

Easy Knowledge and Other Epistemic Virtues

Brian Weatherson, Rutgers/Arché

February 4, 2010

This paper has three aims. First, I'll argue that there's no good reason to accept any kind of 'easy knowledge' objection to externalist foundationalism. It might be a little surprising that we can come to know that our perception is accurate by using our perception, but any attempt to argue this is impossible seems to rest on either false premises or fallacious reasoning. Second, there is something defective about using our perception to test whether our perception is working. What this reveals is that there are things we aim for in testing other than knowing that the device being tested is working. I'll suggest that testing aims for sensitive knowledge that the device is working. Testing a device, such as our perceptual system, by using its own outputs may deliver knowledge, but it can't deliver sensitive knowledge. So it's a bad way to test the system. The big conclusion here is that sensitivity is an important epistemic virtue, although it is not necessary for knowledge. Third, I'll argue that the idea that sensitivity is an epistemic virtue can provide a solution to a tricky puzzle about inductive evidence. This provides another reason for thinking that the conclusion of section two is correct: not all epistemic virtues are to do with knowledge.¹

1 Responding to the Easy Knowledge Objection

I often hear it argued that scepticism is very intuitive. It seems to me that this is completely wrong. I'm currently looking out the window and seeing the rain. It seems to me that I thereby come to know that it is raining. The sceptic says that this isn't at all true. Put that bluntly, it's hard to imagine anything more counterintuitive.

What is intuitive about scepticism is that there are intuitively plausible principles that entail, fairly directly, sceptical conclusions. So although the sceptic's conclusions are counterintuitive, her premises are very plausible. And the implication of her conclusion by her premises is so

¹The references here are obviously incomplete – this should be thought of more as a long blog post than a draft article.

immediate that it sometimes seems that her conclusion is intuitive as well.

You might think that if this was all that was going on with scepticism, it would be easy to argue against the sceptic. We just have to identify those premises, show that they are false, and declare scepticism defeated. The problem is that the sceptic doesn't have to be pinned down to any particular set of premises. She could use premises related to infallibilism, as Descartes does. She could use anti-circular-reasoning premises, as Hume does. Or she could use premises related to sensitivity, as Nozick does. Or, if she's a confused philosophical neophyte, she could bounce between these three arguments without clearly distinguishing them. The result is that any objection to one of the arguments will just leave the sceptic, or the sceptical-friendly neophyte, feeling that the core reasons for being sceptical have not been touched.

I think we see a similar dialectic in the debate over the 'Easy Knowledge' objection to certain kinds of externalist foundationalism. The worry is easy enough to state. Let's assume our target theory says that for some privileged class of cognitive processes, if that process reliably produces true beliefs, then it produces justified beliefs. We'll assume, for ease of exposition, that the privileged class includes visual perception and introspection of what she is perceiving.

Then we imagine an agent who doesn't antecedently know that her vision is working, though in fact it is. She looks at a bunch of things, and (a) sees what is there, and (b) introspects what she is seeing. For instance, she looks out the window, sees it is raining, and introspects that she is seeing this. She believes it is raining on the basis of the perception, and that she sees it is raining on the basis of the introspection. By a simple entailment, she comes to believe her visual system is working correctly. And she can repeat this trick indefinitely many times. So she concludes that her visual system is highly reliable.

Many people intuit directly that something has gone wrong here, and so there is something wrong with our foundationalist externalist.² I think we should be very careful about such intuitions. After all, the following principle seems very plausible. If we know what the output of a machine is, and we can see that this output matches reality, and our visual system is working correctly, then we have good reason to believe that the machine is accurately reflecting reality. And that implies that our agent can know that the machine that is her visual system is working. The reasoning of

²It's tempting to think that the easy knowledge problem is a distinctive problem for externalists, or for those who assign a special role for perception in epistemology. But the structure of the problem is so general that it can arise for anyone. We could use the same structure to raise a problem for a theorist who gives introspection a special role in epistemology. Right now it seems to me that I'm looking at a red carpet, and it seems to me that it seems to me that I'm looking at a red carpet, so my introspection tells me that my introspective access to how things seem to me is accurate. This seems like just as bad a way to check on the accuracy of introspection as is checking on perception using perception.

the last few sentences might fail, but we need a good reason to say why it fails.

Now there are several such reasons that we might propose. In fact, I'll identify 6 such reasons below. I don't think any of these reasons work. But I think they mostly fail for independent reasons. And that's part of what makes the easy knowledge problem so difficult. Just like with scepticism, a not particularly intuitive conclusion is supported by a wide variety of distinct, and intuitive, considerations. When you knock down one, it is easy for the anti-foundationalist, or anti-externalist, to reply using one of the other reasons. So it's best to try to deal with all the reasons at once. I'll make a start on that here.

First, a small disclaimer. Several of the points I make will be familiar from the literature. A chunk of this section is my collecting my thoughts on easy knowledge arguments, not advancing the debate. And since this is basically a blog post in PDF form, not a scholarly article, I haven't cited anyone. But that doesn't mean I'm claiming much of this is original - except perhaps for the broad framing. Having said that, here are the six reasons I think one might find easy knowledge problematic. The last of these is the segue to the discussion of testing in section 2.

1.1 Sensitivity

Objection: If you use perception to test perception, then you'll come to believe perception is accurate whether it is or not. So if it weren't accurate, you would still believe it is. So your belief that it is accurate will be insensitive, in Nozick's sense. And insensitive beliefs cannot constitute knowledge.

The obvious reply to this is that the last sentence is false. As has been argued at great length, e.g. in Williamson (2000: Ch 7), sensitivity is not a constraint on knowledge. We can even see this by considering other cases of testing.

Assume Smith is trying to figure out whether Acme machines are accurate at testing concrete density. She has ten Acme machines in her lab, and proceeds to test each of them in turn by the standard methods. That is, she gets various samples of concrete of known density, and gets the machine being tested to report on its density. For each of the first nine machines, she finds that it is surprisingly accurate, getting the correct answer under a very wide variety of testing conditions. She concludes that Acme is very good at making machines to measure concrete density, and that hence the tenth machine is accurate as well.

We'll come back to the question of whether this is a good way to *test* the tenth machine. It

seems that Smith has good inductive grounds for knowing that the tenth machine is accurate. Yet the nearest world in which it is not accurate is one in which there were some slipups made in its manufacture, and so it is not accurate even though Acme is generally a good manufacturer. In that world, she'll still believe the tenth machine is accurate. So her belief in its accuracy is insensitive, although she knows it is accurate. So whatever is wrong with testing a machine against its own outputs, if the problem is just that the resulting beliefs are insensitive, then that problem does not preclude knowledge of the machine's accuracy.

1.2 One-Sidedness

Objection: If you use perception to test perception, then you can only come to one conclusion; namely that perception is accurate. Indeed, the test can't even give you any reason to believe that perception is inaccurate. But any test that can only come to one conclusion, and cannot give you a reason to believe the negation of that conclusion, cannot produce knowledge.

Again, the problem here is that the last step of the reasoning is mistaken. There are plenty of tests that can only produce knowledge in one direction only. Here are four such examples.

Brown is an intuitionist, so she does not believe that instances of excluded middle are always true. She does, however, know that they can never be false. She is unsure whether Fa is decidable, so she does not believe $Fa \vee \neg Fa$. She observes a closely, and observes it is F . So she infers $Fa \vee \neg Fa$. Her test could not have given her a reason to believe $\neg(Fa \vee \neg Fa)$, but it does ground knowledge that $Fa \vee \neg Fa$.

Jones is trying to figure out which sentences are theorems of a particular modal logic she is investigating. She knows that the logic is not decidable, but she also knows that a particular proof-evaluator does not validate invalid proofs. She sets the evaluator to test whether random strings of characters are proofs. After running overnight, the proof-evaluator says that there is a proof of S in the logic. Jones comes to know that S is a theorem of the logic, even though the failure to deliver S would not have given her any reason to believe it is not a theorem.

Grant has a large box of Turing machines. She knows that each of the machines in the box has a name, and that its name is an English word. She also knows that when any machine halts, it says its name, and that it says nothing otherwise. She does not know, however, which machines are in the box, or how many machines are in the box. She listens for a while, and hears the words 'Scarlatina',

‘Aforetime’ and ‘Overinhibit’ come out of the box. She comes to believe, indeed know, that Scarlatina, Aforetime and Overinhibit are Turing machines that halt. Had those machines not halted, she would not have been in the right kind of causal contact with those machines to have singular thoughts about them, so she could not have believed that they are not halting machines. So listening for what words come out of the box is one-sided in the

Adams is a Red Sox fan in Australia in the pre-internet era. Her only access to game scores are from one-line score reports in the daily newspaper. She doesn’t know how often the Red Sox play. She notices that some days there are 2 games reported, some days there is 1 game reported, and on many days there are no games reported. She also knows that the paper’s editor is also a Red Sox fan, and only prints the score when the Red Sox win. When she opens the newspaper and sees a report of a Red Sox win (i.e. a line score like “Red Sox 7, Royals 3”) she comes to believe that the Red Sox won that game. But when she doesn’t see a score, she has little reason to believe that the Red Sox lost any particular game. After all, she has little reason to believe that any particular game even exists, or was played, let alone that it was lost. So the newspaper gives her reasons to believe that the Red Sox win games, but never reason to believe that the Red Sox didn’t win a particular game.

So we have four counterexamples to the principle that you can only know p if you use a test that could give you evidence that $\neg p$. The reader might notice that many of our examples involve cases from logic, or cases involving singular propositions. Both of those kinds of cases are difficult to model using orthodox Bayesian machinery. That’s not a coincidence. There’s a well known Bayesian argument in favour of the principle I’m objecting to, namely that getting evidence for p presupposes the possibility of getting evidence for $\neg p$. (See ... for further discussion of this.) I haven’t discussed that objection here, because I think it’s irrelevant. When dealing with foundational matters, like logical inference, Bayesian modelling is inappropriate. We can see that by noting that in any field where Bayesian modelling is appropriate, the objection currently being considered works!³ It’s crucial to the defence of foundationalist externalism I’m offering here that this objection not work, so it’s crucial that we not be able to model some features of learning by perception using orthodox Bayesian techniques. If perception is genuinely foundational, like logical reasoning, that shouldn’t be too much of a surprise, given the well-known troubles Bayesians have with representing logical reasoning.

³Or see the large literature on the Problem of Old Evidence.

1.3 Sensitivity

Generality: Assume we can use perception to come to know on a particular occasion that perception is reliable. Since we can do this in arbitrary situations where perception is working, then anyone whose perception is working can come to know, by induction on a number of successful cases, that their perception is generally reliable. And this is absurd.

I'm not sure that this really is absurd, but the cases already discussed should make it clear that it isn't a consequence of foundationalist externalism. It is easily possible to routinely get knowledge that a particular F is G , never get knowledge that any F is not G , and no way be in a position to infer, or even regard as probable, that all F s are G s.

For instance, if we let F be *is a Turing machine in the box Grant is holding*, and G be *halts*, then for any particular F Grant comes to know about, it is G . But it would be absurd for her to infer that all F s are G s. Similarly, for any Red Sox game that Adams comes to know about, the Red Sox win. But it would be absurd for her to come to believe on that basis that they win every game.

There's a general point here, namely that whenever we can only come to know about the F s that are also G s, then we are never in a position to infer inductively that all, or even most, F s are G s. Since even the foundationalist externalist doesn't think we can come to know by perception that perception is *not* working on an occasion, this means we can never know, by simple induction on perceptual knowledge, that perception is generally reliable.

1.4 Circularity

Objection: It is impossible in principle to come to know that a particular method delivers true outputs on an occasion by using that very method. Since the foundationalist externalist says you can do this, her view must be mistaken.

The first sentence of this objection seems to be mistaken. Consider the way we might teach an intelligent undergraduate that (1) is a theorem.

$$((p \rightarrow q) \wedge p) \rightarrow q \tag{1}$$

We might first get the student to assume the antecedent, perhaps reminding them of other uses of assumption in everyday reasoning. Then we'll note that both $p \rightarrow q$ and p follow from this assumption. If we're ambitious we'll say that they follow by \wedge -elimination, though that might be too ambitious. And then we'll draw their attention to the fact that these two claims together imply q , again noting that this rule is called \rightarrow -elimination if we're aiming high pedagogically. Finally, we'll note that since assuming the antecedent of (1) has let us derive its consequent, then it seems that if the antecedent of (1) holds, so does the consequent. But that's just what (1) says, so (1) is true. For a final flourish, we might note the generalisation of that reasoning to the \rightarrow -introduction rule. Or we might call it quits at proving (1); undergraduates can only handle so much logic at one time.

That's a pretty good way, I think, of teaching someone that (1) is true, and indeed why it is true. That is, it is a way the student can learn (1), i.e., come to know (1). Indeed, I think it's more or less the way that I came to learn (1). But note something special about (1). It's obviously closely related to one of the rules we used to prove it, namely \rightarrow -elimination. Indeed, when Achilles first tries to do without that rule in Lewis Carroll's famous dialogue, he appeals to an instance of (1). Given how closely related the two are, you might think using \rightarrow -elimination to come to learn (1) would be viciously circular. But in fact it isn't; the argument we just gave is a perfectly good way of learning (1). I don't know whether we should say it is circular but not viciously so, or simply non-circular, but however we classify the argument it isn't a bad argument. So I conclude that this rule can be used to come to learn that the rule is truth-preserving.

It might be objected at this point that the undergraduate must have already known (1) if they consent to \rightarrow -elimination. I think this turns on an unrealistic picture of what undergraduates are capable of. It looks to me that (1) is in fact a very complicated proposition. It has a conditional embedded in the antecedent of a conditional! Such sentences are not familiar in everyday life, and most introductory logic students have problems parsing such sentences, let alone knowing they are true. So I don't think this kind of Socratic objection works.

It might also be objected that perception is different to logical reasoning. Such an objection might even be made to work. But it isn't a way of defending the objection we opened this subsection with; it relies on a principle that's simply false.

1.5 A Priority

Objection: Assume it is possible to come to know that perception is reliable by using perception. Then before we even perceive anything, we can see in advance that this method will work. So we can see in advance that perception is reliable. That means we don't *come* to know that perception is reliable using perception, we could have known it all along. In other words, it is a priori knowable that perception is reliable.

This objection misstates the foundationalist externalist's position. If perception is working, then we get evidence for this every time we perceive something, and reflect on what we perceive. But if perception is not working well, we don't get any such evidence. The point is not merely that if perception is unreliable, then we can't possibly know that perception is unreliable since knowledge is factive. Rather, the point is that if perception is unreliable, then using perception doesn't give us any evidence at all about anything at all. So it doesn't give us evidence that perception is reliable. Since we don't know antecedently whether perception is reliable, we don't know if we'll get any evidence about its reliability prior to using perception, so we can't do the kind of a priori reasoning imagined by the objector.

This response relies heavily on an externalist treatment of evidence. An internalist foundationalist is vulnerable to this kind of objection. As I've argued elsewhere, internalists have strong reasons to think we can know a priori that foundational methods are reliable. Some may think that this is a reductio of internalism. (I don't.) But the argument crucially relies on internalism, not just on foundationalism.

1.6 Testing

Objection: It's bad to test a belief forming method using that very method. The only way to learn that a method is working is to properly test it. So we can't learn that perception is reliable using perception.

This objection is, I think, the most interesting of the lot. And it's interesting because in some ways I think the first premise, i.e. the first sentence in it, is true. Testing perception using perception is bad. What's surprising is that the second premise is false. The short version of my reply is that in testing, we aim for more than knowledge. In particular, we aim for sensitive knowledge. A test can be bad because it doesn't deliver sensitive knowledge. And that implies that a bad test can

deliver knowledge, at least assuming that not all knowledge is sensitive knowledge. Defending these claims is the point of the next section.

2 The Virtues of Tests

I think the focus on *knowledge* obscures some features of the easy knowledge problem. The easy knowledge problem arises in the context of testing measuring devices. The most common versions of the problem have us testing our own perceptual faculties, but the problem can arise for testing an external device. We can't properly test whether a scale is working by comparing its measurement of the weight of an object to the known weight of the object if the only way we know the weight of the object is by using that very device. Strikingly, this is true even if we do know that the scale is working, and so *knowledge* isn't even at issue. So I'm going to argue that the problem of easy knowledge really should tell us something about *testing*, not something about *knowledge*.

Although I think the easy knowledge problem doesn't, after all, have much to tell us about knowledge, I think it has a lot to tell us about epistemic virtues. The story I'm going to tell about testing is driven by the tension between these two theses.

Fallibilism We can know *p* even though *p* could, in some sense, be false.

Exhaustion Knowledge exhausts the epistemic virtues, so if S knows that *p*, S's belief that *p* could not be made more virtuous.

I think **Fallibilism** is true, and that suggests **Exhaustion** must be false. After all, **Fallibilism** implies that S's belief that *p*, which actually amounts to knowledge, could be strengthened in some ways. It could, for example, be made infallible. So **Exhaustion** is false.

The falsity of **Exhaustion** opens up a conceptual possibility that I don't think has been adequately explored. It is commonly agreed these days that a belief need not be *sensitive* to be known, as we noted in section 1. But this doesn't mean, and in fact doesn't even suggest, that sensitivity is not an epistemic virtue. Indeed, it seems to me that sensitivity is an epistemic virtue, just as infallibility is. Neither of these virtues is necessary for knowledge, but they are virtues nonetheless.

The core hypothesis of this section is that tests of measuring devices aim at *sensitive* beliefs about their accuracy. That is, we aim to draw a true conclusion about the accuracy of the device, and arrive at that conclusion through a method that would have yielded a different verdict were the device's level of accuracy different. And this, I think, explains at least a part of the appeal of the

easy knowledge objection. The way the objection is usually set up, someone is testing a measuring device by a method that clearly cannot yield sensitive conclusions. We conclude, correctly, that this means it is a bad test. I suspect that at least some people infer from that that the test cannot deliver knowledge that the device is inaccurate. This inference isn't, in most cases, conscious, but I suspect it is part of what drives people who find easy knowledge implausible. If what I've said so far is correct, that last step *may* be a non sequiter. Good tests deliver sensitive conclusions, and knowledge doesn't require sensitivity, so it is conceptually possible that a bad test can produce knowledge. To be sure, our knowledge is usually sensitive, so if our belief that the device is accurate is insensitive, that is some evidence that that belief does not amount to knowledge. But it is not a proof that the belief does not amount to knowledge. So a belief grounded in an insensitive, and therefore bad, test may amount to knowledge.

A lot of what I want to say about testing comes from reflection on the following case. It is a slightly idealised case, but I hope not terribly unrealistic.

In a certain state, the inspection of scales used by food vendors has two components. Every two years, the scales are inspected by an official and a certificate of accuracy issued. On top of that, there are random inspections, where each day an inspector must inspect a vendor whose biennial inspection is not yet due. Today one inspector, call her Ins, has to inspect a store run by a shopkeeper called Sho. It turns out Sho's store was inspected just last week, and passed with flying colours. Since Sho has a good reputation as an honest shopkeeper, Ins knows that his scales will be working correctly.

Ins turns up and before she does her inspection watches several people ordering caviar, which in Sho's shop goes for \$1000 per kilogram. The first customer's purchase gets weighed, and it comes to 242g, so she hands over \$242. The second customer's purchase gets weighed, and it comes to 317g, so she hands over \$317. And this goes on for a while. Then Ins announces that she's there for the inspection. Sho is happy to let her inspect his scales, but one of the customers, call him Cus, wonders why it is necessary. "Look," he says, "you saw that the machine said my purchase weighed 78g, and we know it did weigh 78g since we know it's a good machine." At this point the customer points to the certificate authorising the machine that was issued just last week. "And that's been going on for a while. Now all you're going to do is put some weights on the scale and see that it gets the correct reading. But we've done that several times. So your

work here is done.”

There is something deeply wrong with Cus’s conclusion, but it is surprisingly hard to see just where the argument fails. Let’s lay out his argument a little more carefully.

1. The machine said my caviar weighed 78g, and we know this, since we could all see the display.
2. My caviar did weigh 78g, and we know this, since we all know the machine is working correctly.
3. So we know that the machine weighed my caviar correctly. (From 1, 2)
4. By similar reasoning we can show that the machine has weighed everyone’s caviar correctly. (Generalising 3)
5. All we do in testing a machine is see that it weighs various weights correctly.
6. So just by watching the machine all morning we get just as much knowledge as we get from a test. (From 4, 5)
7. So there’s no point in running Ins’s tests. (From 6)

Cus’s summary of how testing scales works is obviously a bit crude⁴, but we can imagine that the spot test Ins plans to do isn’t actually any more demanding than what the scale has been put through while she’s been standing there. So we’ll let premise 5 pass.⁵ If 3 is true, it does seem 4 follows, since Cus can simply repeat his reasoning to get the relevant conclusions. And if 4 and 5 are true, then it does seem 6 follows. To finish up our survey of the uncontroversial steps in Cus’s argument, it seems there isn’t any serious dispute about step 1.⁶

So the contentious steps are:

- Step 2 - we may deny that everyone gets knowledge of the caviar’s weight from the machine.
- Step 3 - we may deny that the relevant closure principle that Cus is assuming here.

⁴For a more accurate account of what should be done, see Edwin C. Morris and Kitty M. K. Fen, *The Calibration of Weights and Measures*, third edition, Sydney: National Measurement Institute, 2002, or David B. Prowse, *The Calibration of Balances*, Lindfield: Commonwealth Scientific and Industrial Research Organization, 1985.

⁵If you’d prefer more realism in the testing methodology, at the cost of less realism in the purchasing pattern of customers, imagine that the purchases exactly follow the pattern of weights that a calibrator following the guidelines of Prowse (1985) or Morris and Fen (2002) would put on the machine.

⁶It may be epistemologically significant that human perception does not, in general, work the same way as balances in this respect. We may come to know p through perception before we know that p was the output of a perceptual module. Indeed, an infant may come to know p through perception even though she lacks the concept of a perceptual module. This may be of significance to the applicability of the problem of easy knowledge to human perception, but it isn’t immediately relevant here.

- Step 7 - we may deny that the aim of the test is (merely) to know that the machine is working.

One way to deny step 2 is to just be an inductive sceptic, and say that no one can know that the machine is working merely given that it worked, or at least appeared to work, last week. But that doesn't seem very promising. It seems that the customers do know, given that the testing regime is a good one, and that the machine was properly tested, that the machine is working. Perhaps there is some reason to think that although the *customers* know this, the *inspector* does not. We'll come back to that. But for now it seems like step 2 is good.

In recent years there has been a flood of work by philosophers denying that what we know is closed under either single-premise closure, e.g. Dretske (2005), or multi-premise closure, e.g. Christensen (2005). But it is hard to see how kind of anti-closure view could help here. We aren't inferring some kind of heavyweight proposition like that there is an external world. And Dretske's kind of view is motivated by avoidance of that kind of inference. And Christensen's view is that knowledge of a conjunction might fail when the amount of risk involved in each conjunct is barely enough to sustain knowledge. But we can imagine that our knowledge of both 1 and 2 is far from the borderline.

A more plausible position is that the inference from 1 and 2 to 3 fails to *transmit* justification in the sense that Crispin Wright (2000, 2004) has described. But that just means that Ins, or Cus, can't get an initial warrant, or extra warrant, for believing the machine is working by going through this reasoning. And Cus doesn't claim that you can. His argument turns entirely on the thought that we already know that the machine is reliable. Given that background, the inference to 3 seems pretty uncontroversial.

That leaves step 7 as the only weak link. I want to conclude that Cus's inference here fails; even if Ins knows that the machine is working, it is still good for her to test it. But I imagine many people will think that if we've got this far, i.e., if we've agreed with Cus's argument up to step 6, then we must also agree with step 7. I'm going to offer two arguments against that, and claim that step 7 might fail, indeed does fail in the story I've told, even if what Cus says is true up through step 6.

First, even if Ins won't get extra knowledge through running the tests on this occasion, it is still true that this kind of randomised testing program is an epistemic good. We have more knowledge, or at least more reliable knowledge, through having randomised checks of machines than we would get from just having biennial tests. So there is still a benefit to conducting the tests even in cases

where the outcome is not in serious doubt. The benefit is simply that the program, which is a good program, is not compromised.

We can compare this reason Ins has for running the tests to reasons we have for persisting in practices that will, in general, maximise welfare. Imagine a driver, called Dri, is stopped at a red light in a quiet part of town in the middle of the night. Dri can see that there is no other traffic around, and that there are no police or cameras who will fine her for running the red light. But she should, I think, stay stopped at the light. The practice of always stopping at red lights is a better practice than any alternative practice that Dri could implement. I assume she, like most drivers, could not successfully implement the practice *Stay stopped at red lights unless you know no harm will come from running the light*. In reality, a driver who tries to occasionally slip through red lights will get careless, and one day run a serious risk of injury to themselves or others. The best practice is simply to stay stopped. So on this particular occasion Dri has a good reason to stay stopped at the red light: that's the only way to carry out a practice which it is good for her to continue.

Now Ins's interest is not primarily in welfare, it is in epistemic goods. (She cares about those epistemic goods because they are related to welfare, but her primary interest is in epistemic goods.) But we can make the same kind of point. There are epistemic practices which are optimal for us to follow given what we can plausibly do. And this kind of testing regime may be the best way to maximise our epistemic access to facts about scale reliability, even if on this occasion it doesn't lead to more *knowledge*. Indeed, it seems to me that this is quite a good testing regime, and it is a good thing, an epistemically good thing, for Ins to do her part in maintaining the practice of randomised testing that is part of the regime.

The second reason relates to the tension between **Fallibilism** and **Exhaustion** I mentioned at the top. It may be that the aims of the test are not exhausted by the aim of getting knowledge that the machine is working. We might also want a sensitive belief that the machine is working. Indeed, we may want a sensitive belief that the machine has not stopped working since its last inspection. That would be an epistemic good. In some sense, our epistemic standing improves if our belief that the machine has not stopped working since its last inspection becomes sensitive to the facts.

This idea, that tests aim for sensitivity, is hardly a radical one. It is a very natural idea that good tests produce results that are correlated with the attribute being tested. When we look at the actual tests endorsed in manuals like Prowse (1985) or Kitty and Fen (2002), that seems to be one of their central aims. But 'testing' the machine by using its own readings cannot produce results that are correlated with the accuracy of the machine. If the machine is perfectly accurate, the test will say it

is perfectly accurate. If the machine is somewhat accurate, the test will say it is perfectly accurate. And if the machine is quite inaccurate, the test will say that it is perfectly accurate. The test Ins plans to run, as opposed to the ‘test’ that Cus suggests, is sensitive to the machine’s accuracy. Since it’s good to have sensitive beliefs, it is good for Ins to run her tests.

So I conclude that step 7 in Cus’s argument fails. There are reasons, both in terms of the practice Ins is part of, and in terms of what epistemic goods she’ll gain on this occasion by running the test, for Ins to test the machine. That’s true even if she knows that the machine is working. The epistemic goods we get from running tests are not restricted to knowledge.

3 Evidence and Epistemic Virtues

Timothy Williamson has argued in several places that our evidence consists of all and only the things we know. This view, usually called $E=K$, has become quite popular in recent years. It’s tempting to think it counts much to much as evidence. Here’s one quick objection to $E=K$ that doesn’t work. Some of our knowledge is not self-evident. But any time a piece of knowledge is part of our evidence, it is self-evident. So $E=K$ is false. But the middle premise there fails. For something to be self-evident, it isn’t just true that it must be part of our evidence, it must also be that our knowledge of it is not *based* on anything else. And $E=K$ does not imply anything about basing.

There are other ways to argue for a similar conclusion. If I observe that E , and inductively infer H , and thereby come to know H , it seems a little odd to say that H is part of my *evidence*. If we ask ourselves in such a case, what possibilities are inconsistent with my evidence, then intuitively they include the $\neg E$ possibilities, but not the $\neg H$ possibilities. On the other hand, consider an extension of the case where H provides some support for a further conclusion H' . If we say that H' is probable, and someone asks us what evidence we have for this, it is often tempting to reply that our evidence is H . These examples have been rather abstract, but I think you’ll find if you fill in real propositions for E, H and H' , you’ll find that both intuitions have some support. Sometimes we intuitively treat inductively drawn conclusions as part of our evidence, sometimes we don’t. Intuition here doesn’t unequivocally support $E=K$, but it doesn’t totally undermine it either.

Still, the idea that inductive inferences form part of our evidence does create a puzzle. Here’s one way to draw out the puzzling nature of $E=K$.

Jack is inspecting a new kind of balance made by Acme Corporation. He thoroughly

inspects the first 10 (out of a batch of 1,000,000) that come off the assembly line. And each of them passes the inspection with flying colours. Each of them is more accurate than any balance Jack had tested before that day. So Acme is making good balances. He knows, by observation, that the first 10 balances are reliable. He also knows, by induction, that the next balance will be reliable. It's not obvious that he knows the next one will be phenomenal, like the ones he has tested, but he knows it will be good enough for its intended usage. But he doesn't know they all will be that good. Surprisingly, it will turn out that every balance made in this assembly line will be reliable. But you'd expect, given what we know about assembly lines, for there to be a badly made machine turn up somewhere along the way.

So Jack knows the first 11 are reliable, and doesn't know the first 1,000,000 are reliable. Let n be the largest number such that Jack knows the first n are reliable. (I'm assuming such an n exists; those who want to hold on to $E=K$ by giving up the least number theorem are free to ignore everything that follows.) For any x , let $R(x)$ be the proposition that the first x are reliable. So Jack knows $R(n)$. Hence by $E=K$ $R(n)$ is part of his evidence. But he doesn't know $R(n+1)$. This is extremely odd. After all, $R(n)$ is excellent evidence for $R(n+1)$, assuming it is part of his evidence. And $R(n+1)$ is true. Indeed, by many measures it is safely true. So why doesn't Jack know it?

It seems to me there is a mystery here that, given $E=K$, we can't explain. If we have a more restrictive theory of evidence, then it is easy to explain what's going on. If, for instance, evidence is perceptual knowledge, then Jack's evidence is simply $R(10)$. And it might well be true, given the correct theory of what hypotheses are supported by what evidence, that $R(10)$ supports $R(84)$ but not $R(85)$. That explanation isn't available to the $E=K$ theorist. And we might well wonder what explanation could be available. Similarly, we can't say that it is too much of a stretch to infer $R(85)$ from Jack's evidence. After all, his evidence includes $R(84)$. If that's not a sufficient evidential basis to conclude $R(85)$, we're getting very close to scepticism about enumerative induction. If this is not to appear to be a concession to the sceptic, we need a story about why this case is special.

I think the best explanation of why the case is special turns on the idea that there are different epistemic virtues, and some of these cross-cut the presence of absence of knowledge. I don't think this is the *only* explanation of what's going on, but I think it tracks the phenomenon better than

the alternative.⁷ Let's say that evidence, like knowledge, can be of better or worse quality. If you don't know p , then p is of no evidential use to you. That's the part $E=K$ gets right. But even if you do know p , how much evidential use it might depend on how you know it. For instance, if you infallibly know p , then p is extremely useful evidence. More relevantly for current purposes, if you have sensitive knowledge that p , then p is more useful than if you have insensitive knowledge that p . That's because, as we argued in the previous section, sensitive knowledge has epistemic qualities over and above insensitive knowledge.

Let's go through how this plays out in Jack's case. Although he knows $R(11)$, this knowledge is insensitive. If $R(11)$ were false, he would still believe it. Had the production system malfunctioned when making the 11th balance, for instance, then the 11th machine would have been unreliable, but Jack would have still believed it. The only sensitive evidence he has is $R(10)$. By the time he gets to $R(n)$, his knowledge is extremely insensitive. There are all sorts of ways that $R(n)$ could have been false, in many fairly near worlds, and yet he would still have believed it.

This suggests a hypothesis, one that dovetails nicely with the conclusions of the previous section. The more insensitive your evidence is, the less inductive knowledge it grounds. If Jack had sensitive knowledge that $R(n)$, he would be in a position to infer, and thereby know $R(n+1)$. The reason he can't know $R(n+1)$ is not that he doesn't have enough evidence, but rather that the evidence he has is not of a high enough quality. That's an explanation for why Jack can't infer $R(n+1)$ that neither leads to inductive scepticism, nor violates the letter of $E=K$. Sometimes the way Williamson presents $E=K$ suggests he wants something stronger, namely that what we know determines not only what our evidence is, but which things are supported by our evidence. I think that isn't true. Indeed, the best explanation of Jack's predicament is that it isn't true. But I also don't think it is something that the original arguments for $E=K$ support, so we shouldn't feel any hesitation about dropping it.

⁷The kind of alternatives I have in mind are ones where we appeal to other knowledge that Jack has. For instance, we might note that he knows he knows $R(10)$, but doesn't know that he knows $R(n)$. This explanation will have to explain why this second-order knowledge is relevant to inductive inference. After all, why should facts about what Jack knows be relevant to whether the $n+1$ 'th machine will function well? I think similar explanations will suffer similar defects. I'm grateful here to conversations with Jonathan Weinberg.