lead to an account of justification governed by the following two principles:

**FA:** S has foundational justification for believing that  $P$  only if: (i) S is acquainted with a fact or set of facts  $f$  that epistemically support  $P$  and (ii) S is acquainted with the fact that  $f$  epistemically supports  $P$ .

**IA:** S has justification for believing that  $Q$  on the basis of her justified belief(s) that  $P$  only if S is *non-conceptually* aware of, or directly acquainted with, the fact that  $P$  epistemically supports  $Q$ .

These principles follow Goldman's (1979) suggestion that we need to provide a recursive analysis of justification. FA provides the base clause and IA provides the recursive clause.<sup>36</sup>

---

<sup>35</sup> This isn't to say that moving up a level to a belief about a basic source's reliability *never* improves one's perspective on  $P$ 's truth. For instance, if we were to conclude that a basic source is reliable on the basis of an inductive argument that included independent checks on the source's output, then this would improve my perspective on  $P$ 's truth. In this case I would actually become aware of additional evidence that different source's outputs cohere with one another. This provides me with evidence for the basic source's reliability of which I was *not* even weakly aware when first forming my belief via that basic source. Moving up a level merely *via bootstrapping* doesn't improve one's perspective on the first-level belief. The evidence for the conclusion in the case of bootstrapping is just the epistemic support relation itself, of which one was already aware when one (justifiably) formed the initial belief via the basic source.

<sup>36</sup> FA and IA only state necessary conditions. In order to provide a full account of justification that gives necessary and sufficient conditions we would have to add a no defeater clause. Developing a no defeater clause that covers both foundational (the base clause) and inferential (the recursive clause) justification is a bit difficult since one cannot employ the notion of justification in one's account of defeaters, as this would introduce conceptual circularity into the analysis. However, characterizing the nature of defeat without appealing to the notion of justification isn't impossible. The general idea would just be that there is no larger set of facts  $f^*$  with which one is acquainted such that  $f^*$  doesn't epistemically support the belief.

## IS JUSTIFICATION EASY OR IMPOSSIBLE?

I'll refer to this analysis of justification (FA+IA) as the moderate acquaintance theory. It's moderate because it's weaker than some well-known acquaintance theories. For instance, what we might call the strong acquaintance theory of foundational justification requires acquaintance with a truth-maker for one's belief and with the truth-making relation itself.<sup>37</sup> Insofar as such theories only require *acquaintance* with the truth-making relation, rather than a meta-belief about the existence of such a relation, I think such views can avoid bootstrapping worries in much the same way that I have suggested.<sup>38</sup> However, requiring acquaintance with a truth-maker for foundational justification is tantamount to accepting that foundational justification requires *infallible* evidence, i.e. one couldn't have foundational justification for a false belief. My alternative moderate acquaintance theory rejects any such requirement on foundational justification. By merely requiring acquaintance with something that epistemically supports my belief (and acquaintance with the epistemic support relation itself), the moderate acquaintance theory makes room for the possibility of probabilistic foundations. This gives my moderate theory an advantage insofar as it allows for the possibility of foundationally justified perceptual and memorial beliefs.

My moderate acquaintance theory of inferential justification is also weaker than inferential internalism. Inferential internalism is a view that, in cases of inferential justification, requires that a subject have a justified meta-belief that the epistemic support relation is present.<sup>39</sup> On such views, *mere* acquaintance with the epistemic support relation is insufficient for inferential justification. As such, inferential internalism falls prey to the objections of over-intellectualization and Carroll style regress. By moving to my moderate acquaintance theory we can avoid these worries.

---

<sup>37</sup> See Fumerton (1995) for an example of a version of a strong acquaintance theory.

<sup>38</sup> Of course, there are stronger acquaintance theories that require a justified meta-belief about the truth-making relation even at the foundational level. See, for instance, Bonjour (2003) and Fales (1996). We might call these "super-strong" acquaintance theories. Such theories fall prey to the vicious regress objection. See Bergmann's (2006) discussion of both Bonjour and Fales' views for a particularly nice argument that these particular views fall prey to vicious regress.

<sup>39</sup> See Fumerton (1995, 2006a) and Leite (2008). Leite defends inferential internalism but makes no appeal to the traditional notion of acquaintance. He also attempts to defend an inferential internalism that is consistent with either a foundationalist or coherentist account of the structure of justification.

## IS JUSTIFICATION EASY OR IMPOSSIBLE?

However, my moderate acquaintance theory is also stronger than what we can refer to as weak acquaintance theories. According to weak acquaintance theories, justification for believing that *P* only requires that a subject be acquainted with something that in fact epistemically supports believing *P*; *S* needn't even be acquainted with the epistemic support relation itself. Weak theories clearly fall victim to the worries associated with bootstrapping. If perceptual or memorial experiences can justify the beliefs used in the bootstrapping without any awareness of the fact that these perceptual or memorial experiences epistemically support the relevant beliefs, then when we go through the bootstrapping reasoning we arrive at justification for a conclusion where we are not aware of anything supporting this conclusion.

My moderate acquaintance theory holds on to the motivation for traditional internalist views that appeal to direct acquaintance while also providing a nice compromise between strong and weak acquaintance theories. Moreover, such a view is in a unique position to avoid both vicious regress and implausible gains in justification via bootstrapping. Justification is neither easy nor impossible.

### BIBLIOGRAPHY

- Ballantyne, N. (2012). Acquaintance and Assurance. *Philosophical Studies*, 161 (3), 421-431.
- Bergmann, M. (2006). *Justification without Awareness*. Oxford: Oxford University Press.
- Bonjour, L. (2000). Toward a Defense of Empirical Foundationalism. In M. DePaul (Ed.), *Resurrecting Old-Fashioned Foundationalism* (pp. 21-38). Lanham: Rowman and Littlefield.
- Bonjour, L. (2003). A Version of Internalist Foundationalism. In L. Bonjour, & E. Sosa, *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues* (pp. 3-96). Malden, MA: Wiley-Blackwell.
- Goldman, A. (1979). What is Justified Belief? In *Justification and Knowledge* (George Pappas ed., pp. 1-23). Dordrecht, Holland: D. Reidel.
- Bonjour, L. (2010). *Epistemology: Classic Problems and Contemporary Responses* (2nd ed.). Lanham: Rowman and Littlefield.
- Goldman, A. (1986). *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Goldman, A. (1993). Epistemic Folways and Scientific Epistemology. In A. Goldman, *Liasons: Philosophy Meets the Cognitive and Social Sciences*. Cambridge: Cambridge University Press.

## IS JUSTIFICATION EASY OR IMPOSSIBLE?

- Goldman, A. (2008, April 21). *Reliabilism*. (E. Zalta, Ed.) Retrieved from The Stanford Encyclopedia of Philosophy:  
<http://plato.stanford.edu/entries/reliabilism/>
- Leite, A. (2008). Believing One's Reasons Are Good. *Synthese*, 161 (3), 419-441.
- Lewis, D. (1986). *On the Plurality of Worlds*. Malden, MA: Wiley-Blackwell.
- Lyons, J. (2009). *Perception and Basic Beliefs*. Oxford: Oxford University Press.
- Gupta, A. (2008). *Empiricism and Experience*. Oxford: Oxford University Press.
- Carroll, L. (1895). What the Tortoise Said to Achilles. *Mind*, 4(14), 278-280.
- Cohen, S. (1984). Justification and Truth. *Philosophical Studies*, 46(3), 279-295.
- Cohen, S. (2002). Basic Knowledge and the Problem of Easy Knowledge. *Philosophy and Phenomenological Research*, 65(2), 309-329.
- Cohen, S. (2005). Why Basic Knowledge is Easy Knowledge. *Philosophy and Phenomenological Research*, 70(2), 417-430.
- Fales, E. (1996). *A Defense of the Given*. Lanham: Rowman and Littlefield.
- Feldman, R. (2004). The Justification of Introspective Beliefs. In E. Conee and R. Feldman, *Evidentialism: Essays in Epistemology*. Oxford: Oxford University Press.
- Foley, R. (1983). Epistemic Conservatism. *Philosophical Studies*, 43(2), 165-182.
- Fumerton, R. (1985). *Epistemological and Metaphysical Problems of Perception*. Lincoln: University of Nebraska Press.
- Fumerton, R. (1995). *Metaepistemology and Skepticism*. Lanham: Rowman and Littlefield.
- Fumerton, R. (2004). Epistemic Probability. *Philosophical Issues*, 14(1), 149-164.
- Fumerton, R. (2005). Speckled Hens and Objects of Acquaintance. *Philosophical Perspectives*, 19 (1), 121-138.
- Fumerton, R. (2006a). *Epistemology*. Malden, MA: Blackwell.
- Fumerton, R. (2006b). Epistemic Internalism, Philosophical Assurance, and the Skeptical Predicament. In T. Crisp, M. Davidson, and D. Vander Laan (ed.), *Knowledge and Reality: Essays in Honor of Alvin Plantinga*. Springer.
- Fumerton, R. (2010). Poston on Similarity and Acquaintance. *Philosophical Studies*, 147 (3), 379-386.
- Greco, J. (1999). Agent Reliabilism. *Philosophical Perspectives*, 13, 273-296.
- Greco, J. (2010). *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge: Cambridge University Press.
- Hasan, A. (2013). Phenomenal Conservatism, Classical Foundationalism, and Internalist Justification. *Philosophical Studies*, 162, 119-141.
- Hawthorne, J. (2002). Deeply Contingent a Priori Knowledge. *Philosophy and Phenomenological Research*, 65(2), 247-269.
- Huemer, M. (2001). *Skepticism and the Veil of Perception*. Lanham: Rowman and Littlefield.
- Huemer, M. (2006). Phenomenal Conservatism and the Internalist Intuition. *American Philosophical Quarterly*, 43(2), 147-158.
- Huemer, M. (2007). Compassionate Phenomenal Conservatism. *Philosophy and Phenomenological Research*, 74(1), 30-55.
- Keynes, J. M. (1921). *A Treatise on Probability*. London: Macmillan.
- McCain, K. (2008). The Virtues of Epistemic Conservatism. *Synthese*, 164(2), 185-200.

IS JUSTIFICATION EASY OR IMPOSSIBLE?

- McCain, K. (2012). Against Hanna on Phenomenal Conservatism. *Acta Analytica*, 27(1), 45-54.
- McGrew, T., & McGrew, L. (2007). *Internalism and Epistemology: The Architecture of Reason*. New York: Routledge.
- Moser, P. (1989). *Knowledge and Evidence*. Cambridge: Cambridge University Press.
- Poston, T. (2007). Acquaintance and the Problem of the Speckled Hen. *Philosophical Studies*, 132 (2): 331-346.
- Poston, T. (2010). Similarity and Acquaintance: A Dilemma. *Philosophical Studies*, 147 (3), 369-378.
- Pryor, J. (2000). The Skeptic and the Dogmatist. *Nous*, 34(4), 517-549.
- Russell, B. ([1912] 1997). *The Problems of Philosophy*. Oxford: Oxford University Press.
- Schroeder, M. (2008). Having Reasons. *Philosophical Studies*, 139 (1), 57-71.
- Skene, M. (2013). Seemings and the Possibility of Epistemic Justification. *Philosophical Studies*, 163(2), 539-559.
- Sosa, E. (2009). *Reflective Knowledge: Apt Belief and Reflective Knowledge*. Oxford: Oxford University Press.
- Steup, M. (2004). Internalist Reliabilism. *Philosophical Issues*, 14(1), 403-425.
- Titelbaum, M. Tell Me You Love Me: Bootstrapping, Externalism, and No-Lose Epistemology. *Philosophical Studies*, 149 (1), 119-134.
- Tucker, C. (2010). Why Open-Minded People Should Endorse Dogmatism. *Philosophical Perspectives*, 24(1), 529-545.
- Tucker, C. (Ed.). (2013). *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*. Oxford: Oxford University Press.
- Turri, J. (2011). Contingent A Priori Knowledge. *Philosophy and Phenomenological Research*, 83 (2), 327-344.
- Vogel, J. (2008). Epistemic Bootstrapping. *Journal of Philosophy*, 105 (9), 518-539.
- Weisberg, J. (2010). Bootstrapping in General. *Philosophy and Phenomenological Research*, 81 (3), 525-548.