

There is no easy bootstrapping problem

Peter Kung, Pomona College

Masahiro Yamada, Claremont Graduate University

May 25, 2009

Under review at *Australasian Journal of Philosophy*

Abstract

Some epistemological views accept that certain belief forming mechanisms can deliver justified beliefs even if the epistemic agent lacks justification to believe that the mechanism is reliable. The bootstrapping objection charges that such views make it too easy to acquire justified belief that these belief-forming mechanisms are reliable by inductive reasoning. We offer a novel way to diffuse this charge. An independently plausible constraint on inductive reasoning is also available to those epistemological views targeted by bootstrapping argument. The constraint arises from consideration of clear cases of bad inductive reasoning and does not require commitment to potentially controversial epistemological views.

Keywords: epistemology; foundationalism; skepticism; bootstrapping; easy knowledge

Suppose you think skepticism is false. How do you go about answering the skeptical challenge? A foundationalist theory might do the trick. Knowledge rests on a foundation of noninferentially justified beliefs that are *prima facie* justified, or entitled, such that this justification does not derive from the justification for any other beliefs. Foundationalist theories vary along numerous dimensions; different foundationalist theories will have different kinds of beliefs in the foundation. One attractive version holds that foundational beliefs do not depend on having justified beliefs about the reliability of

belief-forming methods, so it is possible to form a belief in a certain way, e.g. perceptually, without having antecedent justification for believing that the way of forming belief in question, e.g. perception, is reliable. Let us say that a belief is *basic* if it is foundational and its justification does not require antecedent, independent justification that the way it is formed is reliable. If these basic beliefs are sufficiently rich they provide precisely the fulcrum the foundationalist needs to dislodge the skeptic.¹ Let us call a view that accepts the existence of basic beliefs so characterized a *basic belief theory*.

There is a worry. Yes, basic beliefs are the foundationalist's answer to the skeptical challenge. But is it too easy an answer? Do basic beliefs license *too much* knowledge, knowledge that is intuitively unearned? Stewart Cohen (2002: §III) and Jonathan Vogel (2000) charge that it does. In particular, they argue that if one accepts basic beliefs, it is too easy to bootstrap our way to knowledge of the reliability of our cognitive faculties: the bootstrapping problem of easy knowledge.²

The problem in a nutshell is that subjects can use their pool of basic beliefs as the basis for an inductive inference to the conclusion that the mechanism responsible for producing the basic beliefs is reliable. Once it is clear how the subject gathers her pool of basic beliefs the absurdity of the inductive inference is manifest: what it means, in effect, is that just by taking a few glances around a subject can know that her perceptual system

¹ Famously Moore (1939) holds this position. Recent advocates include Alston (1989), Audi (1993), and Pryor (2000, 2004).

² The problem of easy knowledge, as Cohen labels it, has a second component: the closure problem (Cohen 2002: §II). That problem is not our focus.

is reliable. We agree that this is absurd. *If* it is possible to show that basic belief theory cannot explain why bootstrapping is unacceptable, one could launch a *reductio ad absurdum* against basic belief theory: it has absurd consequences, so it must be false.

We argue in this paper that this *reductio* fails. The fault lies not with basic beliefs but rather with a flawed inductive inference. The *reductio* argument seems appealing only insofar as we fail to notice (or just ignore) some important restrictions on cogent inductive reasoning. As we will argue, everybody needs to accept some constraints on cogent induction to avoid problems akin to easy bootstrapping that have nothing to do with basic belief theory. We propose a constraint that does the job. Moreover, a basic belief theorist can appeal to this constraint to explain what is wrong with bootstrapping. So there is no bootstrapping problem of easy knowledge.

Of course, just because this particular *reductio* against basic belief theory fails does not mean that there are no other objections to basic belief theories.³ There are, and our points in this paper do not address those objections. Our point is simply that even on the assumption there are basic beliefs — beliefs justified without the need for antecedent independent justification for the reliability of the relevant belief forming processes — one can explain what is wrong with bootstrapping without appeal to ad hoc principles.

I Preliminaries

The bootstrapping problem of easy knowledge can be adapted to different basic belief theories. As formulated by Vogel, the problem is one about knowledge, and he directs

³ Thanks to Stewart Cohen for emphasizing this point.

his argument against reliabilism. According to reliabilism about perceptual beliefs, a perceptual belief is knowledge just in case it has been formed in a reliable way.⁴ The pool of basic beliefs that license bootstrapping would then be drawn from this perceptual knowledge. As Cohen notes, Vogel's argument can easily be generalized to any theories that accept basic beliefs since the crucial aspect of reliabilism for the purposes of his argument is that reliabilism does not require antecedent knowledge that perception is reliable — any theory that does not require antecedent knowledge that perception is reliable and accepts induction as a way of extending knowledge will succumb to Vogel's arguments. Vogel and Cohen present the problem as one about the structure of knowledge but the problem can straightforwardly be recast as one about justification.⁵ For instance, dogmatism is a theory of justification, rather than knowledge, which holds that a belief formed in response to a perceptual state is *prima facie* justified in the absence of evidence for the *unreliability* of perception on that occasion.⁶ On this view these perceptually justified beliefs constitute the pool of basic beliefs and lead to an unacceptable increase in your justification to believe that perception is reliable. The problem is not simply that you can become justified in believing the conclusion; rather the problem is that bootstrapping leads to *any* increase in your justification at all, regardless of whether the increase is sufficient for justified belief. In this paper we focus on the more general bootstrapping problem of easy justification, but all points should

⁴ The locus classicus for reliabilism is Goldman (1979).

⁵ Wright advances a similar line of objection to basic justification; see e.g., Wright (2003).

⁶ For a statement and defense of dogmatism see Pryor (2000) and (2004). See also Davies' (2004) discussion of negative entitlement.

apply equally to easy knowledge. For brevity we'll simply refer to "the easy bootstrapping problem."

It is natural to understand the easy bootstrapping problem as a challenge to foundationalism in general. But as Cohen notes (2002, 310n4) there is nothing in the idea of foundationalism that prevents beliefs about the reliability of one's cognitive faculties being part of the foundationally justified beliefs. Such foundationalist views do not succumb to the easy bootstrapping problem. For this reason, we follow Cohen in describing the easy bootstrapping problem as one for basic beliefs rather than for foundationalism.

Finally, to make the discussion more concrete, throughout the paper we focus on perception. Our example of a basic belief theory is one that takes some simple *perceptual* beliefs as basic beliefs. However, nothing that we or Cohen or Vogel say is specific to such views.

2 Stating the Problem

What is the problem of easy bootstrapping? It comes in two steps and it is easiest to see in an example. Suppose Kofi is shown a series of colored panels. For each panel, he forms two beliefs. He notes and thereby forms a belief about the panel's color. Then he carefully introspects his perceptual experience and forms a belief about his experience.

Let us represent the sequence in the following way:

	Experience as of	Perceptual belief	Introspective belief
1.	Red	panel is red	I'm having an experience as of a red panel
2.	Green	panel is green	I'm having an experience as of a green panel
...
10.	Blue	panel is blue	I'm having an experience as of a blue panel

Assume that the introspective beliefs are justified. If each of the perceptual beliefs is also justified, as a basic belief theory suggests,⁷ then each time Kofi is justified in thinking: "The panel is color C, and I'm having an experience that it is color C." He continues in this fashion until he has looked at 1000 panels. He now goes through his data and realizes that on the n occasions when the panel was green he also had experience that the panel is green. Kofi performs a simple enumerative induction:

Of n panels that appeared to be green, all were green and none were not green.

Therefore, my visual system is reliable when it comes to the color green.

This inductive inference might look unobjectionable provided n is high enough. Even when n is low the inductive base might look to provide *some* evidence for the conclusion. Kofi repeats this with other colors that appeared in his trials and concludes that his perceptual system is reliable with respect to the color red, with respect to the color yellow, and so on. Once he is satisfied that he is reliable with respect to enough colors, he concludes that his color vision is reliable.

Surely, it cannot be so easy for Kofi to generate *any* evidence for the reliability of his visual system. He cannot even gain any evidence that his color vision is reliable as far as

⁷ As we will emphasize for the last time, not every basic belief theory suggests this. We use basic belief theories that accept perceptual beliefs as basic as an example. For other basic belief theories, the example can be modified accordingly.

the color green is concerned. He has merely taken a few glances around; he hasn't actually done any investigation. This is blatant bootstrapping. Basic belief theories seem unable to locate a flaw in Kofi's reasoning. This is the problem of easy bootstrapping.

So what *is* wrong with Kofi's procedure?⁸ It is clear that there is a problem with the enumerative induction that we displayed above. Kofi does not gain any justification for believing his vision reliable with respect to colors in general because he is not even gaining any justification for believing that it is reliable with respect to a single color, like green. (Henceforth we'll use Kofi's reasoning about his green color vision as our primary bootstrapping example.) But what exactly is wrong with the enumerative induction? When we examine Kofi's reasoning leading up to the inductive generalization that his green color vision is reliable, there are several places where we might try to block his reasoning. One might insist that justification of his perceptual beliefs — e.g. that the panel in front of him is green — depend on an antecedent justification for belief in reliability of his color vision. On this explanation, the problem with Kofi is that his reasoning is circular. This explanation is, of course, not available to basic belief theory. It is the defining feature of basic belief theory that perceptual beliefs

⁸ Alston (1986) holds that there is in fact nothing wrong with it. More recently, Davies (2004) and Markie (2005) have also suggested that the reasoning itself is unproblematic. According to Markie, the apparent problem arises from the fact that the reasoning could not persuade someone who is skeptical of the conclusion; that is, the problem is one of dialectical inefficacy and not of bad reasoning. Alston gives a similar diagnosis. However, this leaves the puzzle why such an argument is dialectically ineffective. Surely, the central reason is that there is indeed something wrong with Kofi's reasoning and there needs to be an account of what is wrong with it. As Cohen (2005) responds to Markie, the reasoning appears to be a bad piece of reasoning whether or not it is presented in a dialectical context in which the conclusion is in question. A similar objection can be raised against Davies' suggestion that the reasoning is dialectically ineffective in *settling the question* — see Davies §VII.

do not depend for their justification on justification to believe perceptual faculties are reliable. One might hold that even though the perceptual beliefs and introspective beliefs are individually justified, Kofi is not justified in believing the conjunction that the panel is green *and* that he has an experience as of a green panel. This would be to block conjunction introduction: somehow, conjunction introduction here does not extend Kofi's justification. While this position is consistent with basic belief theory, we do not recommend this as a response. Yet another possibility is to hold that while Kofi is justified in believing that he got it right on any given occasion, he is somehow not justified in believing in his track record that experience matched the fact that the panel is green on *n* occasions and didn't fail to match on a single occasion. This move also blocks conjunction introduction and is consistent with basic belief theory but, again, we do not recommend this move either, for reasons that we will explain. Finally, one could argue that even though Kofi is justified in believing in his track record, he cannot use this as a basis for cogent inductive generalization. This is the position we recommend the basic belief theorist take.

In arguing that it is the final step of inductive generalization that is wrong with Kofi, we are arguing that the following is false:

(Crude Induction) If *S* is justified in believing that of the examined *n* *A*s,
all are *B*, then this increases *S*'s justification for believing that all *A*s are *B*.

So for Kofi we are denying: If Kofi is justified in believing that of the examined *n* panels that appeared to have color *C* are color *C*, then this increases Kofi's justification for believing that all panels that appear to be color *C* are color *C* (where *C* ranges over color

properties). First thing to note is that everybody needs to reject Crude Induction for if it were true, there would be nothing wrong with bootstrapping on *any* non-skeptical epistemology. Take some non-skeptical non-basic belief theory of perceptual justification according to which to have justified perceptual beliefs one needs antecedent reason to believe that one's perceptual system is reliable. If skepticism is false, then we have such antecedent reason, and so we have justified perceptual beliefs. We can then justifiably compile a positive track record and perform the inductive reasoning just like Kofi does. If Crude Induction were true, this would increase our justification to believe in the reliability of perception. It better not be so easy to confirm our own reliability. So rejecting Crude Induction is not some peculiar commitment that basic belief theory needs to make.

Second, while the above point could be accommodated by adding to Crude Induction that the justification for the track record must not depend on justification for the conclusion, there is another obvious problem with Crude Induction that such a clause could not handle. For instance, if Crude Induction were true, there would be nothing wrong with blatant sampling bias. Suppose we go through a whole bunch of objects. A fully automated machine discards all but the black ones by using a scanner that is sensitive only to color. We then examine the leftover objects one by one. Some of them are ravens. We notice that of the examined ravens all are black. We are justified in believing this. This, however, doesn't increase our justification in believing that all ravens are black and the problem is not one of circularity: even if we were convinced that there are pink ravens out there, we would still be justified in believing that all the

examined ravens are black. But Crude Induction delivers the absurd result that we do increase our justification for believing that all ravens are black. There is ample evidence that we need to constrain Crude Induction regardless of whether we accept basic belief theory.

In this paper we will focus on examples that exhibit a yet further flaw. The flaw is distinct from circularity or sampling bias and it appears in a range of cases where basic belief theory is *not* assumed. Hence there are independent reasons for regarding the flaw we point out as the source of a constraint on Crude Induction. Crucially for our purposes, it turns out that the flaw is something that basic belief theory can appeal to in explaining what is wrong with Kofi's reasoning. Moreover, the examples will show why it would not be a good idea to block Kofi's reasoning by trying to block some earlier step in his reasoning.

In identifying a flaw in Kofi's reasoning that basic belief theory can appeal to, we are not claiming that it is the *only* problem with Kofi. A non-basic belief theorist might insist that Kofi's reasoning is circular since the justification for his premises depends on justification for his conclusion. We are not arguing that Kofi's reasoning should not be counted as circular even by a non-basic belief theorist. Our recommendation to the non-basic belief theorist is to acknowledge at least *two* problems with Kofi: circularity and the one we point out. The latter is a problem shared with certain flawed but non-circular reasonings and can be acknowledged by basic belief theory. Whether Kofi has one, two or even more problems is not something we aspire to settle in this paper. Our ambition is only to show that there is at least one problem that even basic belief theory can point

out. This suffices to show that the easy bootstrapping *reductio* against basic belief theory that we are interested in fails.

3 Bootstrapping without basic beliefs

Here are two cases that share important features with Kofi yet where circularity is not the problem. First, consider Thant, who wants to sell his house in a volatile crashing market. Suppose a trustworthy friend tips Thant to a website, *insiderbroker.com*, that lists current buyers, the type of home they are interested in, and their very highest buying price. The friend emphasizes that he can vouch for *today's* information only. Resting on his friend's testimonial justification and the website, Thant forms justified beliefs about today's current buyers, what type of home they want, and the highest prices they are willing to pay for it. It's a goldmine of information; but rather than sell today to the highest bidder, Thant wants to monitor buyers' highest prices for a few days, to see whether buying prices are trending up or down. Worried that the site won't continue to provide good information tomorrow, Thant compiles an inductive sample.

	belief about website	belief about real estate market
1.	the website says that the Hammarskjölds are willing to pay \$2.2 million for a 4BR ranch with pool	the Hammarskjölds are willing to pay \$2.2 million for a 4BR ranch with pool
2.	the website says that the Waldheims want a 20-acre castle with dungeon for at most \$17.8 million	the Waldheims want a 20-acre castle with dungeon for at most \$17.8 million
...
100.	The website says the de Cuéllar's need a 2BR loft apartment for not more than \$925,000	the de Cuéllar's need a 2BR loft apartment for not more than \$925,000

Thant reads the website to form his belief about what the website says and bases his belief about the buyer's market by accepting what the website says at face value. Each time he of course concludes the website gets it right: "the website says that the Hammarskjölds are willing to pay \$2.2 million for a 4BR ranch with pool, and the Hammarskjölds are in fact willing to pay \$2.2 million for a 4BR ranch with pool — the website got that one right." He uses this belief about the website's track record to bootstrap his way up to the conclusion that the website will be reliable *tomorrow*. Clearly this is illegitimate; he can't verify the website's accuracy today by checking it against *itself* (he's not even double-checking), nor can he justifiably increase his confidence that the website will continue to be accurate in the future. Thant's reasoning has the very same bootstrapping problem as Kofi's.

Thant's problem is certainly not one of circularity: the conclusion is that the website is reliable *tomorrow* and that is not one of the premises, nor does justification for his premises depend on justification for the conclusion. Thant depends only on justification

for the belief that the website is reliable *today*. In fact, suppose Thant's friend warned him that the site's information would *not* be good tomorrow. Thant would be justified in believing that the website is reliable today while also being justified in believing that the website is *unreliable* tomorrow. This shows that the justification Thant has for his various beliefs today do not depend on justification for his conclusion. Moreover, while it is possible for Thant to gather evidence that his friend is *wrong* about tomorrow, he cannot do so by checking the website against itself. So there must be some problem with Thant's procedure that isn't circularity.

More generally, this case illustrates that every theory must block bootstrapping to the conclusion that our measuring equipment is reliable under circumstances beyond those for which we already possess reason for believing the equipment is good. And any theory that does not start with foundational material sufficient for *prima facie* justification that our perceptual faculties are reliable under all circumstances must find a way to block bootstrapping to the conclusion that our perceptual faculties are reliable under all circumstances.

If one were a skeptic about inductive reasoning, one could complain that Thant's procedure displays sampling bias: he checks only today to draw conclusions about tomorrow. How is that any better than examining only people in Minneapolis to draw conclusions about people in Azerbaijan (say, that most Azerbaijanis are native speakers of English)? We take it that is not an interesting way to diffuse the easy bootstrapping problem. If one does not want to be a skeptic about inductive reasoning, one better accept that one can gather evidence today to draw conclusions about tomorrow. Thant's

problem is not that he is using data from the past and present to draw conclusions about the future.

Now consider a second case.⁹ Boutros never did very well in chemistry but does remember a few things:

(RA) If a liquid turns a piece of litmus paper red, it is acidic.

(RC) If a liquid turns a piece of litmus paper red, it corrodes iron.

Boutros takes a liquid, exposes the litmus paper to it, observes the paper turn red, and deduces (comes to believe) that it is acid, and that it corrodes iron. Given that Boutros is justified in believing RA, RC and that the litmus paper turned red, he is certainly justified in believing the conjunction that the liquid is both an acid and corrosive to iron.

Suppose he wants to find out whether all acids corrode iron. RA and RC do not entail that all acids corrode iron, so as far as Boutros is concerned it is an open question whether all acids corrode iron. He could find out by doing some research: collect a bunch of liquids and see if the acids among them corrode iron. If they all do, he can generalize to the claim that all acids corrode iron. But there is a small problem. RA only gives him a sufficiency test for acidity. If he relies on RA to determine whether a liquid is an acid, he might, for all he knows, miss out on all kinds of acids that do not turn litmus paper red. Luckily, his neighbor is a chemical engineer and tells Boutros that

(AB) A liquid is acid only if it bubbles up when baking soda is added to it.

Boutros can now test whether all acids turn litmus paper red. He collects a bunch of

⁹ We consider this case in Kung & Yamada (forthcoming).

liquids and for those that don't turn litmus paper red, he adds baking soda. None of them bubble up so none of them satisfy the necessary condition for being acids. He concludes

(RA+) A liquid turns a piece of litmus paper red if and only if it is acidic.

Boutros now goes back to his original question. Do all acids corrode iron? He collects a large sample of liquids making sure there is a wide variety among them. Boutros takes a liquid, exposes the litmus paper to it, observes the paper turn red, and deduces (comes to believe) that it is acid (via RA), *and* deduces that it corrodes iron (via RC). He repeats this with several different liquids.

	Litmus paper	A	B
1.	Red	liquid is acidic	liquid corrodes iron
2.	Red	liquid is acidic	liquid corrodes iron
...
100.	Red	liquid is acidic	liquid corrodes iron

In each case, Boutros comes to believe that the liquid is acidic and that it corrodes by *deducing* this from the color of the litmus paper. He observes the color of the litmus paper directly. He believes RA and RC via testimony and memory. Assuming that perception, testimony, memory, and deduction confer justification, Boutros is justified in thinking: "This is an acidic liquid that corrodes iron; that is one occasion where acidity accompanies corrosiveness." After numerous trials, Boutros is justified in thinking: "There were one hundred acidic liquids that corrode. Since RA+ is true, there was no acidic liquid that I failed to check whether it corrodes iron. So, there were one hundred acidic liquids that corrode and none that do not." Can Boutros reason in the following

way?

Of 100 examined acidic liquids, all of them corrode iron.

Therefore, all acidic liquids corrode.

There is obviously something fishy about this inductive inference. Certainly Boutros does not *increase* his justification for thinking that all acidic liquids corrode.

RA+ and RC entail the conclusion that all acids corrode iron. Let us grant that Boutros has justification to believe the conclusion once he gains justification to believe RA+. The problem is that even if he is justified in believing that all acids corrode, his little “research” does nothing to *increase* his justification for the conclusion. To be clear, Boutros could easily strengthen his justification for believing that all acids corrode iron: he could throw in a piece of iron each time a liquid turns out to be acid and see what happens. If the iron is corroded each time, that would strengthen his justification for believing that all acids corrode iron. So the problem is not due to the fact, if it is a fact, that Boutros has justification for his conclusion even before his little “research.”¹⁰ The problem is also not circularity since although Boutros is justified in believing RA+ and RC, and hence has justification for believing the conclusion, he is not *depending* on the justification in drawing his conclusion. His reasoning relies on justification provided by

¹⁰ Boutros might fail to see the entailment from RA+ and RC to the conclusion that all acids corrode iron. Depending on the mechanism behind the failure, it seems to us that it could be that Boutros doesn't have justification for believing that all acids corrode iron. After all, we are not justified in all deductive consequences of our beliefs. We could also complicate the entailment (e.g., by adding a few more starting premises, or changing a modus ponens into a modus tollens) to make it more plausible that Boutros would not see it. We constructed the case this way because our point is that *even if* we grant that he has antecedent justification for the conclusion, his reasoning does not *depend* on that justification.

observation, on justification for RC which is provided for by testimony, and on justification for RA+ which in turn depends on observation and justification for RA and AB both of which are provided by testimony; none of this depends on justification for believing that all acids corrode iron.¹¹ Basic belief theory or not, there must be some problem with Boutros's procedure and that problem is not a form of circularity. And the problem is also not a form of sampling bias: he has justification to believe that he has checked all the acids in his sample and the problem has nothing to do with how large and varied his collection of liquids is.

We will now turn to our diagnosis of the problem displayed by both Thant and Boutros.

4 Rigging

One striking feature of cases like Thant and Boutros is that they are proceeding in such a way that the result they get is a forgone matter. *Of course* Thant is going to conclude that the website is going to be reliable tomorrow if he proceeds in the way he does, and *of course* Boutros is going to conclude that all acids corrode iron. This is obvious even before they embark on their research and this is why their reasoning is flawed. We think that this intuition, when properly cashed out, explains what is wrong with easy bootstrapping. But we need to proceed carefully in spelling out the intuition. We will begin by examining several attractive proposals, and show why each turns out to be

¹¹ What we mean by 'dependence' is very close to what Pryor (2004) classifies as Type 5 dependence. The justification for S's belief that p *depends* on D just in case D is one of the factors that *make* S justified in believing that p. D is often another justified belief but it could also be other things such as perceptual and memory states.

inadequate. This will help us work towards our own account.

In the Boutros case, there are in fact no acids that do not corrode iron. That is a physical necessity so in this sense he is bound to conclude that of all the examined acids, all of them corrode iron. But this cannot be the reason why there is something wrong with his reasoning. If it were, there would be no cogent way of inductively arriving at the conclusion that all acids corrode iron, or any other nomologically necessary truths.

But Boutros is not merely bound to conclude that all the examined acids corrode iron, he is in a position to *know* this once he knows RC and RA+. Maybe this is the problem? Perhaps his reasoning is flawed because he knows that all acids corrode iron even *before* carrying out the procedure designed to answer whether all acids corrode iron. More generally, the idea would be that if one already knows that *p*, then one cannot cogently reach the result that *p* in an inductive way. Even if the procedure does not suffer from circularity, maybe there is something wrong with knowing the answer in advance. But this cannot be right.

So long as knowledge does not require absolute certainty, it can make sense for Boutros to double check whether acids *really* corrode iron even if he knows that all acids corrode iron. Boutros would be *increasing* his justification to believe that all acids corrode iron. But if Boutros's knowing that all acids corrode iron is what makes his reasoning bad, there could be no cogent way of double checking whether all acids corrode iron once one knows that acids corrode iron. This is simply not true. For instance, instead of deducing that a given acid corrodes iron via RC, he could check by throwing in a piece of iron. This would be a cogent way of double checking whether an

acid corrodes iron and thereby increasing his justification for believing it.

Moreover, the diagnosis does not apply to Thant's case. Thant concludes that the website will be reliable tomorrow but he does not know that the site will be reliable tomorrow before going through his inductive reasoning (and does not know it afterwards, either). So whatever the problem is with Thant, it cannot be due to his knowing the correct answer in advance. We would have to give a different diagnosis of what is wrong in Thant's case. But Boutros and Thant do seem to share a problem and if there is a way of diagnosing the cases that treats them in a uniform fashion, that is to be preferred.

Perhaps the problem isn't that Thant knows that the website will be reliable tomorrow in advance, but that he knows in advance what he will conclude about the track record of the website: he will conclude that the website got it right every single time he checked. The same seems true of Boutros and Trygve. Boutros knows in advance that he will conclude that all the acids he checked corrode iron and Trygve knows that he will conclude that his advisors agreed with him every time he checked with them. While attractive, this cannot be right, either. Consider the Kofi case. He does not know in advance that he will conclude that every time there was a green panel that it looked green to him. After all, for all he knows in advance, no green panels might turn up. Nevertheless, the way he reaches the conclusion that his green color vision is reliable is unacceptable. Insofar as there clearly seems to be a common problem with Kofi, Thant, Boutros and Trygve, the common problem cannot be antecedent knowledge of what one is going to believe about the track record. What is the problem they have in common?

We are trying to spell out the intuition that what is wrong with Thant and Boutros is that they proceed in ways that make it obvious what they are going to conclude even before going through the reasoning. The first proposal was that the problem is that it is obvious that one concludes that p because p is some kind of necessary truth. This does not work. The second proposal was that the problem is that one is engaged in reasoning to determine whether p while already knowing that p . This does not work, either. What does work?

The third proposal was that the problem lies with the agent's knowing what she will conclude about the track record. This, too, is not quite right. But it is close.

We propose that the problem with easy bootstrapping is *rigging*. Rigging is a feature of *procedures*, rather than of the subject (*a fortiori* rigging isn't a matter of what the subject knows in advance). In our investigations we want to use procedures that give us the right answer. We wonder, "is p true?" Our procedure should enable us to recognize whether p by answering " p " if p turns out to be true and "not- p " if p turns out to be false.¹² Let us say that a procedure is rigged if it does not enable us to recognize whether p . Notice that even if the procedure is rigged, one might have other means of recognizing whether p . But the procedure is rigged insofar as *it* does not enable us to recognize whether p . For ease of exposition, we will say that a procedure cannot recognize p just in case *it* does not enable the agent using it to recognize whether p (even if the agent has other means of recognizing p at her disposal). For instance Trygve leads a global enterprise and is responsible for synthesizing vast amounts of complex data and

¹² Or remain silent if p is not true.

rendering important decisions. Unfortunately Trygve handpicks his inner circle of advisors principally based on whether they are sycophantic yes men and women. They always agree with him and tell him what he wants to hear. When Trygve solicits his advisor's opinions, they will always tell him that he made the right decision. So the procedure of asking his advisors cannot recognize any bad decisions that he might be making. Their advice is rigged.

It is pretty clear that generating inductive evidence using a rigged procedure is problematic. Trygve accumulates track record data about his own business judgment by checking with his sycophantic advisors: "I think we should acquire at \$2.50/share. My advisors agree. Another right for me." Trygve cannot use this track record data to increase his justification for believing that he's got sound business judgment. We will be cautious and say that you cannot use an *obviously* rigged procedure to generate inductive evidence.

Some procedure are rigged but are not obviously rigged. For instance, one might be relying on a device to detect whether or not a liquid is an acid which is in fact malfunctioning and always tells one that the liquid is not an acid independently of the liquid's acidity. A procedure for detecting whether a liquid is acid that depends on this device is rigged. But that it is rigged need not be obvious: one could have reason to believe that the device works and have no readily available means of checking the device. Merely understanding the procedure itself does not reveal that it is rigged. Many rigged procedures are like this. However, there are other rigged procedures that can be seen to be rigged if one simply understands the procedure. For instance, Boutros's

procedure above is obviously rigged. One does not need any knowledge of chemistry to see that Boutros's procedure cannot produce the verdict that a liquid is an acid but non-corrosive to iron even if the liquid in question is in fact a non-corrosive kind of acid. It is clearly bad to use a procedure which is obviously incapable of delivering the verdict that something is an A but not a B even if it is, since in such cases it is *obvious* that the procedure will not recognize counterexamples to the claim that all As are B by using the procedure even if there are some counterexamples.

When a procedure is obviously rigged, the agent does know certain things about the verdicts the procedure will or will not deliver. For instance, Boutros knows that his procedure will not deliver the verdict that a given liquid is an acid but not corrosive to acid. But he does *not* know that the procedure on any given occasion will deliver the verdict that it is both acidic and corrosive to iron: the procedure can deliver the verdict that the liquid is not an acid. So even if a procedure is obviously rigged, the agent using it need not know what the conclusion of the procedure will be on any given occasion. This is a crucial point that distinguishes our proposal from the third proposal above that the problem with bootstrapping is that the agent knows in advance what she will conclude about the track record.

How does the fact that their procedures are rigged interfere with the cogency of Boutros's and Thant's reasoning? One might suspect that the fact that one cannot recognize a case in which an A is not a B shows that one cannot gain justification for believing that an A is B. In fact, the Trygve case might seem to support this. On one way of telling Trygve's story, it should be obvious to him that, given how he has chosen his

advisors, agreement from his advisors is *no* evidence that he got it right on a single instance. It isn't at all clear that Trygve has the usual testimonial justification to believe what his advisors say, given he knows (or ought to know) that they are responding to him rather than the *truth*. And should we not say the same about Boutros and Thant? Maybe Boutros is not even justified in believing that any particular liquid is both an acid and corrosive to iron. That is, perhaps what explains what is wrong with cases like Boutros and Thant is the following condition:

(Anti Rigging Strong) If it ought to be obvious to S that her procedure cannot recognize that an A that is not B even if there are any such cases, then S's justification for believing that an A is B does not increase.

Anti Rigging Strong would explain what is wrong with Trygve, Thant and Boutros as follows: It would say that in each of these cases, the subjects have not gained justification to believe that a single A is a B, let alone any justification for the track record that many As are B, and hence they lack any justification for their basis of inductive generalization. This is a tempting thought but Anti Rigging Strong must be rejected.

Suppose you notice that someone donated a large amount of money to the Democratic Party. You would have evidence that the person is both wealthy and sympathizes with Democratic Party policies. However, your procedure for determining whether this conjunction is true is obviously rigged. You have only a single piece of evidence and use it as basis for supporting each of the conjuncts. So it is obvious you will not end up believing one of the conjuncts true but the other false. If Anti Rigging Strong were correct, the fact that someone donated a large amount of money to the

Democratic Party could not be used to raise the level of justification for believing that the person is both wealthy and sympathetic with Democratic Party policies. Surely, this is incorrect. Similarly, when someone who believes RA and RC finds a liquid turning litmus paper red, surely this gives him justification to believe that the liquid is acid and corrosive to iron: a chemist doesn't need to conduct two distinct tests to come to believe justifiedly that a liquid has two properties, e.g. that it is both an acid and corrosive to iron. These considerations show that Anti Rigging Strong is unacceptable. Therefore there is no reason to deny the otherwise plausible thought that Boutros and Thant are justified in single instances.¹³

Perhaps one would want to raise questions about the track record beliefs. For instance, one might think that even though Boutros gains justification to believe that any given liquid is both acid and corrosive, he somehow cannot gain justification for believing that there were n liquids that were both acid and corrosive and no liquids that were acid but not corrosive for n much larger than 1. That is, one might think the following is true:

(Anti Rigging Medium) If it ought to be obvious to S that her procedure cannot recognize that an A that is not B even if there are any such cases, then S 's justification for believing that n A s are B , where $n \gg 1$ (or some

¹³ Recall that the bootstrapping argument takes the form of a *reductio*. Objectors assume basic belief theory and derive an absurd bootstrapping result; that absurd result must be derived using plausible general principles. We have shown that Anti Rigging Strong is not a plausible general principle. One might still insist that basic belief theory delivers an absurd verdict on Kofi in a single case: Kofi simply *cannot* be justified in believing that even one panel was really the color it appeared to be. This reaction amounts to insisting on the closure problem of easy knowledge, which we do not address in this paper. See footnote 2.

other threshold), does not increase.

This would locate the problem with Boutros and Thant at a later point than Anti Rigging Strong. The individual conjunctions making up the track records are justified but somehow one cannot turn those into a justified belief in the track record. But Anti Rigging Medium has problems as well. Anti Rigging Medium blocks the application of conjunction introduction to get to a justified belief in the track record. While there can be reasons to block conjunction introduction in certain cases,¹⁴ inductive reasoning *requires* justified beliefs in track records and the only way to get to justified track record beliefs is by conjoining the results of individual trials. So on pain of inductive skepticism we better not end up with a general rejection of conjunction introduction as a way of getting to a justified track record belief. Are there any specific reasons to think that conjunction introduction does not work in cases like Boutros and Thant? There do not seem to be any. If Boutros checks n liquids and they all turn litmus paper red, he gains justification for believing that n liquids are acid. And he also gains justification for believing that n liquids are corrosive. He then gains justification for believing that n liquids are both acid and corrosive. What could possibly motivate denying any of this? What he does is a perfectly good way to come to believe that there are precisely n liquids, as opposed to $n+1$, $n-1$ or any other number of liquids, that are both acid and corrosive to iron. Anti Rigging Medium would deny this. Or suppose a fundraiser for the Democratic Party notices a bunch of people individually donating large sums of money to the Party. Surely that can raise justification for believing that there is a bunch of people who are

¹⁴ The preface paradox might be such a reason.

both wealthy and sympathetic to Democratic Party policies. Anti Rigging Medium would deny this, too. Anti Rigging Medium needs to be rejected.

This leads to our final, and we believe, correct propose:

(Anti Rigging) If it ought to be obvious to S, by inspection of the procedure alone, that it cannot recognize counterexamples to the claim that all As are B (i.e. cannot recognize that an A is not a B) even if there are any, then S's justification for believing that all As are B does *not* increase.

Anti Rigging does explain what is wrong with both Thant and Boutros. Thant is checking whether a given claim made by the website is correct and given the way he checks, it is obvious that the procedure won't recognize any instance of a mistake made by the website even if there are any. In Boutros's case, as already noted, it is obvious that he will not recognize any acids that do not corrode iron. Similarly, Trygve cannot use the fact that his advisors agree with him as evidence that he has sound business judgment and this is also explained by Anti Rigging: since his advisors are sycophantic yes-men, Trygve will not recognize a case in which he makes a bad decision by relying on his advisors as they never disagree with him.

Anti Rigging places the problem with the attempt to *use* track record beliefs that are arrived at in certain ways, viz. by using procedures obviously incapable of recognizing counterexamples even if there are any, as a basis for inductive generalization. There is nothing wrong with the track record beliefs themselves in the case of Boutros and Thant. The trouble is that given the way they get to the track record, they can't use the justified

belief in their track record as a basis for inductive generalization. Anti Rigging also points out *one* problem Trygve has. Given the way the case was described, it is plausible that Trygve is not even justified in believing that he got it right on a single occasion. This is a problem. However, Trygve also has another problem. Even if we massaged the case so that Trygve is not obviously unjustified in believing that he got it right on a single or even numerous occasion, Anti Rigging still tells us there is something wrong. Trygve has at least *two* problems and one of them is the one pointed out by Anti Rigging.

5 Dissolving the Easy Bootstrapping Problem

Consider again Kofi and the colored panels. His introspective belief, his belief that he is having an experience as of a green panel (A), is based on his green-panel experience. Kofi's green-panel experience is also the basis for his believing that the panel is green (B). His procedure is obviously rigged. One only needs to understand how the procedure works to see that the procedure cannot recognize cases in which he is having an experience as of a green panel (A) but the panel is *not* green (not-B). That is, it is obvious that the procedure cannot recognize counterexamples to the claim that his vision always gets the color green right even if there are any counterexamples. Thus, Anti Rigging tells us that Kofi's justification for believing that his vision is reliable with respect to the color green does not increase even if he is justified in his track record belief that on numerous occasions his green color vision got it right. So Kofi cannot increase his justification for believing that his color vision is reliable with respect to all colors in the way described above. If you are a non-basic belief theorist, you take this to mean that

even though Kofi is justified in believing his visual system to be reliable his level of justification remains unchanged; if you are a basic belief theorist, you take this to mean that Kofi still lacks justification for believing that his perception is reliable. In either case, there is no easy bootstrapping.

Anti Rigging only blocks the final inductive generalization in Kofi's reasoning. So our proposed explanation leaves intact Kofi's justification for believing on any given occasion that a panel is color C and that he has a visual experience of the panel's being that color. It also leaves intact his justification for believing that on numerous occasions when the panel was color C he also had the visual experience of the panel's being that color. This is as it should be. Our arguments against Anti Rigging Strong and Anti Rigging Medium show that parallel moves in the case of Boutros and Thant are unacceptable so it would be ad hoc either to reject that Kofi is justified in believing that he got it right on any particular occasion or to reject that Kofi is justified in believing that he got it right on numerous occasions without failure.

The problem accounted for by Anti Rigging need not be the only problem Kofi has. No doubt, if one is a non-basic belief theorist, one will charge Kofi with circular reasoning. But our point is that even a non-basic belief theorist should acknowledge that Kofi, just like Trygve, has at least *two* problems. The fact that Kofi's procedure is rigged is a *further* problem that is not solved even if it turns out that Kofi's reasoning is not circular.

Anti Rigging is available to basic belief theory to explain what is wrong with bootstrapping as the explanation in terms of it does not require that the reasoning

depend on the justification for the conclusion. Thus, the *reductio* argument that basic belief theory is false because it cannot block bootstrapping fails.¹⁵

6 References

Alston, W. (1986). Epistemic circularity. *Philosophy and Phenomenological Research*, 47, 1–30.

Alston, W. (1989). *Epistemic justification*. Ithaca, NY: Cornell University Press.

Audi, R. (1993). *The structure of justification*. Cambridge: Cambridge University Press.

Cohen, S. (2002). Basic knowledge and the problem of easy knowledge. *Philosophy and Phenomenological Research*, 65, 309–29.

Cohen, S. (2005). Why basic knowledge is easy knowledge. *Philosophy and Phenomenological Research*, 70, 417–430.

Davies, M. (2004). Epistemic entitlement, warrant transmission and easy knowledge. *Proceedings of the Aristotelean Society, Supplementary Volume 78*, 213–45.

Goldman, Alvin. (1979). What is justified belief? In G. Pappas (Ed.), *Justification and knowledge* (pp. 1–23). Dordrecht: Reidel.

Kung, P & Yamada, Y. (forthcoming). A neglected way of begging the question. *American Philosophical Quarterly*.

Markie, P. (2005). Easy knowledge. *Philosophy and Phenomenological Research*. 70, 406–416

¹⁵ Many thanks to our Claremont colleagues Alex Rajczi, Dion Scott-Kakures, Peter Thielke, Rivka Weinberg, and Charles Young. This paper arose from early conversations we had with Paul Boghossian, and we have also benefited from discussion with Peter J. Graham. Special thanks to Stewart Cohen and Yuval Avnur for pressing objections that led to what we think are substantial improvements in the paper.

Moore, G. E. (1939). Proof of an external world. *Proceedings of the British Academy*, 25, 273–300.

Pryor, J. (2000). The skeptic and the dogmatist. *Noûs*, 34, 517–49.

Pryor, J. (2004). What's wrong with Moore's argument? *Philosophical Issues*, 14, 349–78.

Vogel, J. (2000). Reliabilism leveled. *Journal of Philosophy*, 97, 602–23.

Wright, C. (2002). Anti-sceptics simple and subtle: Moore and McDowell. *Philosophy and Phenomenological Research*, 65, 330–48.

Wright, C. (2003). Some reflections on the acquisition of warrant by inference. In S. Nuccetelli (Ed.), *New essays on semantic externalism and self-knowledge* (pp. 57–77). Cambridge, MA: MIT Press.