

Draft of 7/2006

Intuition & Calibration

Jonathan M. Weinberg (jmweinbe@indiana.edu)
Stephen J. Crowley
Chad Gonnerman
Stacey Swain
Ian Vandewalker

Indiana University – Experimental Epistemology Laboratory

Abstract: The practice of appealing to esoteric intuitions, long standard in analytic philosophy, has recently fallen on hard times. Various recent empirical results have suggested that philosophers are not currently able to distinguish good intuitions from bad. This paper evaluates one possible type of approach to this problematic methodological situation: calibration. Both critiquing and building on an argument from Robert Cummins, the paper explores what possible avenues may exist for the calibration of philosophical intuitions. It is argued that no good options are *currently* available, but leaves open the real possibility of such a calibration in the future.

1. The Challenge for Philosophical Intuitions

That analytic philosophers typically – perhaps even stereotypically – rely on intuitions is hardly a claim in much need of defense.¹ But what is coming to require something of a defense is that reliance itself.² Our interest here is the exploration of one possible means of defense: calibration. But first, we wish to summarize the state of play,

¹ Though see Pust (2000) if such a defense is needed.

² See, e.g., Weinberg, Nichols, and Stich (2001); Cummins (1998); and Miller (2000) on the attack; and Jackson (1998), Bealer (1998), and Weatherson (2003) on the defense.

since without a picture of what the challenge to intuitions is supposed to be, it will be hard to evaluate what responses to that challenge will or will not work.

1.1. What's an intuition?

One runs into immediate difficulties here, in that there is little consensus among those philosophers who discuss intuitions as to just what it is that they are discussing. (See Sosa (forthcoming) and Lynch (forthcoming) for a recent exchange attempting to analyze “intuition”.) Perhaps intuitions are *sui generis* propositional attitudes with their own distinctive phenomenology (Bealer (1996, 2000)); perhaps they are a species of belief, or of disposition to belief (Sosa (1996)). Perhaps their propositional content must have abstract or modal content (Sosa (1998); Bealer (1993); Pust (2000); perhaps not (Kornblith (forthcoming); Goldman (forthcoming)). Perhaps they are grounded in our concepts (Peacocke (2001); Goldman and Pust (1998); Henderson and Horgan (2001); Jackman (2005)) or acquired empirical knowledge (Levin (2004); Devitt (forthcoming)) or perhaps they are a direct glomming on to some aspect of reality (Katz (1998); Bonjour (1998)). And of course this diversity can only be located among those who actually write explicitly about intuitions – the vast majority of philosophers who use intuitions spend no time at all setting forward what they have in mind by that term. Those who would challenge intuitions thus risk finding themselves wrestling with a protean opponent, and with a dilemma: operate without a clear definition of intuition and risk the charge of obscurity, or contend under the rubric of some precise formulation of intuitionhood, and risk being responded to with a chorus of “oh, well, *that's* not what *we* had in mind”. We propose to perform an end-run around this problematic by beginning our discussion not

in terms of the intuitions themselves, but in terms of the practice of analytic philosophical appeals to them. For, looking at how that practice is structured, we can determine what intuitions would have to be, in order that such a practice could even in principle be well-formed.

In the practice of appeal to intuitions as philosophical evidence, one can cite in defense of one's claims one's application or withholding of a concept from a given case, usually a hypothetical one. This citation requires no further argumentation, and in particular no empirical backing, because one is presumed to have stipulated all the relevant contingencies. Such citations are thus, in one sense, foundational: although they are used to provide evidence, one does not, and need not, provide further evidence for them. However, they are not generally taken to be incorrigible or indubitable (BonJour (1998); Bealer (1996, 2000); Weatherson (2003)). In particular, they can be both overturned (e.g., by uncovering sufficient evidence in the opposite direction) and undercut (e.g., by discovering that one has not considered the case sufficiently closely). In this way, the practice of appeal to intuitions is similar to the practice of appeal to perception, or memory.³ It has critical dissimilarities with both of those standard practices, however. For our perceptual practices operate under a rich set of norms, developed over millennia of ordinary usage, guiding how much weight to place on any such deployment: more weight is granted when deployed in good lighting, when others can engage in the same perceptual activities, and so on. Similarly with memory (though we suspect that the norms with memory are not as successful as those with perception; to take just one recent example out of literally hundreds, see Loftus (2003)). We have no such well-developed ordinary norms for the appeal to intuition – which surely is part of

³ This is a point frequently made by proponents of intuition; see, e.g., Bealer (1998), Sosa (1998)).

the explanation for the noisiness of the philosophical literature on the topic. Furthermore, these norms impose fairly strict constraints on what sorts of propositions might be defended by such appeals: memory can only be cited concerning matters that one has previously had some encounter with or testimony about, for example, whereas it is vital to our practices with intuition that we can cite it for judgments about scenarios we have never before encountered. Standard philosophical practice, indeed, seems to require that we can cite it in the evaluation of propositions far outside the scope of our ordinary lives (e.g., twin earths and demon deceivers and the like). For perception and memory, we must have or have had some sort of appropriate causal contact with the entities about which we are judging; not so with intuition.⁴

So, what intuitions must be – what analytical philosophical practice requires that intuitions be – are first of all propositional seemings. They must be propositional, in order to serve as citable evidence at all, and their status as seemings distinguishes them from ordinary beliefs, which cannot usually be cited as evidence without opening oneself up to the possible obligation to offer further defense. Furthermore, because they are seemings – that is, they present their contents as true – they are distinct from imaginings or supposings, which lack such a character (though this does not distinguish them from delusions or hallucinations). Moreover, intuitions must be a fundamentally different sort of mental happening than memories or percepts, since they do not have the causal limitations of those other seemings – they can concern matters to which we are causally connected, but they do not need to. They must also be distinguished somehow from

⁴ The obvious exception being, of course, memory of prior intuitions.

hunches, premonitions, and other sorts of disreputable seemings.⁵ Beyond these minimal requirements of philosophical practice, we will not build anything more into the notion of “intuition” in play here.

1.2. The nature of the challenge

Our goal is to explore the viability of one particular form of response to the challenge to the appeal to intuitions. So we must first canvass just what that challenge is. One can start with the history of philosophy itself, which presents a wealth of instances of intuitions gone bad. A classic example here is Kant’s claim, early in the *Critique of Pure Reason*, that space is necessarily Euclidean. Another is his notion of analytic necessity: to claim that gold is necessarily yellow metal is to rely on an intuition now recognized to be faulty.⁶ Leibniz based a good deal of his work on the principle of sufficient reason, thereby presenting a ripe target for such criticisms as Voltaire’s anti-panglossian riffs in *Candide*. Descartes based his (to more contemporary eyes) spurious causal principle on intuition: “...it is manifest by the natural light that there must be at least as much reality in the efficient and total cause as in the effect of that cause” (Descartes 1984, p. 41). More recently, the naïve comprehension axiom presents an intuition both powerfully felt and substantively disconfirmed. (Levin (2004) also surveys some further cases along these lines.) Cognitive psychologists have demonstrated further systematic shortcomings of our intuitive capacities, from probabilistic reasoning (e.g., Kahneman, Slovic, and

⁵ Which is not to say, of course, that what present themselves to our minds as intuitions are not *sometimes* memories, hunches, etc. that are mistaken for intuitions, or hunches. All our philosophical practices seem to require is that these are real distinctions, and that we do not generally make such misidentification. We will grant, for the purposes of this discussion, that these particular requirements are met.

⁶ Of course, neither of these are intuitions in Kant’s usage of “intuition” – or, rather, “Anschauung” – but that usage is not the standard one for contemporary analytic philosophers.

Tversky (1982)) to intuitions about necessity and logical validity (Evans and Over (1996)).

Other recent work has focused directly on the sorts of intuitions appealed to by philosophers, and raises serious worries about such intuitions. In a series of experiments, Weinberg, Nichols, and Stich (2001) tested the intuitions of various groups on such paradigmatic philosophical thought-experiments as the Gettier case (1963), Dretske's 'zebra' case (1970), and Lehrer's 'Truetemp' case (1974). They found significant differences in intuition across both ethnic and socio-economic lines. Further research has demonstrated similar differences across ethnic lines in intuitions concerning semantics (Machery, Mallon, Nichols and Stich (2004)), and has suggested mismatches between the intuitions of philosophers and those of untrained folk (Nichols, Stich and Weinberg (2003)); Nahmias *et al.* (forthcoming)).⁷ Still other research suggests that intuitions about such centrally philosophical theses as compatibilism and incompatibilism can be influenced by such metaphysically irrelevant factors as the emotional tenor of the scenario (Nichols and Knobe (forthcoming)). There is also recent evidence that intuitions of individual persons may be malleable, depending too much on which other hypothetical cases have recently been considered (Swain, Alexander, and Weinberg (2006)).

This troublesome historical and empirical evidence presents a challenge to the philosophical practice of appeal to intuitions: there is significant extant evidence that

⁷ Must the troublesome evidence be understood as a case of the *a posteriori* threatening the *a priori* -- which some philosophers, such as Bonjour, have argued is an impossibility? We have two main responses to that question. First, the putative impossibility would be one of the empirical showing the *a priori* to be *completely* untrustworthy, resulting in a strong form of skepticism about intuitions. But our worries are more restrained: we do not take the troublesome evidence to show that *no* intuitions have *any* evidential value, but only that *esoteric* intuitions have much *less* evidential value than philosophers have presumed (see below). Second, one can view the situation as one not of *science* threatening intuitions, but rather as one in which the science is simply informing us that the *intuitions* are impeaching each other. The science is not interfering illicitly from the outside, but rather is informing us of a difficulty that exists already within the house of reason.

such intuitions are not trustworthy. A number of prominent philosophers have attempted to respond to this challenge with the observation that an argument for skepticism about intuition will inevitably become an argument for skepticism more generally. One way that such skepticism would spread is by way of parity of argument: if intuition is fallible, so too is perception, so any argument against intuition on the ground of its fallibility must just as well be turned against perception (Sosa (1998); Goldman (forthcoming)). *Mutatis mutandis* for providing a fully non-circular justification for the use of intuition or perception (Bealer (1998)). Intuition-skepticism might also lead to a more general skepticism if those intuitions play an ineliminable role in most of our knowledge about the world, for example, providing the justification for the inferences that take us beyond the immediate here-and-now of perception (BonJour (1998); Williamson (2004a)). It thus seems that one can infer from the skepticism about intuitions to a more radical and global skepticism, and so one needs only be committed to the rejection of the latter sort of skepticism to *modus tollens* to the rejection of the former sort.

We concur with the general thrust of these anti-skeptical arguments, and have no desire to traffic in the sorts of skepticism that they warn against. And so it is important to us that our construal of the challenge here does not lead to a general skepticism about all intuitions, nor thereby to a global skepticism. In order to show that our concerns here are not skeptical in that way – and hence that they may perhaps need to be taken seriously in a way in which a more obviously overreaching set of concerns might not be – it will be useful for us to clarify three aspects of the challenge. First, we need to articulate which intuitions are challenged, if it is not all of them; next, just how challenged those intuitions

are; and third, why it might not be possible to apply the same kind of challenge to perception.

1.3. Which intuitions are challenged?

We take the checkered history of philosophical appeals to intuitions, together with the more recent experimental results, to impugn the reliability of one class of intuitions most particularly: those that involve the application of concepts to situations that lie significantly outside the ordinary practices governing them. Applications of the term "chair" to the typical office chair are not under attack here, for we take such practice-typical usages to be well-regulated by those practices. But our practices are simply not prepared to adjudicate cases like Wittgenstein's appearing-and-disappearing chair-like object (*Philosophical Investigations*, §80); and, indeed, we take it that it was at the heart of Wittgenstein's point here (and in his preceding remarks on blurriness) that such cases fall outside the sphere of our practices, yet those ordinary practices themselves are not threatened by this fact.

Experts' application of their expertise often manifests as intuitions, such as the doctor's swift emergency diagnosis, or the chess master's sense of which moves are worthy of close consideration. We do not take these intuitions to be challenged, either. Or, perhaps better, we take it that these intuitions have the means to meet any similar challenge. For example, experts can point to their own excellent track record in their domain of expertise, and we can test such claims against independent standards of success, like the saved patient or the victorious endgame.

Williamson has also argued in his (2004b) that there is no reason to be nervous about our modal cognition in particular, which is implicated in many of the most mundane sorts of planning and reasoning; and that anyways there are real Gettier cases in our own possible world.⁸ And we will not gainsay his arguments here, and happily acquiesce in his claim that we are generally very good at such reasoning. But the distinction between ordinary and esoteric cross-cuts the distinction between the actual and the counterfactual. The problem with Gettier cases, actual or not, is that they lie well outside the typical cases for which we consider applying or withholding the term “knows”. (And Gettier cases are some of the less wacky ones that epistemologists monger about; compare to, e.g. Lehrer’s gypsy lawyer (1971) or Lehrer and Cohen’s new evil demon (1983) or Lackey’s double-lesioned truth-telling liar (forthcoming).) And we take such results as Weinberg, Nichols, and Stich (2001) to have shown exactly that philosophers need to be much more careful about crediting their own judgments about Gettier cases (especially philosophers of European descent who mostly talk to other philosophers of European descent; namely, most of the profession).

1.4. How challenged are they?

If we are right that these intuitions about esoteric cases – and we will stipulatively label such intuitions “philosophical intuitions” (PI’s), to contrast them with ordinary intuitions⁹ – should be viewed with suspicion, one might reasonably ask just how much

⁸ Williamson does not discuss our modal cognition under the rubric of modal intuitions (see §1.6 below), but his claims about the general success of our modal cognition do not rely on his antipathy to intuition-talk.

⁹ Which is not to say that philosophers do not at least sometimes appeal to ordinary intuitions, too. But it is perhaps unsurprising that mostly they don’t. After all, the ordinary cases will rarely help decide between rival philosophical theories, because theories that violate too many of our ordinary-case intuitions are unlikely to survive long enough to be advanced in the literature.

suspicion we have in mind. It would be too radical to claim that our intuitions about such cases are devoid of *any* justificatory force *whatsoever*; a fairly mild form of conservatism or Reidian self-trust would prevent even that restricted a version of intuition skepticism,¹⁰ as would perhaps a sort of induction from our general success on the more ordinary sorts of cases.¹¹

But just as the challenge is not basically a skeptical challenge, so too is it not addressed by a merely anti-skeptical response. Even if esoteric intuitions are not completely devoid of epistemic potency, are they capable of bearing the tremendous evidential weight that philosophers rest on them? Such intuitions are frequently cited not merely as having some minor bearing on a given philosophical debate, but as central to the argument, even dispositive (Pust (2000); Weatherson (2003)). The key starting point for our paper, then, is that current analytic philosophical practice (especially in much of epistemology and metaphysics) presupposes that the evidential status of such intuitions is very strong; and that so long as the challenge stands unanswered, philosophers are not entitled to that presupposition.

We should emphasize that we do not take intuitions to *start off* as challenged, and will not object to those who, like Bealer, construe intuitions as starting with a presumptive positive epistemic status. As noted above, we do not find it legitimate to require that, for intuitions to count as evidence, they antecedently have an independent verification of their reliability. But we follow Goldman in thinking that “it seems reasonable to substitute a weaker condition as a ... requirement on evidencehood. This weaker condition is a “negative” one, viz., that we *not* be justified in believing that the putative

¹⁰ Tidman (1996)

¹¹ See also our discussion of Bealer in §5.2.3. below.

source is *unreliable*” (Goldman (forthcoming); emphasizes original). And thus when the reliability of a putative source of evidence *is* challenged, those who would continue deploying that source must respond to the challenge, or adapt their use of the source so that the challenge no longer applies.

We offer this principle not as an intuition (which might leave us open to a response of *tu quoque* or worse), but as a generalization from both ordinary and scientific practice. If you would cite a given magazine’s reportage in an argument, and your opponent points out a pattern of past errors in that magazine, then you cannot simply continue appealing to that publication without, say, showing that the alleged past errors were not such; or that there has been a change in management since those mistakes; or the errors occurred only under one journalist's byline, and this article is by a different author; and so on. But our ordinary practices do not allow for you to simply keep relying on that magazine without saying something significant in its defense. To take just one example from the sciences, electron microscopy was not well-trusted until biologists could assure themselves that various potential sources of error (like the violence that had to be done to specimens, in order to get thin enough slices) were under control (Bechtel (2000)). And introspection was rejected as a method in psychology for decades, following the failures of Wundt and Titchener. It was only rehabilitated decades later under the careful external controls of protocol analysis, and then only in a comparatively weaker form than deployed by its earlier adherents. And indeed, we find the development of protocol analysis by Newell, Simon and others to constitute an excellent example of a practice at one time being unable to meet a challenge of unreliability, but ultimately being salvaged by changing the practice to meet the critics’ worries. What is good enough for our

ordinary epistemic lives and for scientific practice is surely good enough for philosophical methodology; and the application of this challenge to PI's is not a skeptical break from our ordinary practices, but is a consistent application of them.¹²

1.5. Is there a parallel challenge to perception?

One might test our construal of the challenge, by evaluating how well a parallel challenge might apply to sense perception. Surely we are aware of as many -- more than as many -- instances in which perceptions have diverged from each other, and from the truth. So how does this argument *not* entail an embargo on perception as well?

Pace Sosa, the epistemic disanalogies between the appeal to perception and the appeal to philosophical intuitions are numerous. We have a very robust set of checking and correcting practices for perception, against both each other's perceptions and other established standards. We have an excellent working theory -- mostly a folk theory, but more recently supplemented by the deliverances of scientific psychology -- of where perception will and will not work reliably. We know pretty well what facts perception can reveal to us (e.g., the color of surfaces in good lighting) and what they can't (e.g., the content of another person's soul). Moreover, perception itself gives us a good sense of when we are operating within good circumstances for its use. Perception works poorly in dimly-lit conditions, and fortunately for us, such conditions are seen *as* dimly-lit, and not *as* strangely gray and blurred.

None of this applies to esoteric intuitions of the sort commonly used by philosophers. These intuitions have not been checked against the intuitions of a wide range of

¹² This brief argument can also be taken as pointing towards one kind of response-by-example to the worry that, if we give up on the method of intuitions-about-funny-cases, philosophizing will become impossible.

individuals (as has indeed been revealed by such results as Weinberg, Nichols, and Stich and Machery *et al.*). Beyond the roughest guidelines, like not being intoxicated or heavily distracted, we have no good sense of the circumstances that are better or worse for having reliable intuitions. Nor do we have much of a working theory as to which propositions are tracked successfully by our intuitions of them. Sosa claims that we have a "practical sensitivity" (1998, p. 266) to the bounds of the reliable deployment of intuition, and that we can somehow know to rely upon intuitive beliefs only within those bounds. We agree, insofar as one is concerned with the sorts of intuitions that are relied upon in our everyday lives, such as those involved in basic arithmetic. But there are zones where this practical sensitivity breaks down, and the preliminary state of the evidence suggests that PI's generally lie beyond that boundary.

Given these substantial differences between the epistemic resources available to perception on the one hand, but absent in PI's on the other, one can safely challenge philosophers' reliance on esoteric intuitions without fearing that skepticism about perception will threaten as well.

1.6. The state of play & the goal of this paper

So, here is our take on the state of play: intuitions about esoteric cases are in a state of challenge, in that significant evidence against their reliability has been amassed to which no response has yet been found. Philosophers must therefore respond to that challenge on their behalf, perhaps by modifying their intuition-deploying practices so that the challenges no longer apply; or perhaps by placing an embargo on such deployments. The rest of this paper proceeds from that starting point, and aims to evaluate one obvious

candidate for addressing the challenge: calibration. To be clear, we do not think that calibration is the *only* possible response to the challenge, but rather the most natural first option for what to do with a putative source of evidence whose trustworthiness is in doubt. Other responses may be available. For example, one might construe the relation between intuitions and the facts that they report on not as evidential, but as constitutive (Weatherson (2003); Jackman (forthcoming)), and explore from there the consequences for the challenge. Williamson has recently suggested that doubts about standard philosophical methodology arise in large part from a misconception that intuitions are needed at all, and he argues that we should take the relevant facts themselves to be our evidence (Williamson (2004a)). Perhaps this re-envisioning of philosophical practice will allow it to escape the challenge, although we have our doubts.¹³ But we shall not canvass such moves or their shortcomings here. Our project is to evaluate the prospects of calibration as *a* way of trying to save the practice of appeal to esoteric intuitions.

Furthermore, we do not expect all readers to be persuaded, as yet, that PI's are really in a state of challenge. We do hope, however, that even for those readers, this initial discussion suffices to convince them that such a challenge is a real possibility, and that it need not be skeptical in nature, and thus cannot be ignored along with the really skeptical arguments. Because laying out the prospects for the calibration of PI's is our primary interest here, we invite such readers to treat this initial section as a hypothetical: *suppose* that philosophical intuitions were in a state of challenge – could calibration be a way to save them?

¹³ [identifying reference suppressed]

2. The process of calibration

We take it that the most standard way to restore trust to a putative source of evidence whose reliability has been impeached is to *calibrate* it. Fans of intuition might plausibly hope that this process, so valuable in the sciences, can rescue intuitions from an ignominious end. Our project here is to evaluate the prospects for such a calibration, and to do so, we will begin with a general though by no means complete picture of calibration.

Calibration is a process of regulating a putative source of evidence, by inspecting it and, if needed, adjusting it to render it accurate. Calibration is divisible into three parts: *testing*, *diagnosis*, and *correction*.

Testing is a matter of checking the reliability or accuracy of the target device.¹⁴ During testing, we use the device to make a series of observations of phenomena within the device's intended domain of application. We then attempt to evaluate the accuracy of those observations that can be checked against another source.¹⁵ For example, a scientist might determine that a mercury thermometer is faulty if its readings do not compare well with those of an already trusted platinum thermometer.¹⁶

¹⁴ We will henceforth use "device" instead of "putative source of evidence", trusting that nothing here is lost by doing so. We recognize that not all sources of evidence can be usefully assimilated to devices – to do so with testimony would be to lose our grasp on the epistemic agency of our fellows, for example. But there is no such loss at risk with intuitions. Moreover, we take it that calibration will simply be a non-starter for someone who takes there to be a radical gulf between PI's on the one hand and scientific instruments on the other. So this paper is aimed more at those willing to endorse a fundamental similarity between the two.

¹⁵ It should be stressed, though, that testing does not always require an independent basis of comparison. Although we would not want to trust a device's assertions of its own trustworthiness, sometimes the device will impeach itself with a set of self-incriminating observations. If the device produces a significant number of mutually incompatible observations, then we know something must be wrong. This 'theory rich' approach to calibration will be discussed below.

¹⁶For a rich and thoughtful account of the full range of challenges that must be overcome in order to calibrate a scientific instrument, with a particular focus on thermometers, see Chang (2004).

If the process of testing reveals that the target device is sufficiently accurate, then there is no need to proceed to the other two stages of calibration. On the other hand, if an undesirable amount of deviation is found – and if we want to continue using the device – then the next step in a calibration will be to determine in what way the device goes awry. Optimally, we would have a worked-out theory of the source of error; for example, physics enables us to understand the circumstances in which a pocket compass will fail to point north, such as when it is near a large iron deposit or a coiled electrical wire. Diagnosis may still go forward in the absence of such a theory, to whatever extent we can tell at least in what kinds of cases the device tends to be more reliable, and where less.

If testing has revealed an undesirable amount of deviation in our device, and we have a diagnosis of this deviation, then we finally face the question: how do we go about correcting for this deviation? There are at least two methods for correction: *restriction* and *rehabilitation*. With restriction, as the name suggests, we simply ignore the device's results in circumstances when we expect it to be unreliable. With rehabilitation, we tweak the problematic device itself, rendering it accurate across its entire domain.¹⁷ In order to make these two methods clearer, we will consider their application in a case borrowed from science: the astronomer's problem of seeing.

'Seeing' is a technical term in astronomy, referring to the blurring and twinkling of an image caused by turbulence in the Earth's atmosphere.¹⁸ Albeit romantic in nonscientific contexts, this phenomenon is a menace to astronomers. Although we have good astrophysical reasons to suspect that celestial bodies are themselves neither blurry nor

¹⁷Of course, in correcting an information-gathering device, some combination of restriction and rehabilitation might be called for; we might call this sending the device for some much-needed 'R&R'.

¹⁸For more on 'seeing', see e.g. Fix (1995).

twinkly, our naïve telescopic observations would lead us to believe that they are both. The testing stage, thus, indicates something amiss with our practice for using the device.

We can turn to optics to diagnose the source of these errors. Each time light from a celestial object hits a pocket of air with an index of refraction that is different from its surroundings, the light from the celestial object is bent, ultimately producing a blurred image as well as presenting astronomers with the frustrating fact that they rarely acquire images as sharp as the diameter of their telescopes allow. Astronomers can correct for the undesirable effects of seeing via both restriction and rehabilitation. First, astronomers can *restrict* their usage of the telescope to circumstances within which it will produce sharp, untwinkling images – for example, by observing only large and proximate objects, like the moon. Astronomers dissatisfied with such a restriction can correct for the effects of seeing by *rehabilitating* the telescope itself, for example, by changing it to one that uses adaptive optics and flexible mirrors, or, as with the Hubble Space Telescope, by relocating the telescope outside the atmosphere altogether.

With this understanding of calibration in hand, we can now hope to evaluate its application to PI's. If intuitions can be calibrated in anything like the straightforward way that scientific instruments are calibrated then there is no need for serious worry about the employment of intuitions in philosophy.

First, however, we must address a well-known argument in the literature that suggests that PI's may be calibrated only in circumstances that make such calibration a waste of time. Given the benefits of calibration, it seems that one *ought* calibrate if it turns out that one *can* calibrate. Not so, argues Robert Cummins. For it turns out that the ability to calibrate PI's would be parasitic upon an understanding of the target domain that

provides *better* access to that domain than even the best calibrated intuitions. If Cummins is correct about the methods available to us for the calibration of PI's then such calibration will indeed be possible only if also a waste of time, at least from the perspective of gaining insight into the domains in question. So, before discussing the prospects for the calibration of PI we must make a case for the possible value of the project. We begin by rehearsing Cummins' argument for the pointlessness of calibrating PI's (§3), and then present our criticisms of that argument (§4). We contend that the argument is an enthymeme, and addressing the gap in his argument will reveal possible resources available to the would-be calibrator of intuitions. We then return to our central project of assessing the prospects for the calibration of PI (§5).

3. Cummins' argument against calibration

In his (1998), Cummins argues that PI's can be calibrated only if we already have trusted non-intuitive access to the facts of its intended domain of application. And we have this sort of access, it seems, only when we have a trusted, non-intuitive theory covering that domain. But with this sort of theory in hand, there is no epistemic work left for PI's to do. Therefore, Cummins concludes, calibration of philosophical intuition is possible only when it would be useless:

philosophers could have no possible use for intuitions in a context in which the relevant theory was well enough settled to form the basis of a credible calibrated test. Philosophical theory in such good shape is ready to bid the Socratic midwife farewell and strike out on its own in some other department. Philosophical intuition, therefore, is epistemologically useless, since it can be calibrated only when it is not needed (Cummins 1998: 117-8).

Cummins' argument is brief, and thus somewhat elliptical. In particular, he does not argue for the required premise that any independent theory sufficient to calibrate intuition

on the one hand, would thereby render those intuitions useless on the other. Although we think that this premise is, in fact, false, we take it to be quite suggestive. In the remainder of this section, we will try to clarify what we take to be a basic plausible connection between the required existence of independent evidence, and the possible usefulness of the calibrated intuitions. We will term this connection the *novelty requirement*. Once we see how the novelty requirement works, we will be able to see how Cummins is too quick in presuming that the calibration of PI's will violate it (§4), and from his error discover how better to apply it to make trouble for fans of intuition (§5).

The novelty requirement insists that in order for a particular instance of calibration to be epistemically worthwhile – to be able to extend our knowledge to hitherto unevaluated propositions – one's trusted independent access cannot cover the entire domain of the device being tested. If it does, there is nothing to be learned from the device, no matter how well-calibrated it is. There is, in other words, nothing which the device can teach us that we do not already know through other means.¹⁹

On Cummins' view, there must be some subset of the device's domain for which we have trusted independent access, in order for calibration to even occur. The subset forms a *certified basis* for calibration, from which an *extrapolative* inference to device's reliability over its entire domain can be made.²⁰ As we read him, Cummins thinks that when the certified basis comes from a trusted, independent theory, the extrapolative inference is unnecessary: the certified basis will cover the entirety of the domain of the

¹⁹There is no need to deny that one can still calibrate a device when the novelty requirement is not fulfilled, for various *non-epistemic* reasons. It may be useful, for example, if the calibrated device will ultimately be more convenient to use than other available means. For some sufficiently large look up table and some sufficiently small calculator there is complete overlap in domain but it is clear that the calculator is pragmatically superior. In such instances, however, calibration will not be *epistemically* worthwhile: it will not extend our knowledge.

²⁰ These notions will be discussed further in sections 4.2 and 5.

device being calibrated. As he claims, “philosophers could have no possible use for intuitions in a context in which the relevant theory *was well enough settled to form the basis of a credible calibrated test.*” (Cummins 1998: 118; emphasis added). Put in our terms, he claims that possession of a certified basis from a trusted, independent theory will always entail a violation of novelty. To reach this conclusion, Cummins assumes that an independent theory capable of providing a certified basis for calibration will also provide answers to all questions within the relevant domain. He offers no argument for this assumption, however, and it will prove the weak link around which we will organize our criticism of the argument in the following section.

4. Problems with Cummins’ argument

Although we have some sympathies with Cummins’ general conclusion, we find his argument unsatisfactory. Our concerns with Cummins’ critique of the calibration of philosophical intuitions are of two kinds. First, we think Cummins’ story ignores the possibility of an important form of calibration. Second, we think Cummins is too hasty in asserting that PI’s may only be calibrated by some theory which renders those PI’s irrelevant.

4.1. Partial Certified Bases

Recall that Cummins claims that calibration of philosophical intuitions is possible only when we can compare them to a full relevant theory of their domain. For example, intuitions about space and time would be compared to our best theory of space and time, and intuitions about representational content to our best theory of cognition. Perhaps for

these cases it is reasonable to think the independent theory provides a certified basis that exhausts the target domain, making any extrapolative inference unnecessary. Such an inference would not be required when the basis does exhaust the domain, but in such circumstances the device clearly would fail the novelty requirement. Cummins errs, however, in neglecting to discuss the possibility of a certified basis that only *partially* covers the target domain. A full evaluation of the prospects for calibration must consider the viability of partial certified bases. In the rest of this section, we present some ways in which partial certified bases might be achieved for philosophical domains.

How might a partial certified basis be of use in calibration – in particular in calibration that is meant to satisfy the novelty requirement? Any form of domain-oriented calibration requires a certified basis of information about the device’s domain. If the device is to be epistemically useful, that basis must not exhaust the domain over which we are planning to use the device. Therefore, we would have to make an extrapolative inference. We can take the partial certified basis and calibrate the device over that part of its target domain, and then it may be possible to infer to its reliability over the entire target domain, which will include propositions whose truth we do not already possess a means of evaluating. This extrapolative step may be a very standard sort of inductive inference, from an observed set to a larger population, and so we can easily establish some conditions for a successful extrapolation: we must have reason to think that the certified basis is large enough to justify the extrapolation, and that the certified basis does not represent a biased sample of the target domain.²¹ Should those

²¹ We acknowledge that sometimes this inference will be less straight-forward, particularly in cases where it can be hard to say what would or would not comprise a biased sample of a domain. But since we’re trying to be maximally fair to the prospects of calibration of PI’s, we want also to note that this inference will not *necessarily* be problematic.

conditions be met, we would thereby have evidence that the device will make reliable deliverances about matters outside of the certified basis; and thus we would have an instance of successful calibration that respects the novelty requirement.

The possibility of calibrating from a partial certified basis remains a bare possibility without some suggestion as to where a philosopher might acquire such a basis. One perhaps under-appreciated possibility here is experimental data. Assume, for example, that humans are sensitive to the fairness of economic distributions; moreover, there is good evidence that we display this sensitivity behaviorally. In concrete terms, we humans object if we regard our portion as unfair, a lack of complaint may therefore be regarded as evidence for the fairness of a distribution. In this situation, we may be able to design experiments to find out about fairness by simply operationalizing it in terms of the relevant behavior patterns. By experimentally instantiating various distributions and seeing how subjects react, we could get non-theoretical access to the target domain without relying on reports of intuitions about fairness. Perhaps research using the 'ultimatum game' would be useful here; see, e.g., Skyrms (1996); Nowak et al. (2000). Moreover, it seems unlikely that such experimental access will be possible for all cases where issues of fairness arise. We may take it therefore, at least provisionally, that this sort of access might only be partial, thereby providing a partial certified basis from which the extrapolative inference could be made (assuming, of course, that PI's can be successfully calibrated against that partial basis in the first place).

In addition, a partial certified basis might be achieved by appeal to an independent theory provided that the relationship between that theory's domain and the target device's domain is other than simple exhaustion. If we have a sufficient condition on the concept

in question, our theory of that condition might well suffice to provide a partial certified basis. For example, we expect that proved mathematical theorems are a species of knowledge. Our theory of theoremhood is basically independent of our more general intuitions about knowledge, depending only on a subset of our theory of the relations between formal symbols. If mathematical theorems were to provide us with a broad enough set of cases in which our intuitions about knowledge are congruent with our account of theoremhood, then we might have enough information to license an extrapolative inference to the reliability of such intuitions in areas outside of mathematics.

This example was just a quick illustration of how, in principle, an extrapolative form of calibration might apply within philosophy; however, we do suspect that in fact the set of mathematical theorems would provide a sample that is too badly biased for the calibration of our knowledge intuitions. Good scientific inference might be a better starting point, but the danger there would be that an account of such inferences would rest so substantively on antecedent intuitions about knowledge, that the resulting basis would not count as sufficiently independent. Demonstrating the hole in Cummins' argument requires that we articulate possibilities here – we do not claim that they are easily-actualized ones!

4.2. An overlooked form of calibration?

As Cummins conceives it, calibration requires independently acquired information about the phenomena upon which the target device is supposed to report. As he puts it, "... an invariable requirement is that there be, in at least some cases, access to the target

[that is, the phenomena upon which the device in question is supposed to report] that is independent of the instrument or procedure to be calibrated” (*ibid.* 117). We contend that one may calibrate a device (such as PI’s) *without* any such independent access. As described so far, calibration works off of some independent information about the target domain. But that is but one mode of calibration, which we will term *domain-oriented calibration*. There is at least one other mode, which we will call *instrument-oriented calibration*.

We use instrument-oriented calibration any time we create a new type of device that is sensitive to previously inaccessible facts. Our earlier story of the calibration of the telescope is of this type. As far as towers are concerned, domain-oriented calibration could suffice, since we have non-telescopic access to their shapes and configurations. But when it comes to the moons of Jupiter or the rings of Saturn, no such independent information is available. Calibrating a telescope for astronomical purposes relies more on information about the workings of the device than it does information about the facts in the target domain; hence “instrument-oriented” as opposed to “domain-oriented”.

Employing our earlier three-stage framework, we present the story of instrument-oriented calibration as follows. Begin with a device whose intended target domain may include propositions not otherwise within our epistemic reach. When such independent access is indeed unavailable, we cannot test the device by looking to see how well it reports on known phenomena in the target domain. We must evaluate the device’s accuracy by examining the device’s output on different occasions and in different situations. For example, we might evaluate the device’s accuracy by comparing its various outputs over time. If the output is consistent, then no reason has been presented

to doubt the device's accuracy.²² If the device produces inconsistent results, then the device has a problem, and we might therefore pursue a diagnosis of its difficulties.

In the case of instrument-oriented calibration, diagnosis requires identifying a way of separating good from bad outputs of the device, most typically by identifying factors that interfere with the correct functioning of the device. For example, we may identify some feature of the device's environment that co-varies with a particular kind of output from the device. Ideally we would then provide a theoretical account of the influence of the identified feature on the device. E.g., a dance troupe rehearsing in the next room may correlate with a particular kind of reading from a very sensitive seismic sensor.

However the interfering factor is diagnosed, the process of correction is straightforward. We can rehabilitate the device by adapting it with some means of inhibiting the meddlesome environmental factor. For example, if the device is highly sensitive to magnetic fields and we do not want it to be, the addition of a magnetic shield might be called for. Or we might also calibrate using restriction, using only those reports of the device that do not co-occur with the troublesome factor -- simply disregard any reports generated during dance class.

In order to make instrument-oriented calibration clearer, consider a case inspired by 18th century navigation. Navigators during this period often brought clocks aboard set to Greenwich Time, in order to calculate their longitude by comparison with local time. Now imagine an 18th century navigator in some exotic locale, with access to local time (presumably via sextant). We will equip our hypothetical navigator with a set of two

²²There is some similarity between this position and that of Tidman (1996). Just how our view differs from Tidman's and why this is crucial to the prospects for the calibration of intuition will be taken up in section 5.1.

clocks, set to GMT. For our purposes, view this complex as a device that provides information about GMT, its target domain. At some point along the journey, many weeks out from port, to the navigator's surprise and dismay he finds that the clocks report different times. In terms of the epistemology of instruments, the device has impeached itself. The presupposition of reliability no longer holds.

Domain-oriented calibration is not a possibility for our lost navigator. He does not have any other access to the target domain of GMT. But diagnosis may be possible with instrument-oriented calibration, if the navigator can identify a factor responsible for the device's incorrect functioning. Suppose after investigating the issue the navigator succeeds. With one of the clocks, he discovers one of a number of possible scenarios to be the case: it went unwound for a week, or the spring broke, or rats removed key parts of the mechanism. The navigator is now in position to account for the clocks' inconsistency. With such an account, the navigator corrects the error by using only the reports from the first clock, and thus restores the device to proper functioning and discovers his location.²³

In summary, although we sympathize with his conclusion, we find Cummins' argument overlooks a potential form of calibration (i.e., instrument-oriented calibration), and underestimates the variety of resources potentially available for domain-oriented approaches to calibrating PI's. So we cannot preemptively foreclose on the possibility of an epistemically-fruitful calibration of PI's. In order to determine whether such a calibration may in fact be possible, we will have to consider the particular facts on the

²³ Thanks to John Johnson for a very useful discussion of these matters.

ground regarding PI's, how they are produced, how they are used, and who has which ones.

5. Options for calibrating the appeal to philosophical intuitions

We argued in §1 that PI's currently stand in a state of challenge, and the result of the immediately preceding discussion is that it should be at least in principle possible for this challenge to be met by means of calibration. In this section of the paper we critically evaluate that possibility. Our view is that the prospects for some form of calibration of PI's are currently dismal, but with some glimmers of hope for the future. In §5.1 we investigate the prospects for theoretical calibration of PI's and in §5.2 we address the possibilities for calibrating PI's where there is limited independent access to the target domain.

5.1. Instrument-oriented calibration of philosophical intuitions

Analytic philosophy is rife with cases of conflicting intuitions. Such conflicts demonstrate the need for some form of correction of the intuiting mechanism. Unfortunately, the likelihood of gaining independent access to information about, say, twin-earth, is vanishingly small, so there is no way to get even a partially certified basis. Domain-oriented calibration may thus be impossible. Unfazed, defenders of intuition often suggest what we take to be a form of theoretical calibration. They claim to identify intuition-corrupting influences at work in certain circumstances, and based on this they reject those intuitions which generate the contradiction. Just such an approach is set out explicitly in Goldman and Pust (1998). This strategy is an example of instrument-

oriented calibration, taking whatever mechanism or mechanisms that produce PI's to be the instrument in question.

Instrument-oriented calibration of PI's works analogously to the case of the sea-farer and his clocks. Recall that our sailor has no independent access to Greenwich time; to correct for the contradictory output of his clocks, he has to rely on a theory of clockwork. With such a theory, he is able to identify those factors which might interfere with the proper working of the clocks. Similarly, the philosopher attempts to identify those factors which might influence the proper working of the intuition-producing mechanisms. The problem for the philosopher is that he must do this with little knowledge of how those mechanisms work. The philosopher faces the challenge of attempting instrument-oriented calibration without actually having a theory of the instrument that she hopes to calibrate.

Perhaps this problem for theoretical calibration will be made clearer through a familiar example. Going back at least to Plato and Kant, and continuing more recently through such theorists as Kohlberg (1969), many ethicists have valued judgments untainted by emotion, for fear that irrational emotion may corrupt rational apprehension of the moral good. As uncontroversial as this sounds, it actually involves unfounded assumptions about the nature of the intuition-producing mechanisms. Given that we don't know where intuitions come from, what reason do we have for thinking emotional attachment to a case would adversely affect one's intuitive mechanism? It could just as easily be the case that affect *positively* influences the accuracy of one's intuitions.²⁴ The problem is that without a rich, constructive account of philosophical intuitions – presumably as part of a theory of intuitions more generally – there isn't any reason to

²⁴ See, e.g., Gibbard (1990); Blackburn (1998); Nichols (2004).

suppose that any particular factors are corrupting those intuitions. Without an account of the device's proper functioning, it isn't possible to identify which cases are the ones where something is interfering with that proper functioning.²⁵

Despite this serious problem, there may still be hope for instrument-oriented calibration of PI's. In addressing the corrigibility of rational insight in his *In Defense of Pure Reason*, Lawrence BonJour advocates what we might take to be type of instrument-oriented calibration. For BonJour, rational insight is the ability, "simply to see or grasp or apprehend that the proposition [being considered] is *necessary*," (BonJour 1998 p. 106; italics original) which we take to be a form of intuition. Admitting the fallibility of rational insight, BonJour offers a method for identifying those intuitions which are faulty.²⁶ BonJour defends the possibility of identifying faulty intuitions through careful inspection of the intuitions in question (see his section 4.5).²⁷ This careful reflection provides the diagnosis which allows for correction, by simply rejecting those intuitions which are faulty.

BonJour's strategy seems to offer something like a form of instrument-oriented calibration, in that it proposes to separate good intuitions from bad in circumstances devoid of any independent access to their domain. But, unlike the sort of instrument-oriented calibrations discussed thus far, this strategy makes no appeal to a rich,

²⁵ Weinberg, Nichols and Stich (2001) and Goldman and Pust (1998) consider a number of ways that the intuitions relevant to philosophical practice might be delimited. Such methods may be considered as proto accounts of the circumstances under which PI's are faulty. The thing to note is that, to the extent that these notions are taken to be accounts of error, they contain implicit empirical claims about how PI's works – claims which have received up to now no formal investigation.

²⁶ Here "intuition" and BonJour's preferred "rational insight" will be used interchangeably. Note that BonJour is concerned explicitly with fending off skeptical worries, and not with calibrating intuition in a non-skeptical setting. So it should not be taken as a mark against BonJour's discussion if, as we argue here, his discussion does not provide sufficient means for such a calibration.

²⁷ BonJour actually mentions an additional form of diagnosis – assessing the coherence of the insight in question with other beliefs. The appeal to reflection however is the central form of diagnosis.

constructive account of intuition. Bonjour distinguishes between mistakes that are externally and internally correctable, where the latter require only thinking harder about or re-examining the problem. Bonjour claims it is possible to identify faulty intuitions by “appeal[ing] to the fact that many such errors, and perhaps all of them, are correctable ‘from the inside’ via *further reflection*” (Bonjour 1998 p. 117, emphasis added).

If it is true that further reflection is all that is ever needed to identify which intuitions need to be rejected—if their falseness becomes manifest through simple re-examination, then Bonjour has provided analytic philosophy with all it may need by way of addressing the troublesome evidence. But Bonjour’s primary evidence for his claims are themselves based on intuition. We grant Bonjour that many cases of divergent intuitions can be rectified by closer inspection, such as “most routine mistakes of reasoning or calculation, which yield to equally routine corrections” (*ibid.*, 116-117). But as he notes, “[t]he important question is whether *all* mistakes of apparent rational insight, all cases in which something seems necessary that is not really necessary, are mistakes that are internally correctable in this sense. Such a thesis *seems plausible enough on an intuitive basis*, though there is no apparent way to argue for it...” (*ibid.*, 117, latter emphasis added).

Bonjour’s defense of the claim that further reflection reveals which intuitions to discard is fraught with problems. The reason we can take ourselves to be able to identify faulty intuitions through reflection, according to Bonjour, is that we have the intuition that we can. But we would contend that *that* intuition is itself contested, and not one that can be relied upon at this stage. It is obviously not an intuition shared by opponents of

intuition, and indeed the authors of this paper feel compelled to report that we fail to share it, too.²⁸

Moreover, we disagree with Bonjour's assertion that this claim is not one for which only intuition could serve as evidence. Numerous empirical tests are available. One obvious sort of test is induction from past successes and failures; despite a number of local successes, the persistent divergence of intuitions across philosophers (all of whom are surely 'thinking hard' on their thought-experiments!) would seem to indicate that this is not a method with a high rate of success in weeding out incorrect PI's. One could also conduct controlled experiments in which subjects were given varying incentives; or different instructions; or that tested the subjects intrinsic motivation to think hard, known in the psychology literature as 'need for cognition'.²⁹ Just how easily and reliably one can get subjects to perform at ceiling on which sorts of intuitive judgment tasks – of more or less esoteric sorts as well – is an empirical question, and one which should not be begged with a bald assertion that it falls under intuition's purview. Surely *some* tasks can be facilitated by just thinking about them more clearly.³⁰ But there is also evidence that thinking too hard can also cloud one's judgment and 'lock in' an early mistake – especially in conditions in which the relevant information is neatly packaged, as is often the case with philosophical thought-experiments (i.e., one is supposed to think only about the contents of the scenario, and only with regard to the question of whether it is or is not a case of knowledge, free will, or the like).³¹ Which kind of influence would be the dominant influence on PI's is simply an empirical question that the philosophical

²⁸ In fact, at least some of us find the opposite intuitive: that we can't generally identify faulty intuitions via intuition itself.

²⁹ Cacioppo and Petty (1982); Cacioppo, Petty, and Kao (1984)

³⁰ Stanovich and West (2000); Epley and Gilovich (2005)

³¹ Petty *et al.* (2001).

community currently lacks any information on, and thus we cannot at this time assume that the ‘think harder’ approach will bear much calibrational fruit.

Instrument-oriented calibration is an important option for the defender of intuition. However, we have argued that instrument-oriented calibration requires a rich, constructive account of intuition of a sort that is not yet available. We suggest that the proponents of intuition begin searching – presumably in a manner that includes empirical psychological research – for the kind of robust theory that can ground theoretical calibration and render philosophers’ reliance on intuitions legitimate. Given that such an account is not currently available, we now turn to consider the case of calibration where independent access to the target domain is possible.

5.2. Domain-oriented Calibration

As we have seen, domain-oriented calibration may satisfy the novelty requirement so long as the calibrator can both (i) attain a partial certified basis on some proper subset of the target domain, and (ii) after calibrating the device on that basis, perform an extrapolative inference from the device’s new-found reliability over the subset to its reliability over the whole domain. So, if philosophers could have sufficient partial access, independent of intuition, to the truth about some domain of philosophical inquiry, then perhaps intuitions across the whole domain could be calibrated using that partial access. We now proceed to examine the prospects for the major forms of domain-oriented calibration of PI’s. Two possible bases for extrapolation have already been mentioned: experiments and independently justified theory. We will canvass these slightly more, and then further consider the calibrational resources of non-philosophical

intuitions, the judgments of experts in the domain, and demonstrated consensus of intuition. We intend to show that none of these last three options can be used as a basis for calibration. On the other hand, both experiments and independently justified theory might serve as such a basis (though we will raise some caveats).

5.2.1 Independently justified theory

In our presentation of the possibility of certification from a partial certified basis, we considered one way in which an independent theory can make possible domain-oriented calibration that satisfies the novelty requirement: namely, if the available theory covered only a proper subset of the target domain. One primary difficulty here, however, is in finding a good candidate for such an independent-from-intuitions-but-also-partially-overlapping-with-them theory in a philosophical domain. Scientific theories are an obvious place to look for such candidates. If one could take a key category from a philosophical domain to be a natural kind, and then conduct a scientific investigation on it (rather than on people's reactions to it), one might begin to get such a theory.³² But in how many cases will this solution will be applicable? For most areas of philosophy, one simply does not, and probably will never, find empirical theories that overlap the target domain. There does not look to be an empirical science of, say, metaphysical necessity, or the moral good. Even where an empirical theory speaks to a matter of philosophical concern, it will not automatically help with the project of calibration – for the theory may so radically disagree with our intuitions about the matter that it yields not a certified basis

³²A paradigm example of this approach is Hilary Kornblith's (2002) work on knowledge. While we do not necessarily endorse the details of Kornblith's account of knowledge, we encourage investigation into the viability of researching objects of philosophical inquiry as natural kinds. If such approaches work, they could provide welcome relief from reliance on evidentially suspect PI's or a robust certified basis for calibration.

for, but rather a thorough decertification of the relevant intuitions. The lesson from modern physics is not, after all, that our intuitions about the structure of space and time are reliable, but that we're better off leaving our intuitions behind.

Another possibility here is that we may have a theory about a domain that is distinct from our target domain, but is in some way substantively connected to or related to it. (Our toy example from earlier, of using our theory of proofs to help calibrate our intuitions about knowledge, might be of this sort.) For this to be a viable basis for calibration we would need an account of why the two domains are related. For example, research on animal subjects often allows us to draw conclusions, albeit somewhat cautiously, about human subjects, because we have excellent biological and evolutionary reasons to think the two domains are closely related. One might construe the *Principia Mathematica* as itself a work on this model, using results from logic and set theory in order to shed light on arithmetic – two domains which, again, we have every reason to think are intimately related. However, we are not currently aware of any analogous pairs of domains that would license such a move in analytic philosophy. But we cheerfully admit that this may be because no one has tried yet to exploit this move, and so no one has been looking seriously for such pairs, either.

5.2.2 Experiments

As we suggested in §4, intuitions (and, of course, beliefs) are not the only part of human psychology which may be sensitive to the facts about moral goodness, knowledge, and so on. To the extent that various aspects of our biology may track matters of philosophical interest, we might be able to exploit that relation in order to gain a partial

certified basis for the calibration of intuitions. Because such a relationship would not be consciously accessible, it would have to be determined experimentally. We gave the earlier example of using the ‘ultimatum game’ experiment to explore a connection between our decision-making systems and distributive justice, but that’s just one possible link. Our affective and normative systems are tightly wound around each other, and so eye dilation, skin conductance, cortisol levels, and even detectable patterns of neural activity³³ may reveal moral responses even when conscious intuition is silent or confused. Subtle linguistic phenomena may reveal further connections: Joshua Knobe has been systematically exploring a surprising asymmetry in our linguistic behavior with “intentionally”, in which it seems that we are willing to say someone performed an action intentionally when the consequences of that action are negatively valenced, even when we are not willing to say so when the consequences are positive (Knobe (2003, 2004); Knobe and Burra (forthcoming)). It does not seem that subjects are aware that the evaluations are driving their use of “intentionally”, so it may be that this effect operates independently of our ethical intuitions about the cases. As these examples illustrate, morality may be the domain most easily tapped into experimentally, because our moral cognition is so tightly interwoven with so many other psychological systems. But it is not impossible that other domains may be explored this way as well; for example, our visual systems seem to have commitments regarding objecthood, causality, and even other minds.³⁴

Unfortunately, little work has yet been done attempting to put experiments to this kind of philosophical work. And one negative consequence of that fact is that we

³³ See Greene and Haidt (2002).

³⁴ Michotte (1946/1963); Scholl and Tremoulet (2000); Castelli *et al.* (2000); Scholl (2001); Scholl, Pylyshyn, and Feldman (2001); Guski and Troje (2003).

philosophers have yet to wrestle with the murky thorny methodological issues that lurk just below the surface, such as determining the extent to which these other psychological correlations can themselves be trusted – and to what extent they may stand in need of their own calibration. So the calibration of PI’s by experiment cannot *at present* help to save the practice of intuition-driven analytic philosophy. But we think such experiments present a valuable avenue of research, both for what they could tell us about philosophical domains and for the possibility of improving philosophical methodology.

5.2.3. *Ordinary case intuitions*

George Bealer argues that we need not take the psychological evidence impeaching the accuracy of intuition seriously. He thinks that philosophical intuitions, as a type of intellectual seeming, are simply a different animal than the kinds of things that psychologists like Kahneman and Tversky study.³⁵ He is a fallibilist about intuitions: he admits that strongly felt intuitions can be false, e.g., the naive comprehension axiom. However, he thinks that we have evidence that intuitions are “for-the-most-part” correct (Bealer (1996)). He bases his trust of PI’s in spite of these mistakes largely on the “on-balance consistency of our elementary concrete-case intuitions” (Bealer (2000)). It seems, therefore, that he takes such quotidian intuitions about cases to be of a piece with the sorts of more abstruse intuitions more typically appealed to by philosophers. For such an argument to work, he needs the capacity that allows us to tell us which objects are chairs to be fundamentally the same capacity that allows us to tell which actions are morally wrong and which situations are examples of knowledge; and he needs to be able

³⁵ We are not sure what he takes the difference to be. He does distinguish *sensory* intuitions from *rational* intuitions. But he takes intuitions about logic to be paradigm examples of the latter, so Wason and Johnson-Laird’s results about reasoning would seem to bear on rational intuitions rather directly.

to make a clear projection from their patterns of response on the ordinary cases out to the more esoteric cases where the philosophical work gets done.

First, we have reason to worry about the viability of extrapolation from easy cases to hard ones. For example, our intuitions about the comparative cardinality of finite sets – e.g., that taking every other member from a set yields a set half as big – seem unproblematic, and we do not doubt they are generally well-calibrated. But we do doubt that this calibration extends into the area of infinitary mathematics. We know that many principles that hold in finite domains do not hold in infinite ones, and so we may legitimately object that our confidence for judgments about n 's cannot be extended to our judgments about \aleph 's. This, of course, does not stop us from making use of our intuitions in inappropriate ways. For example, BonJour appeals to the intuition that there are half as many even integers as there are integers (BonJour 1998, 209n24). Similar situations arise if we look to move from arithmetic to trickier areas of mathematics. Our wild intuitive success with addition and multiplication cannot justify any expectation that we would meet with similar success in probability theory, where in fact people generally have appallingly bad intuitions. So it does not seem that any general licensing of the difficult by the easy will be forthcoming.

Moreover, recent work in psychology raises worries about the very idea that our performance in one intuitive area can shed light on our performance in others. For many psychologists and philosophers³⁶ have come to believe that we do not have one big domain-general intuition system, but rather a number of distinct domain-specific

³⁶ E.g., Samuels (1998, 2000); various papers in Hirschfeld and Gelman (1994); Sperber (2002); various papers in Carruthers and Chamberlain (2000); Carruthers (2004, 2006); Barrett (2005). Note that the relevant notion of modularity here is only that of functional specialization, and need not include the full suite of Fodorian modular characteristics like automaticity or cognitive impenetrability (see Barrett and Kurzban (forthcoming)).

mechanisms subserving our cognition in these areas—what Steven Pinker has colorfully called “the mind as a Swiss Army knife” (Pinker (1997)). If such ‘massive modularity’ theories are correct, then our verifiable successes in one domain (like the everyday world of middle-sized dry goods) would, at best, be evidence that the intuition-producing mechanism subserving *that* domain can be trusted. But just as evidence that your thermometer is reliable is no evidence that your barometer is working, so too would evidence that the mechanism for one domain is reliable fail to be evidence for the reliability of a different intuition-producing mechanism operating in a different domain. This is not to say that we are committed to any form of massive modularity turning out true, and we recognize that it is still a matter of significant debate in cognitive science. Rather, our objection is that in the absence of compelling reasons to think that all intuitions – both ordinary and esoteric – arise from a unified psychological source, no such easy extrapolative step will be justifiable.

5.2.4 Consensus

Perhaps in some target domain we could use the intuitions, even the esoteric ones, that everyone agrees upon as a certified basis for calibration. This method would first involve finding such a consensus. Some of the results cited in §1 give us reason to think that we may be surprised by the extent of divergence in intuitions across different groups. But it is plausible to think that we will find agreement on many cases in a domain – though it will be a decidedly empirical matter to discern where those points of agreement lie. Surely there are some examples of, say, moral rightness which almost no reasonable person will disagree with.

One might well ask whether increased reliability is automatically to be expected in such cases, or whether some further defense of the reliability of consensus intuitions may legitimately be demanded. Regardless of the answer to that question, using intuitions for which there is consensus makes difficult the extrapolative inference that is necessary, if domain-oriented calibration is to comply with the novelty requirement (and thus have any hope of side-stepping Cummins' objections to it). Recall that calibration involves making the device's reports fit with what is known through the certified basis. Such a comparison can only be done in cases covered by both the device's reports and the certified basis – that is, the cases on which everyone agrees. But with consensus, the certified basis consists of cases where the reports of everyone's device agree. This means that the device's reports will trivially fit with the certified basis, since they are defined to be extensionally identical.³⁷ The result is that such a calibration does not permit the extrapolation to intuitions' accuracy on hard cases.

For example, suppose that a Kantian and a utilitarian were attempting to calibrate their intuitions against consensus. For both, their agreed-upon intuitions are thus trivially certified. But there is no way to extrapolate out from the consensus cases in order to settle which is correct where they disagree. A certified basis made of easy-case intuitions doesn't allow the Kantian or the utilitarian to extrapolate the accuracy of their respective hard-case intuitions, since each has an equal claim to comports well with the basis.

³⁷ Cummins raises the same worry for calibration by consensus (Cummins 1998: 117). Indeed a number of the points that follow regarding calibration by consensus are in agreement with Cummins' analysis of the shortcomings of such an approach.

Consensus cases, by their very nature, will probably fail to select one side over the other³⁸.

We can imagine one sort of scenario in which consensus cases might do the trick: if we restrict our attention only to consensus cases, it may turn out that one theory does a much better job than its rivals of systematizing those cases. Somehow it may turn out that a version of one theory is able to systematize them in a natural and simple way, whereas the other theories must give more contorted versions of themselves in order to capture those intuitions. In such a situation, we may do just fine by calibrating using restriction. But it's impossible to know now whether any such situations could ever be real. We can't tell whether this is a live possibility or not, not least of all because *we just don't know what the consensus cases really are*. This is not a matter that is answerable without significant empirical work, which has generally not yet been done.³⁹

5.2.5 Experts

It might be suggested that we could find out which propositions the experts in a given area of philosophy endorse, and let those form the basis for calibration.⁴⁰ Experts presumably have access to the truths of their area of expertise, so bringing our intuitions in line with experts' endorsements looks like an attractive method of calibration. An examination of the possible evidential sources for experts' judgments, however, reveals

³⁸ A similar point is made regarding the analysis of justification by William Alston in his paper 'Epistemic Desiderata' p. 537 in particular (Alston 1993).

³⁹ Note that these arguments also apply to the appeal to ordinary case intuitions considered in the previous section. If we grant that there is consensus on all the ordinary case intuitions, then the total set of consensus intuitions contains all the ordinary case intuitions as a subset. And so if the total set of consensus intuitions is not up to the job of calibration, it will be unsurprising that its subset of ordinary case intuitions will not meet the fan of intuition's needs, either.

⁴⁰ Despite our disagreements articulated earlier, we are indebted greatly to Cummins (1998) for the discussion in this section.

problems with using them as the basis for calibration. Someone endorsing this move faces a dilemma: experts base their judgments either on their own intuitions, or on something else. If experts base their judgments on their philosophical intuitions about the target domain, then we may legitimately ask how *those* intuitions were themselves calibrated. And the answer may, again, be based either on PI's, or on something else. Given the problem at hand – how to calibrate PI's, given that the practice of appeal to intuitions is now in a state of challenge – this cycle cannot stop at more PI's. For in that case, our answer to the problem of how we might calibrate would presuppose that the problem was already solved.

So at some point in this iteration, experts must base their judgments on something other than intuitions, such as independent theories or experimental data. But if this were so, then there would be no need to appeal to expert intuitions at all, since we might just as well calibrate intuitions off of whatever it is that the experts were basing their judgments on. To play on Wittgenstein, the appeal to expert intuitions in calibration would be at best an unnecessary shuffle.⁴¹

6. Conclusion:

Pessimism for the Armchair Present, Guarded Optimism for the Scientific Future

⁴¹ There is a certain structural similarity between this argument and Cummins' main argument as discussed above, so the reader may wonder whether this argument might be vulnerable to the same kinds of objections we lodged against Cummins. However, a key difference is that Cummins is there aiming to show that PI's can *never* be calibrated except when it would not be useful to do so; whereas here we are trying to show that PI's cannot be calibrated *by appeal to experts*.

We would also note that this entire discussion is setting aside the fundamental question of deciding who counts as a relevant philosophical expert. We can well imagine those who have studied closely the texts of Descartes or Russell claiming an expert status on the topic of knowledge, for example, that students of Peirce or Austin might wish to deny. Who gets to count as an Officially Licensed Intuitor? If the epistemic status of philosophical intuitions turns out to be a matter of who can win the relevant political battles over credentialing, then the postmodernists will truly have won.

We may now offer an answer to the question as to whether PI's may in fact admit of calibration. Or, since the short answer is "no, at least not right now", a better question to address would be: under what conditions might PI's in the future admit of calibration? Our explorations so far reveal a number of ways in which future developments might allow for some prospects of calibration.

With regard to domain-oriented calibration, we have argued that expert intuitions and non-philosophical intuitions will be of no help, but the other three possible sources of a partial certified basis all offer at least some potential resources:

Consensus intuitions – If we had a better sense of where different people agreed or disagreed in their intuitions, we would at least thereby come to have a better sense of which particular PI's have been directly impeached, and begin devising ways to adapt our practices to minimize our reliance on them. For reasons stated earlier, we do not think that restricting ourselves to consensus intuitions will be enough of a fix to give intuition-deploying philosophers all they might want, but it is a suggestive start. Moreover, a full account of who has what intuitions under which circumstances would be tremendously valuable in developing a constructive account of the intuition-producing mechanisms (see below).

Independent theory – Using an independent theory that either overlaps or is substantially related to a philosophical domain would be a fairly straightforward matter, *if* we could locate more such theories. Hilary Kornblith's recent attempts at a thoroughgoing naturalism about knowledge as a natural kind (see note 30 above) present one such

possibility, though we are somewhat skeptical both of the account itself, and of the likelihood that it can be extended to other concepts.

We should note that philosophers do not necessarily have to turn to scientists in order to find fairly intuition-free theories of their target domains. Edward Craig's recent account of knowledge in his (1990) does not depend centrally on esoteric intuitions as evidence, in that it appeals more to the role that a concept like knowledge can play in our overall cognitive economy. But his method has not yet been pursued much further by epistemologists.⁴² And we are of course also open to the development of still further novel, non-intuitive methods within philosophy.

Experimental approaches to philosophical domains – Of all the various sources for a partial certified basis for domain-oriented calibration of PI's, we are most sanguine about the possibility of developing experimental, non-intuitive means of access. Part of our optimism here derives surely from the fact that it is *so* very underexplored, but also we expect that there is a wealth of psychological results already out there ripe for philosophical application, and it just remains for philosophers to get sufficiently motivated to bring themselves up to speed on the relevant literatures. Better still, the philosophers may learn how to generate the relevant empirical results themselves.⁴³

And finally, we have our most favorite prospect for any future calibration of PI's:

⁴²Though see Neta (forthcoming) and Weinberg (forthcoming) for some very recent interest in, and applications of, Craig's ideas.

⁴³As is indeed already happening, in a nascent form; see references in §1.2 and §5.2.2 above.

Constructive account of intuition-producing mechanisms – Although we have some limited hope for the prospects for domain-oriented calibration, we are rather more optimistic about instrument-oriented calibration. As we learn more about the different sources of intuitions, discover more of the algorithms and heuristics underlying their production, explore the boundaries of intuitive agreements and disagreements – as our scientific investigation of intuitions extends our knowledge past their veil of introspective opacity, then we may attain much more clarity as to the circumstances under which they can be trusted, and those where they cannot. (An excellent starting place might be Stanovich’s (1999) work on determining which subjects have superior rational judgment in general.) At a minimum such developments should enable some limited forms of restrictive correction. At best, such an understanding of intuition-producing mechanisms may even allow for some degree of rehabilitative correction, such that we may, individually or as a profession, become better intuitors than we are currently. At worst, we will have to radically restrict our appeals to intuitions, but we will have acquired an interesting scientific psychological theory as a booby prize.

Our investigations leave philosophers in a something of a bad-news, good-news situation. The first bad news is that calibration cannot at present help philosophers respond to the growing challenges to intuition-based methods. The good news is that calibration *may* be of some use to the analytic philosopher in the future, if only philosophers can successfully pursue one or more of the avenues listed above. There is a further bad-news punch line here, however: in order to take this route of saving the practice of using intuitions, we must either turn to science for serious help in exploring

the extent of intuitive consensus and, moreover, the structure of our intuition-producing mechanisms; or develop other alternative methods within philosophy that do not themselves rely on intuition. Calibration offers philosophy no way to save its beloved intuitions without, to some significant degree, weaning philosophy off of its reliance on those very intuitions.

Bibliography

- Alston, William. (1993). "Epistemic Desiderata", *Philosophy and Phenomenological Research* LIII (3): 527- 551.
- Bealer, George. (1996). "On the possibility of philosophical knowledge," *Philosophical Perspectives* 10: 1-34.
- (1998). "Intuition and the Autonomy of Philosophy". In DePaul and Ramsey.
- (2000). "A theory of the a priori," *Pacific Philosophical Quarterly* 81:1-30.
- Bechtel, W. (2000). From imaging to believing: Epistemic issue in generating biological data. In R. Creath and J. Maienschein (Eds.), *Epistemology and Biology*, 138-163. Cambridge: Cambridge University Press.
- Blackburn, S. (1998) *Ruling Passions*. Oxford: OUP.
- Boghossian, P. and Peacocke, C., eds. (2001). *New Essays on the A Priori*. Oxford: Clarendon Press.
- BonJour, Laurence (1985) *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- (1998), *In Defense of Pure Reason*. Cambridge University Press: Cambridge U.K.
- Castelli, F., Happé, F., Frith, U. and Frith, C. 2000. Movement and mind: A functional imaging study of perception and interpretation of complex intentional movement patterns. *Neuroimage*, 12, 314 – 25.
- Chang, Hasok (2004), *Inventing Temperature: Measurement and Scientific Progress*. Oxford; Oxford University Press.
- Craig, Edward (1990). *Knowledge and the State of Nature*. Oxford: Clarendon Press.
- Cummins, Robert (1998) "Reflection on Reflective Equilibrium," in DePaul and Ramsey, 113 – 128.
- DePaul, M. and Ramsey, W. (eds.) *Rethinking Intuition: The Psychology of Intuition and its Role in Philosophical Inquiry*. Rowman and Littlefield.
- Descartes, R. (1984) *Meditations on First Philosophy* in John Cottingham, Robert Stoothoff, Dugald Murdoch, (translators) *The Philosophical Writings of Descartes*. Cambridge University Press.
- Dretske, Fred (1970), "Epistemic Operators", *Journal of Philosophy*, 67: 1007-23.

- Evans, J. and Over, D. 1996. *Rationality and reasoning*. Hove: Psychology Press.
- Fix, John D. (1995) *Astronomy: Journey to the Cosmic Frontier*. St. Louis: Mosby.
- Gettier, E. 1963. Is Justified True Belief Knowledge?, *Analysis* 23: 121-123.
- Gibbard, A. (1990) *Wise Choices, Apt Feelings: A Theory of Normative Judgment* Oxford: OUP.
- Goldman, A. and Pust, J. 1998. Philosophical Theory and Intuitional Evidence. In Michael R. DePaul and William Ramsey (eds.) *Rethinking Intuition: The Psychology of Intuition and its Role in Philosophical Inquiry*. Rowman and Littlefield: MD: pp. 179-200.
- Greene, J. and Haidt, J. 2002. How (and where) does moral judgment work? *Trends in Cognitive Sciences*, 6(12), 517-523.
- Guski, R. and Troje, N. 2003. Audiovisual phenomenal causality. *Perception & Psychophysics*, 65(5), 789 – 800.
- Hirschfeld, Lawrence A. and Susan A. Gelman (eds.) (1994) *Mapping the mind: Domain specificity in cognition and culture*. New York, NY: Cambridge University Press.
- Jackman, H. (forthcoming). Intuitions and semantic theory. *Metaphilosophy*.
- Kahneman, D., P. Slovic, and A. Tversky (eds.) (1982) *Judgment under uncertainty: Heuristics and biases*. New York, NY: Cambridge University Press.
- Katz, J. (1998). *Realistic Rationalism*.
- Kohlberg, L. (1969). Stage and sequence: The cognitive-developmental approach to socialization. In D. A. Goslin (Ed.), *Handbook of socialization theory and research*, 347-480. Chicago: Rand McNally.
- Knobe, J. (2003). Intentional action and side effects in ordinary language. *Analysis* 63, 190 – 193.
- (2004). Intention, intentional action, and moral considerations. *Analysis* 64, 181 – 187.
- and Burra, A. (forthcoming). Intention and Intentional Action: A Cross-Cultural Study. *Journal of Culture and Cognition*.
- Kornblith, Hilary (2002) *Knowledge and its Place in Nature*. Oxford University Press: Oxford U.K.

- Lackey, J. (forthcoming) Learning from words. In *Philosophy and Phenomenological Research*.
- Lehrer, K. (1971) How reasons give us knowledge, or, the case of the gypsy lawyer. *Journal of Philosophy*, 68(10), 311-3.
- and Cohen, S. (1983). Justification, truth, and coherence. *Synthese* 55, 191 – 207.
- Levin, J. (2004). The evidential status of philosophical intuitions. *Philosophical Studies*, 121(3), 193-224.
- Loftus, E. (2003). Our changeable memories: legal and practical implications. *Nature Reviews: Neuroscience*. 4(3), 231-4.
- Machery, E., Mallon, R., Nichols, S., and Stich, S. P., (2004), “Semantics, Cross-cultural Style,” *Cognition*, 92(3): B1-B12.
- Miller, R. (2000). “Without Intuitions”. *Metaphilosophy*, 31, 231 – 50.
- Nahmias, E., Morris, S., Nadelhoffer, T, and Turner, J. (forthcoming). “Surveying Freedom: Folk Intuitions about Free Will and Responsibility”. *Philosophical Psychology*.
- Neta, R. (forthcoming) “Epistemology Factualized: New Contractarian Foundations for Epistemology”. *Synthese*.
- Nichols, S. (2004). *Sentimental Rules*. Oxford: Oxford University Press.
- and Knobe, J. (forthcoming) Moral responsibility and determinism: The cognitive science of folk intuitions. In *Nous*.
- , Stich, S., and Weinberg, J. (2003). “Metaskepticism: Meditations in Ethno-Epistemology.” In *The Sceptics*, ed. S. Luper. Burlington, VT: Ashgate.
- Nisbett, Richard and Ross, Lee (1980). *Human Inference: Strategies and Shortcomings of Social Judgment*. Prentice-Hall.
- Nowak, M., Page, K., and Sigmund, K. 2000. Fairness versus reason in the ultimatum game. *Science*, 289(5485), 1773-1775.
- Peacocke, C. (2001). Explaining the a priori: The programme of moderate rationalism. In Boghossian & Peacocke (2001), 255 – 85.
- Pinker, S. (1997) *How the Mind Works*. Penguin Press.
- Pust, J. (2000). *Intuitions as evidence*. New York and London: Garland Publishing

- Samuels, R. (1998). Evolutionary psychology and the massive modularity hypothesis. *British Journal of Philosophy of Science*, 49: 575-602.
- (2000). Massively modular minds: evolutionary psychology and cognitive architecture. In P. Carruthers and A. Chamberlain (eds.), *Evolution and the Human Mind*. Cambridge University Press.
- Sosa, E. (1996) Rational intuition: Bealer on its nature and epistemic status. *Philosophical Studies*, 81(2-3), 151-62.
- (1998). Minimal intuition. In Depaul & Ramsey (1998).
- Sperber, D. (2002). In Defense Of Massive Modularity. In Dupoux, E. *Language, Brain and Cognitive Development: Essays in Honor of Jacques Mehler*. Cambridge, Mass. MIT Press. 47-57.
- Skyrms, Brian (1996). *Evolution of the Social Contract*. Cambridge: Cambridge University Press.
- Stanovich, K. (1999). *Who is Rational?: Studies of Individual Differences in Reasoning*. Mahwah, NJ: Lawrence Erlbaum Associates, Inc.
- Swain, S., Alexander, J. and Weinberg, J. 2006. The instability of philosophical intuitions: Running hot and cold on Truetemp. Presentation to the annual meeting of the Society for Philosophy and Psychology.
- Tidman, Paul (1996), "The Justification of the A Priori Intuitions", *Philosophy and Phenomenological Research* 56:1 pp. 161 – 171.
- Weatherson, B. (2003) What good are counterexamples? *Philosophical Studies*, 115: 1-31.
- Weinberg, J. (forthcoming) What's Epistemology For?: The Case for Neopragmatism in Normative Metaepistemology. In S. Hetherington (ed.), *Epistemology Futures*. Oxford: Oxford University Press.
- , Nichols, S., & Stich, S. (2001). Normativity and epistemic intuitions. *Philosophical Topics*, 29(1&2), 429–459.
- Williamson, T. (2004a). "Philosophical 'Intuitions' and Skepticism about Judgement". *Dialectica*, 58(1), 109 – 53.
- (2004b) "The Presidential Address I—Armchair Philosophy, Metaphysical Modality and Counterfactual Thinking. *Proceedings of the Aristotelian Society* 105 (1), 1-23.