

The Challenge of Sticking with Intuitions through Thick and Thin

Jonathan M. Weinberg (University of Arizona) and Joshua Alexander (Siena College)¹

Our disagreement about the nature and epistemic authority of intuitions is at root a battle for the preservation of philosophy as an autonomous field of inquiry.

- Gary Gutting (1998)

0. Abstract

Philosophical discussions often involve appeals to verdicts about particular cases, sometimes actual, more often hypothetical, and usually with little or no substantive argument in their defense. Philosophers – including ones on both sides of debates over the standing of this practice – have very often called the basis for such appeals “intuitions”. But, what might such “intuitions” be, such that they could legitimately serve these purposes? Answers vary, ranging from “thin” conceptions that identify intuitions as merely instances of some fairly generic and epistemologically uncontroversial category of mental states or episodes to “thick” conceptions that add to this thin base certain semantic, phenomenological, etiological, or methodological conditions. While thin conceptions have received several high-profile endorsements in recent years, thick conceptions have become increasingly popular, in part because they seem to offer some way of responding to recent empirical challenges to our intuition-deploying practices that thin conceptions do not. But this response is not without its costs. Thick conceptions turn out to have their own methodological problems – some actually fail to properly immunize our intuition-deploying practices from the kinds of problems raised by recent empirical challenges, others expose our intuition-deploying practices to different kinds of empirical challenges, and still others leave us in the methodologically untenable position of being unable to determine when anyone is doing philosophy correctly.

1. Introduction

Philosophical discussions often involve appeals to verdicts about particular cases, sometimes actual, more often hypothetical, and usually with little or no substantive argument in their defense. Philosophers – including ones on both sides of debates over the standing of this practice – have very often called the basis for such appeals “intuitions”. But, what might such “intuitions” be, such that they could legitimately serve such purposes? Answers vary, ranging from “thin” conceptions that identify intuitions as merely instances of some fairly generic and epistemologically uncontroversial category of mental states or episodes to “thick” conceptions that add to this thin base certain semantic, phenomenological, etiological, or methodological conditions.

Thick conceptions have become increasingly popular in recent years, in part because they offer a

¹ Authorship is equal. We would like to thank members of the philosophy departments at Florida International University, The University of Illinois at Urbana-Champaign, York University, Marist College, and Siena College, as well as participants at the 2011 Metro Experimental Research Group’s Experimental Philosophy Conference, the Experimental Philosophy Society Meeting at the 2011 Pacific Division Meeting of the American Philosophical Association, and the 2011 Joint Meeting of the Society for Philosophy and Psychology and the European Society for Philosophy and Psychology for valuable feedback on earlier versions of this paper.

way of responding to recent empirical challenges to our intuition-deploying practices. The basic idea is to insulate our intuition-deploying practices from these challenges by arguing that, whatever it is that experimental philosophers have been studying, they haven't been studying the right kinds of mental states or episodes – that is, the kinds of mental states or episodes at play in our intuition-deploying practices. This puts philosophers in the position of being able to argue that the results of these studies shed little insight on the standing of our actual intuition-deploying practices.

Although adopting a thick conception of philosophical intuitions may thus seem to put philosophers in a position to defend against recent empirical challenges, we will contend the apparent strength of thick conceptions is only apparent. Thick conceptions turn out to have their own methodological problems – some actually fail to properly immunize our intuitional deploying practices from the kinds of problems raised by recent empirical challenges, others expose our intuition-deploying practices to different kinds of empirical challenges, and still others leave us in the methodologically untenable position of being unable to determine when anyone is doing philosophy correctly.

2. Philosophical Intuitions

According to one rather influential way of thinking about philosophical practice, philosophical theories are measured (at least in part) by our philosophical intuitions. On this view, we advance philosophical theories on the basis of their ability to explain our philosophical intuitions, defend their truth on the basis of their overall agreement with our philosophical intuitions, and justify our philosophical beliefs on the basis of their accordance with our philosophical intuitions. According to this way of thinking about philosophy, the role and corresponding epistemic status of philosophical intuitions is similar to the role and corresponding epistemic status of perceptions (see, e.g., Bealer 1998 and Sosa 1998). Like perceptions, philosophical intuitions provide a non-inferential, defeasible justificatory foundation in at least the following way: a person may appeal to an intuition as evidence without having to provide further evidence for the intuition.²

Quite possibly the most famous example of this way of thinking about philosophical practice is Edmund Gettier's (1963) use of two thought-experiments to help show that a person's justified true belief need not count as knowledge. Each of Gettier's thought-experiments involve Smith who has deduced a true belief q from a justified false belief that p and, on that basis, formed a justified true belief that q . According to Gettier, despite now having a justified true belief that q , Smith doesn't know that q . Of course, we aren't just supposed to take Gettier's word on this. Instead, it's supposed to seem obvious to us when we consider each case that Smith doesn't know that q . And, these shared judgments – our philosophical intuitions about the two hypothetical cases – are what are supposed to count as sufficient evidence that sometimes a person's justified true belief doesn't count as knowledge.³

² This way of thinking about things leaves open whether treating intuitions (or perceptions) as evidence involves treating psychological states (or propositions about psychological states) as evidence or treating the *contents* of those psychological states as evidence.

³ We are starting off uncommitted about the precise psychological nature of intuitions, but for ease of exposition (and following much extant practice) we will use terms like “intuition”, “intellectual seeming”, and “judgment” more or less interchangeably.

Examples like this abound. In fact, this way of thinking about the epistemic status of philosophical intuitions is so widespread in contemporary philosophy that George Bealer (1998) has called it part of our “our standard justificatory procedure”.⁴ Other philosophers from a wide range of metaphilosophical temperaments have gone further than Bealer, claiming not only that philosophical intuitions are part of our standard justificatory procedure, but also that this fact is part of what defines uniquely *philosophical* methodology. For example, Janet Levin (2005) writes,

This procedure of rejecting or modifying theses in the face of intuitively convincing counterexamples has been characteristic, perhaps definitive, of philosophical argumentation throughout its history.

And Alvin Goldman (2007) writes,

One thing that distinguishes philosophical methodology from the methodology of the sciences is its extensive and avowed reliance on intuition. Especially when philosophers are engaged in philosophical “analysis”, they often get preoccupied with intuitions.⁵

But, what are philosophical intuitions? Answers vary, often reflecting different attitudes about how we should go about trying to answer this question. Some philosophers encourage us to consult our intuitions about intuitions; others recommend paying special attention to what they introspectively seem to be from the first-person point of view; still others advise looking closely at their role in philosophical practice. These different approaches have produced a range of conceptions, ranging from what we will call “thin” conceptions, which identify intuitions as merely instances of some fairly generic and epistemologically uncontroversial category of mental states or episodes, to “thick” conceptions, which add to this thin base certain semantic, phenomenological, etiological, or methodological conditions.

Let’s start with the thinnest conception, according to which philosophical intuitions are simply instances of some fairly generic and epistemologically uncontroversial category of mental states. Timothy Williamson (2007) has been the most visible and vigorous proponent of this way of thinking about philosophical intuitions, arguing that more substantive conceptions leave out much of what is appealed to as evidence when philosophers claim “the intuitiveness of some theories as a virtue and the counterintuitiveness of others as a vice” (2007, 217).⁶ The idea is that

⁴ Bealer not only makes the descriptive, sociological claim that philosophical intuitions are in fact a standard part of our justificatory practices. He also argues for the normative claim that they *must* be. The normative claim is part of his famous argument for the incoherence of empiricism. The basic idea is that empiricists cannot defend their own methodological commitments using any set of justificatory resources that don’t include philosophical intuitions.

⁵ One of the most interesting aspects of Goldman’s view about philosophy’s intuition-deploying practices is his subsequent endorsement of a view he calls “mentalism”. According to this view, our philosophical intuitions tell us less about the world (or at least the non-psychological world) than they do about the ways in which we think about the world.

⁶ David Lewis (1983), Peter van Inwagen (1997), Jonathan Ichikawa (manuscript), Herman Cappelen (forthcoming), and Wesley Buckwalter and Stephen Stich (manuscript) also endorse this kind of maximal inclusiveness. Williamson (2007) actually suggests that the practice of appealing to philosophical intuitions as evidence rests on a misunderstanding of the nature of philosophical evidence, arguing, “Philosophers might be better off not using the word ‘intuition’ and its cognates. Their main current function is not to answer questions about the nature of evidence on offer but to fudge them, by appearing to provide answers without really doing so” (2007, 220). See Alexander

substantive conceptions discriminate, and in doing so fail to reflect the wide variety of things that answer to the name ‘intuition’ in philosophical practice. Williamson encourages a more tolerant view, according to which intuitions are (at most) just judgments or inclinations to judge. This kind of maximal inclusiveness promises to not only better reflect what it is that philosophers actually do, it might also seem to insulate philosophical intuitions from skeptical challenge. The more that counts as intuitional evidence, the more like a dead-end global intuitional skepticism becomes.

The advantages of maximal inclusiveness, however, come at the cost of *evidential diversity*. In the attempt to treat as much as possible as intuitional evidence, we risk discounting important differences between various kinds (e.g., memorial, introspective, perceptual, inferential, testimonial, and even more narrowly construed *intuitional* kinds) of philosophical evidence – differences that figure significantly into questions of reliability, defeasibility, and conflict resolution. Different kinds of evidence have different strengths, and different weaknesses, and a responsible methodology for any field of inquiry should attend closely to those differences, not paper them over. But, even setting aside concerns about a lack of evidential diversity, there is another reason to worry that maximal inclusiveness has dangerous vices to go along with its virtues. As we will see, many contemporary philosophers want to be able to say that not all intuitions *are* created equal – that is, that not everything that is appealed to as intuitional evidence is *genuine* intuitional evidence. And, we might worry that maximally inclusive conceptions leave us without the resources needed to mark this difference.

In light of these sorts of considerations motivating some degree of intra-intuitive discrimination, many philosophers have been tempted to adopt a thicker, more discriminating conception of intuitional evidence. For a time, the most popular way to do this was to argue that genuine philosophical intuitions have a special kind of *propositional content* or distinctive *phenomenology*. Some philosophers argued that philosophical intuitions must have *modally strong* propositional content (e.g., BonJour 1998); other philosophers argued that philosophical intuitions must have *abstract* propositional content (e.g., Sosa 1998); and still others argued that philosophical intuitions must have the *appearance of necessity* (e.g., Bealer 1998). The general idea was to locate some distinctive feature in the content or appearance of genuine philosophical intuitions that could be used to pick them out of the crowd.

In recent years, fashion has shifted away from attempts to identify genuine philosophical intuitions on the basis of specific semantic or phenomenological characteristics, although as we will see phenomenological characterizations are beginning to see a revival of sorts. More popular now are approaches that focus on etiology and methodology – that is, on where philosophical intuitions come from and what we do with them once they are here. Kirk Ludwig (2007) provides an admirably clear and straightforward articulation of the etiological approach, arguing that genuine philosophical intuitions are judgments made only on the basis of *conceptual competence*:⁷

(2010) for a reply to Williamson’s argument against the philosophical significance of philosophical intuitions.

⁷ Others in the etiological camp use somewhat different language, suggesting that genuine philosophical intuitions are judgments (or inclinations to judge) based on *conceptual understanding* (e.g., Sosa 1998), *determinate concept possession* (e.g., Bealer 1998), or what *competent users* of the relevant concepts could say if they considered the

For terminological clarity, I will use “intuition” to mean an occurrent judgment formed solely on the basis of competence in the concepts involved in response to a question about a scenario, or simply an occurrent judgment formed solely on the basis of competence in the concepts involved in it (in response, we might say, to the null scenario). Intuitions in this sense are what we seek to elicit in response to questions about scenarios in thought experiments. (135-136)

Ludwig goes on to argue that not all judgments generated in response to philosophical thought experiments express our conceptual competence. Some, he suggests, are

judgments based on well-entrenched empirical beliefs, judgments based on what would be standardly implicated by the sentence that expresses the judgment, judgments based on reading more or less into the scenario than is intended by the experimenter, judgments based on information carried by the linguistic vehicle for it as opposed to what it means, and judgments made in circumstances in which none should be made because the idealization that the terms involved are semantically complete breaks down dramatically, as in borderline cases of “bald” or “heap” or “person”, or, generally speaking, judgments whose etiology is not solely the competence of the subject in use of the concepts in response to the scenario as described. (137-138)

Since not all judgments generated in response to philosophical thought experiments express our conceptual competence, not all judgments are genuine philosophical intuitions. And, our task becomes to “try to factor out the contribution of competencies from the other factors” (144-145) when determining whether a particular response counts as a genuine philosophical intuition (see, also, BonJour 1998, Sosa 1998, and Henderson & Horgan 2001). This last point suggests a place where the etiological gives way to the methodological. More is needed than simply a call for proper etiological pedigree. What is needed is some way of determining when this call has been met. Antti Kauppinen (2007) suggests that what we need is *reflection*: careful examination and evaluation not only of our judgments about hypothetical cases, but also of the cases themselves and what influence philosophical commitments might have on what details we find relevant in those cases. Only this kind of careful, paradigmatically *philosophical* reflection will help us pick out genuine philosophical intuitions.

3. The Challenge of Sticking with Thin Conceptions

Discrimination always requires some defense, something that shows that our views are *discriminating* rather than *discriminatory*, and some contemporary philosophers defend their turn away from maximal inclusiveness by suggesting that less inclusive conceptions of philosophical intuition seem better able to protect our intuition-deploying practices from a recent empirical challenge to those practices.⁸ Here’s the empirical challenge – which we have elsewhere called

case in *sufficiently ideal conditions* where their judgments were *influenced only by semantic considerations* (e.g., Kauppinen 2007).

⁸ There have been other challenges to our intuition-deploying practices in recent years. Hilary Kornblith (1998) has argued that, while philosophical intuitions count as evidence for something, they don’t count as evidence for anything that out to be of interest to philosophers. Robert Cummins (1998) has argued that philosophical intuitions cannot be treated as evidence because we cannot determine antecedently and independently whether or not they are reliable guides to the truth. And, Michael Devitt (1994) and Catherine Elgin (1996) have argued that philosophical

the *restrictionist challenge* – in a nutshell.⁹ Our intuition-deploying practices are in good standing only to the extent to which philosophical intuitions are *appropriately* sensitive – sensitive only to those things that are relevant (or, at least, those things that philosophers commonly take to be relevant) to the truth or falsity of the claims for which the evidence is supposed to provide evidence. The problem is that recent empirical work suggests that at least some philosophical intuitions are *problematically* sensitive – sensitive to things (e.g., ethnicity, gender, affectivity, and presentation order) that haven't traditionally been thought to be relevant to the truth or falsity of philosophical claims (see, for example, Weinberg et al. 2001, Buckwalter & Stich manuscript, Nichols & Knobe 2007, and Swain et al. 2008). This problematic sensitivity is made worse by the fact that we currently lack the means to either correct for it or predict from the armchair when or where else it will appear. We find ourselves in the untenable epistemic position of suspecting that *some* intuitional evidence isn't trustworthy without being able to reliably predict *what* intuitional evidence isn't trustworthy.

What makes the restrictionist challenge so challenging is that almost any way of responding to it – at least, any way that begins by accepting the empirical results themselves – seems to involve something of a sea change in how we think about the aims or methods of philosophical inquiry. Alvin Goldman (2007), for example, is able to accommodate intuitional diversity by suggesting that philosophical inquiry aims at helping us better understand our individual or shared concepts. On this view, intuitional diversity reflects conceptual diversity, and not all forms of conceptual diversity are necessarily problematic. But, this way of responding to the restrictionist challenge trades a view of philosophy as telling us something about the world for a view of philosophy as telling us something about the ways in which we think about the world, and invites a shift towards certain kinds of philosophical relativism anathema to many traditional philosophers. Nevertheless, resisting this trade comes with its own methodological costs, involving (at the very least) *local* methodological restrictions on our intuition-deploying practices together with a *global* shift in how we think about and approach these practices. Quarantine only works if we understand the dimensions of an outbreak, and even then only education helps restore and maintain public health, so this way of responding to the empirical challenge not only places limits on *what* intuitional evidence can be appealed to but also conditions on *when* intuitional evidence can be appealed to, and here the suggestion is that our intuition-deploying practices must be informed by an understanding of the relevant psychology, cognitive science, and empirically-informed philosophy of mind.

It is perhaps unsurprising, then, that a more popular way of responding to the restrictionist challenge has become to contest the empirical results themselves, or at least to try to minimize their relevance.¹⁰ And, here is where thick conceptions of philosophical intuition might seem to

intuitions cannot be treated as evidence because they are fallible. These challenges have tended to rest on too narrow a conception of philosophy (e.g., Kornblith's challenge), to demand of us the epistemically impossible (e.g., Cummins' challenge), or to be so strong that they deem all putative evidence untrustworthy (e.g., Devitt and Elgin's challenge).

⁹ For more detailed treatments of the restrictionist challenge see Alexander & Weinberg (2007), Weinberg (2007), and Alexander (forthcoming).

¹⁰ Here we focus only on one way of trying to do this; there are others. Some philosophers have responded by questioning the studies themselves (see, e.g., Cullen 2010). But criticizing the experiments and the interpretation of their results is a scientifically substantive project and so this means *more* empirical work for philosophers and not less. And, this means that at least the second part of the restrictionist challenge remains, namely to engage in more

have a dialectical advantage over thin conceptions. If not all putative intuitional evidence is *genuine* intuitional evidence, and if it can be shown that the restrictionist challenge is based on something other than a close study of genuine intuitional evidence, then philosophers are in a position to argue that the restrictionist challenge is actually silent on the standing of philosophy's intuitions-deploying practices.

George Bealer (1998) provides an early version of this way of defending philosophy's intuition-deploying practices from empirical challenge:

Many philosophers believe that the empirical findings of cognitive psychologists...cast doubt on [intuitions'] epistemic worth. But, in fact, although these studies bear on "intuition" in an indiscriminate use of the term, they evidently tell us little about the notion of intuition we have been discussing, which is relevant to the justificatory practices in logic, mathematics, philosophy, and linguistics. As far as I have been able to determine, empirical investigators have not attempted to study intuitions in the relevant sense. (213)

Although not directed at the restrictionist challenge (it was instead directed at earlier empirical challenges stemming from work by folks like Richard Nisbett, Daniel Kahneman and Amos Tversky), this strategy has found new voice among philosophers concerned more directly with the restrictionist challenge. When discussing the methods employed by many experimental philosophers, for example, Kirk Ludwig (2007) writes,

The first point to make is that...responses to surveys about scenarios used in thought experiments are not *ipso facto* intuitions, that is, they are not *ipso facto* judgments which express solely the subject's competence in the deployment of the concepts involved in them in response to the scenario...The task when presented with responses which we know are not (at least all) intuitions is to try to factor out the contribution of competencies from the other factors. This requires an understanding of what the various factors are that may influence responses and enough information about each subject to be able to say with some confidence what factors are at work. It is clear that in the circumstances in which these surveys are conducted we do not have this kind of information. (144-145)

In a similar vein, Antti Kauppinen (2007) writes,

My first criticism of the positive thesis of experimentalism – that surveys are a reliable source of evidence for philosophically relevant claims about folk concepts – can then be formulated as follows: the *actual studies* conducted so far have failed to rule out competence failures, performance failures, and the potential influence of pragmatic

careful empirical study of philosophical intuitions. Other philosophers have mounted what we have elsewhere called the *expertise defense*, arguing that philosophers shouldn't be interested in folk philosophical intuitions (see, e.g., Hales 2006, Ludwig 2007, and Williamson 2007; and, for discussion, see Weinberg et al. 2010). But, who has expertise about what and under what circumstances turns out to be a rather complicated empirical question. It seems that only certain kinds of training help improve task performance and, even then, only for certain kinds of tasks, and there is reason to worry that philosophical training doesn't seem to be the right kind of training nor does philosophical thought-experimenting seem to be the right kind of task.

factors, and as such do not yield the sort of results that could support or raise doubts about philosophical appeals to conceptual intuitions.

Simon Cullen (2010) explains the methodological burden being levied here:

[W]e must distinguish intuitions from survey responses...Survey responses are a kind of behaviour generated by several different kinds of inputs, only one of which might a [philosophical intuition]. Other inputs are the background beliefs a respondent draws on when interpreting a vignette; her beliefs about what the researchers are interested in; her sensitivity to the conversational norms; and so forth.

And, John Bengson (forthcoming) argues that this burden has not been met:

[T]hese attacks neglect a considerable gap between the answers elicited by the relevant empirical studies and the intuitions about which naysayers naysay. It cannot innocently be assumed that subjects' answers expressed how things struck them—what intuitions they had, if any. The point is simple, but not insignificant. For, I will argue, it implies that we are at the present time unwarranted in drawing any negative conclusions about intuitions from the relevant empirical studies.

The common idea seems to be this. Only *very special* kinds of things count as genuine philosophical intuitions and, owing to the limits of their methods, experimental philosophers haven't been in touch with these kinds of things. Or, put more conservatively, experimental philosophers can't be certain that they've been studying the right kinds of things. Since we can't be certain that they've been studying the kind of intuitional evidence that is relevant to philosophy's intuition-deploying practices, the results of these studies shed little insight into the standing of those practices.

4. The Challenge of Sticking with Thick Intuitions

Let's call this way of responding to the restrictionist challenge by appealing to a thick account of genuine philosophical intuition, the *thickness defense*. It is worth taking a minute to unpack the commitments that any proponent of a thickness defense must sign on to, since then we can investigate whether those commitments are ones that are likely to be satisfied. There are generally two kinds of conditions that must hold for any given version of the thickness defense to succeed in providing a response to the restrictionist challenge. The first kind we will call *veritist conditions*, as they concern the propensity of genuine philosophical intuitions to track (or diverge from) the truth. Here, we will consider two different veritist conditions; there may be others. The first we will call the *immunity condition*: genuine philosophical intuitions are not themselves susceptible, to any worrisome degree, to the sorts of problematic effects that form the basis of the restrictionist challenge just rehearsed. The qualifier "to any worrisome degree" is important, since all that needs to be required is a methodologically sufficient immunity to these sorts of errors, and not some sort of quasi-Cartesian absolute immunity. The immunity condition is about the sorts of effects already explored by experimental philosophers, but there may well be other sorts of problematic effects beyond the handful that have been investigated so far. Now, it truly would be a skeptical argument of a fairly traditional sort to attempt to impose a requirement that proponents of the thickness defense rule out even the mere possibility of some other, as-yet-

unexplored and as-yet-unspecified sources of error. And so it is important to see that this is not at all the kind of argument that restrictionists are trafficking in. What we can legitimately require, however, is that the specific sorts of factors that are meant to distinguish genuine philosophical intuitions from “faux” philosophical intuitions – what we will call *thickening factors* – not seem at all likely to generate their own distinct problematic effects. Thickening factors are meant to save traditional armchair philosophical practices; they shouldn’t raise their own new dangers for it. Should a thickening factor seem likely to lead to its own form of troubles, then a version of the thickness defense employing that thickening factor would be ultimately fruitless, merely leading restrictionists to swap brands of ammunition. Any given version of the thickness defense if it is to cure traditional armchair philosophical practices, must thus first do no harm to those practices; we thus call this second veritist condition the *hippocraticity condition*.

Where the immunity and hippocraticity conditions both pertain to the truth-tracking or truth-diverging properties of alleged genuine philosophical intuitions, there is another set of conditions that must be met by any version of the thickness defense. We will call these *methodological conditions* – can the proposed thickening factors, regardless of their veritistic properties, be sufficiently put to work in philosophical practice? There are at least two methodological conditions that must be met by any version of the thickness defense. The first condition concerns whether agents can actually distinguish between genuine and faux philosophical intuitions. After all, here’s a candidate thickening factor that clearly will satisfy any veritist conditions: genuine philosophical intuitions are all and only the ones that are *true*. The problem is that proposing a verity condition as a thickening factor would obviously be a non-starter without some further means at hand for separating true philosophical intuitions from false.¹¹ Thickening factors, thus, need to be *detectable*, and not just in principle or theory, but in practice and – most of all – from the armchair. We call this the *manifestability condition*.

While immunity, hippocraticity, and manifestability are rather substantive conditions of adequacy, it is still possible for a version of the thickness defense to satisfy all three of these conditions while still failing to provide a satisfactory defense of *current* armchair philosophical practices. This is possible in cases where the thickening factor is one that, while veritistically virtuous and potentially manifestable, nonetheless does not currently play a sufficiently universal role in philosophical practice. Suppose it turned out, just to illustrate with a toy example, that intuiting while lying prone in a very brightly lit, bare room provided maximal stability to our intuitions. It would be easy to tell when we are in such a situation, and we could tell from the armchair that we were in it (well, *ex hypothesi* we’d be telling it from the *floor*, but you get the idea). The problem is that since philosophers do not currently pursue this technique, this version of the thickness defense would still require some substantial changes to the norms of philosophical cognition and argumentation. Now, this commitment need not apply to everyone running a thickness defense, since it is open to defenders of armchair philosophical practices to deny that, in defending those practices, they are also looking to defend the *status quo cathedrarum*. Bealer, for example, sometimes protests that he is only defending the *possibility* of armchair philosophy. But we think that many, and perhaps most, proponents of the thickness

¹¹ It is important to note that we aren’t objecting to the idea that truth might be a necessary condition on genuine philosophical intuitions, although we are rather skeptical that it is; our concern is simply that the condition can’t do any work to meet the restrictionist challenge unless we have some way of identifying which of our philosophical intuitions are true.

defense do not take themselves to be defending just some merely possible version of armchair philosophical methodology. They take themselves to be defending what philosophers have been doing in recent decades, and continue to be doing today. These proponents of the thickness defense, at least, are committed to what we would call the *current practice condition*.¹²

Although we will not deploy the current practice condition for any heavy lifting in the discussion that follows, we have two reasons for having wanted to articulate it briefly here. First, we think proponents of traditional armchair philosophical practice should be more cognizant than they often are of the fact that many ways of responding to the restrictionist challenge are rather more revisionary than they seem to commonly take them to be, and we will flag a few instances of this phenomenon below. Second, it helps illustrate that the three main conditions that we are looking to use in our evaluation of the thickness defense – immunity, hippocraticity, and manifestability – are likely not exhaustive. There are probably other substantive empirical commitments that proponents of the thickness defense need to have a right to, if they are to succeed in their response to the restrictionist challenge.¹³ It would be wise for them to start paying attention to these commitments, and articulating and defending them themselves. We hope, even if they don't agree with our critical evaluations below, that they will find our own initial articulation of these commitments helpful. We take it that both restrictionists and the defenders of the armchair share the desire to gain as clear a picture as possible of how well-functioning philosophical practices can work, and thus also share an interest in not merely scoring unilluminating debating victories, but in learning from the moves and countermoves of the debate.

Our project here is to score a number of prominent versions of the thickness defense in terms of these three major commitments: immunity, hippocraticity, and manifestability. We will do so as follows. For each proposal, we will evaluate its status for each commitment using a very rough scale of {highly likely; open, but probable; wide open; open, but improbable; highly unlikely}. Think of the first and last of these values as being likely enough that one should, all things considered, at least *pro tempore* believe the claim or its negation. The middle three values, on the other hand, do not involve probability sufficient to license believing the claim or its negation, although “open, but probable” and “open, but improbable” mark a current tendency in the state of evidence. Since proponents of any version of the thickness defense need all three of these commitments to be true, we will focus our evaluation on the conjunction of all three commitments taken together. Here, we will combine these evaluations not by anything like, say, averaging them, but rather by taking the score of the worst member of the set to provide the score of the set on the whole. After all, the fact that $\Pr(A \& B \& C)$ cannot be any greater than $\text{MIN}(\{\Pr(A), \Pr(B), \Pr(C)\})$ is itself trivial. And on a fairly weak set of further assumptions, this will actually turn out to be a rather generous way to produce an overall evaluation: unless at least two of $\Pr(A)$, $\Pr(B)$, and $\Pr(C)$ are 1 (and none are 0) it will turn out that $\Pr(A \& B \& C)$ will be strictly less than $\text{MIN}(\{\Pr(A), \Pr(B), \Pr(C)\})$. In the end, then, each proposal can itself be evaluated as highly likely, highly unlikely, or as a hypothesis of greater or lesser plausibility.

¹² And we think they are right to do so. Speaking just for ourselves, we are happy to *stipulate* to the metaphysical possibility of a community using the target armchair practices in a methodologically unproblematic way. The extent to which we actually resemble the members of such a community is, of course, the tricky empirical question.

¹³ The same holds true for other attempts to defend armchair philosophical practice from the restrictionist challenge. See, for example, recent discussions of philosophical expertise (Williamson 2007, Weinberg et al. 2010, Williamson 2011).

After scoring the individual proposals, we can consider how things stand in the debate between the restrictionists and the defenders of the armchair. If *any* version of the thickness defense comes out as highly likely, then that would constitute a pretty clear victory for the thickness defense. Conversely, if *all* of them were to come out as highly unlikely, then that would also constitute a pretty telling rejection of this whole line of response to the restrictionist challenge. Perhaps further investigations could revise either such result, but it would take some such further investigations to do so. But what if some versions of the defense come out as being only a more or less plausible sort of hypothesis? And with no versions of the defense doing any better than that? It's important to keep clearly in mind that such an outcome would make for a victory for the restrictionist challenge as well, even if a less satisfyingly resounding one. After all, it is a core part of the restrictionist challenge that the particular boundaries of where our intuition-deploying practices are already unproblematic can only be discerned from outside of the armchair, by means that involve substantial empirical investigation. So, if all versions of the thickness defense stand at this time as, at best, under-substantiated empirical hypotheses, then exactly this kind of further empirical investigation would be what would be needed to try to resolve their status. And the less plausible the hypothesis, the more that this is so, for obvious reasons. But even should there be at least one version of the thickness defense that is modestly plausible – even if still not yet sufficiently confirmed – all that this means is that the defenders of armchair philosophical practice might have an interesting line of investigation that they could (and should!) pursue. That there are paths that *might* lead to results that block the restrictionist challenge does not at all entail that such results really are there to be had. Only a proper investigation would tell.

While there are as many different versions of the thickness defense as there are potential thickening factors, in what follows we will focus on the three most prominent ways of setting up the thickness defense: *phenomenological approaches* that focus on introspectively accessible aspects of intuitional experience; *etiological approaches* that focus on conceptual competence; and *methodological approaches* that focus on philosophical reflection. Each approach is committed to the idea that genuine philosophical intuitions are sufficiently immune from the kinds of problematic effects that have served as the basis of the restrictionist challenge and from the kinds of problematic effects that would only provide experimental restrictionists with a different basis for that challenge, and to the idea that we can distinguish genuine philosophical intuitions from faux philosophical intuitions on the basis of nothing more than armchair reflection. They disagree only about how best to separate the intuitional wheat from chaff, with one focusing on how genuine philosophical strike us, another focusing on where they come from, and yet another focusing on what we do with them once they are here. The question is whether any of these ways of setting up the thickness defense gives us sufficient reason to believe at this time that genuine philosophical intuitions are appropriately immune, hippocratic, and manifest. We'll take each in turn.

A. *Phenomenological Approaches*

Let's start with George Bealer's version of the phenomenological approach, perhaps the most famous version of the thickness defense on offer. According to him, genuine philosophical intuitions have a special phenomenological mark, which he characterizes in terms like "the appearance of necessity". One problem that we immediately run into concerns the conditions of

immunity and hippocraticity; Bealer simply hasn't given us any reason to think that genuine philosophical intuitions, understood in phenomenological terms, are immune from the kinds of problematic effects that form the basis of the restrictionist challenge or from other kinds of problematic effects. That one very particular aspect of a process (its phenomenological feel) is correlated with another (its immunity to order effects, etc.) is a remarkably strong claim, one that would require significant empirical work to substantiate.¹⁴ And, while we invite proponents of this kind of thickness defense to take up this work, we again want to stress that the need for this kind of empirical work is itself a victory of sorts for the restrictionist challenge since part of that challenge is precisely that empirical work is going to be needed in order to establish the standing of philosophy's intuition-deploying practices.

The real issue with Bealer's version of the phenomenological approach, though, is with manifestability. Let us look at George Bealer's version first, perhaps the most famous version of the thickness defense on offer. According to him, genuine philosophical intuitions have a special phenomenological mark, which he characterizes in terms like "the appearance of necessity". The problem is that it turns out that there is a great deal of disagreement, at the very highest levels of the profession and even among philosophers who are friendly to armchair methods, about whether *any* genuine philosophical intuitions display this special phenomenological mark (see, for example, Sosa 1998, Lynch 2006, Goldman 2007, Williamson 2007, and Pust forthcoming). Part of the problem is that it is not particularly clear what it means for a proposition to have the appearance of necessity (see, for example, Pust forthcoming). Does it seem to us that the proposition is necessarily true or does it seem necessary to us that the proposition is true? But, problems would remain even if we were able to make clear what it means for a proposition to have the appearance of necessity. Not all genuine philosophical intuitions, it seems, have a particular way of *seeming*, which means that focusing on the fact that some things strike us one way, while others strike us in other ways, doesn't provide us a way to determine which of these things are genuine philosophical intuitions. So, we have a fairly complete failure of manifestability here. Since there is disagreement from the armchair about whether any genuine philosophical intuitions have the appearance of necessity, it is impossible to use this phenomenological mark to determine from the armchair which philosophical intuitions are genuine. This also explains why it is so hard to evaluate Bealer's version of the phenomenological approach in terms of immunity and hippocraticity: unable to identify the class of cognitions in question, we struggle to be able to evaluate whether cognitions in this class are or are not likely to be susceptible to various sorts of effects! But, even setting this aside, the fact that Bealer's approach runs into problems with manifestability is enough by itself to make the approach highly unlikely to be successful.

What about Bengson's version of the phenomenological approach, according to which genuine philosophical intuitions "strike us in a certain way"? We take Bengson's formulation to pick out something like the idea, fairly common in the literature on intuitions, that intuitions have a phenomenology of a proposition's *seeming* to be true, in a way that feels similar in this regard to perceptual seemings, but not typically similar in terms of the sense-modality aspects of

¹⁴ Depending on what it means for something to "strike us in a certain way", there is some work that suggests that some philosophical intuitions are immune from at least some kinds of problematic effects. So, for example, Jen Wright (2010) has found that confidence is a good indicator of relative immunity from presentation order effects.

perceptual seemings. (For this reason philosophers often refer to them as *intellectual* seemings.) Bengson raises the possibility that the participants in experimental philosophy studies are largely providing answers that are not driven by any such seemings, and clearly intends this possibility to make trouble for restrictionists. It is thus an excellent candidate for consideration in our framework here: Bengson is offering a thickness defense, with the phenomenology of seeming as his candidate thickening factor. And we can thus score his proposal in terms of the three core commitments under consideration here.

First, Bengson's version of the phenomenological approach makes it very likely that genuine philosophical intuitions will be manifest. We take it to be the case that intuitiveness, in the attenuated sense he is considering it, is something that is pretty well accessible to introspection. It seems to us highly plausible that this phenomenology can be distinguished reasonably well from the phenomenology of what Bengson calls "blind or stray answers" – answers that "express hypotheses, guesses, emotional reactions, inferences or other conclusions."¹⁵

Problems arise, however, when we consider the two veritist commitments. Bengson is claiming that, if ordinary subjects restricted their judgments just to those with the appropriate phenomenology of seeming, all these unfortunate effects that form the basis of the restrictionist challenge would disappear. He does offer conjectures about some interesting *possibilities* about ways in which experimental philosophy studies might be going astray. Our core diagnosis of where he goes wrong, though, is in thinking that mere possibilities – even empirically real ones, which we happily grant that Bengson's clearly are – can carry much weight in these sorts of debates. For example, his claims are typically about things that "sometimes" happen; for example, he frames two of his most central claims as, "Sometimes things strike us a certain way; other times they do not, even though we may still be able to give an answer to a question about whether things are that way, when prompted," and "Sometimes things strike us a certain way, even though we may answer that things are not that way, but rather some other way, when prompted." These claims, as "sometimes" claims, are clearly true. But they go no distance towards securing what Bengson actually needs: that these things happen *to a large enough extent and in right distribution* to explain away all of the restrictionist's favorite experimental findings.

Bengson comes close to seeing this when he examines the plausibility of moving from the idea that

[AEI] A subject's prompted answer expresses a subject's intuition

to the idea that

[AEI*] A subject's prompted answer *usually* or *typically* expresses a subject's intuition.

Bengson, who thinks that restrictionists are committed to either AEI or AEI*, gives fairly good reasons for thinking that moving to AEI* isn't particularly helpful, namely, that there's no particular reason, given the experimental designs currently used, to think that "seeming-less"

¹⁵ Actually, there might be reason to worry that this is too generous since it isn't clear how Bengson thinks that we can introspectively peel off "emotional reactions" from seemings more generally, but we don't want to press this worry here.

answers are vastly less common than “seeming-ful” answers. We do not wish to defend AEI* here – we agree with him that it’s not a good commitment to have. The problem for Bengson is both AEI *and* AEI* are much stronger commitments than restrictionists actually need. It seems to us that all they need, in order to respond to Bengson’s version of the thickness defense, is something like this:

[AEIE] The undesirable effects are not themselves completely localized to the to seeming-less, “blind” answers, but also manifest in seeming-ful, intuition-based answers.

AEIE is strictly weaker than AEI*. On the one hand, it is likely entailed by AEI*. If pretty much all the data in the experimental philosophy surveys were coming from Bengsonian intuitions, then it would be highly likely that the observed effects are present in such intuitions. But, on the other hand, AEIE does not entail AEI*. Suppose, for example, that the answers themselves reflect a 50-50 mix of intuitions and strays answers, but the patterning of the unwanted effects is identical in both sets. Then AEI* would be false, even though AEIE would still be true.

Bengson, unfortunately, doesn’t consider AEIE or anything like it, and so doesn’t offer any counter-arguments to it.¹⁶ And AEIE is, at this time, itself highly likely. For the dominant view about intuitive cognition in scientific psychology at this time is that of the “dual process” or “two system” models of reasoning, judgment, and decision-making.¹⁷ One of these systems is commonly described as typically involving fast, apparently effortless cognition, based more on similarity and associations, whereas the other system involves slower, effortful, rule-based cognition. On this model of human cognition, the first system, “System 1” is both the source of (at least many) intuitive seemings, and also susceptible to many effects of the sorts that restrictionists have worried may afflict philosophical intuitions. The fact that these characteristics – intuitiveness and instability – seem very frequently to travel together in the human mind provides solid support for AEIE.

First, it is just not the case that the psychologists investigating System 1 have generally been just considering seeming-less cognitions. For example, in a paper based on his Nobel acceptance address, Daniel Kahneman (2003) observes that, “A core property of many intuitive thoughts is that under appropriate circumstances, they come to mind spontaneously and effortlessly, like percepts” (699), and also frequently analogizes intuitions to sensory impressions. And Jonathan Evans (2003) writes, in an article intended as an introduction to some of these issues, that “Dual-process theorists generally agree that System 1 processes are rapid, parallel and automatic in nature: only *their final product is posted in consciousness*” (454; our emphasis). In our

¹⁶ He does consider a different defense of AEI*, which would involve the claim that “non-intuitive responses such as hypotheses, guesses, emotional reactions, and inferences are (usually are, typically are) random, arbitrary, or nonsystematic.” He contends that that claim is not likely to be true, and we are again in agreement with him. But that is very different from – indeed, pointed in the opposite direction from – the claim that we are considering: that many of the systematic effects manifest in “stray” answers are *also* manifest in seemings. This same problem applies to much of the rest of his discussion there, where he proposes various possible explanations for how non-seeming cognitions could display the restrictionist effects; namely, he does not do anything to make probable that seeming cognitions will not display such effects as well.

¹⁷ Richard Samuels (forthcoming) has recently argued that they are perhaps better viewed as two families of a much larger number of systems, rather than two systems proper. Nothing in this discussion turns on that distinction.

experiences with this literature, these kinds of formulations are fairly common. The workings of System 1 are not conscious, but their products are consciously available.

Second, System 1 has many traits of exactly the sort that Bengson is hypothesizing are *not* involved in genuine intuitions. For starters, System 1 cognition is largely association-based, and for that reason it is also highly susceptible to contextual factors. Kahneman (2003) summarizes a large body of his research concerning the strong presence of framing effects in intuitive cognition, arguing that framing has a big influence on accessibility, and that “intuitive decisions will be shaped by the factors that determine the accessibility of different features of the situation. Highly accessible features will influence decisions, while features of low accessibility will be largely ignored. Unfortunately, there is no reason to believe that the most accessible features are also the most relevant to a good decision” (710). (See also, e.g., Morewedge & Kahneman (2010, p. 437), for a particular discussion of how processes of “associative coherence” lead to both anchoring and framing effects.) We would conjecture that the reliance of System 1 on association, as well as its use on learned prototypes of categories (see Kahneman 2003, 712ff), would also be a potential source of demographic variation: members of different cultural groups may well form subtly distinct associations, as well as form prototypes from somewhat different stockpiles of individual cases. Finally, many canonical descriptions of the processes in System 1 also include that they are sensitive to emotion, perhaps even including an “affect heuristic” (Slovic et al, 2002).¹⁸

In short, according to this widely-held view, the outputs of System 1 processes are typically seeming-ful, and also typically sensitive to exactly those sorts of factors that restrictionists are worried that philosophical intuitions may be sensitive to as well. There just is not much reason at this time to think that a cognitive state or episode’s having the phenomenology of seeming-ness correlates in any useful way with being shielded from the influence of context, demography, or affect. Accordingly, restrictionists do not need at this time to lose any sleep about the status of commitments like AEIE.

It’s important to see why it is things like AEIE, and not Bengson’s AEI or AEI*, that are at stake in this debate. Here’s a totally made-up case that might help to illustrate the basic problem. Suppose someone wanted to run with a version of the thickness defense in which the thickening factor is *being intuited on a Sunday*. Adherents to philosophical sabbaticism (or, perhaps seventh-day rationalism) could make a very strong case that very little of the existing experimental philosophy work has examined what would be, by their lights, genuine intuitional evidence. After all, a very large proportion of existing restrictionist work was done in a classroom setting, during or after an actual meeting of a university class, and it is thus enormously probable that none of that data was collected on a Sunday. In this regard, the sabbaticists are in an even better situation, in terms of the structure of Bengson’s argument, than Bengson himself is. But we hope it is utterly obvious why sabbaticism is of course a complete non-starter as a response to restrictionism: there’s no reason at all to think that our cognition works different in any particular way on Sunday than it does on any other day of the week. In terms of our evaluative framework, sabbaticism would do great in terms of manifestability –

¹⁸ Bengson at one point in his paper *contrasts* heuristic and intuitive cognition (section 5.4). This seems not to be the view held by most researchers in this area, in that many heuristics, especially those not deliberately deployed, are located in System 1.

everyone can tell pretty easily what day of the week it is! – but would be an abject failure in terms of immunity.¹⁹ The moral here is that the success of a version of the thickness defense does not depend solely on whether the allegedly genuine intuitions turn out to be *in some way or other* different from whatever is being studied by the restrictionist experimental philosophers. Instead, specific versions of the thickness defense can be successful only if they identify some factor that distinguishes the allegedly genuine intuitions, and makes it likely that they are psychologically distinct in a way that renders them sufficiently shielded from the sorts of errors that form the basis of the restrictionist challenge.

Now, sabbaticism is offered as an illustration of a problem with Bengson's argument, but we do want to stress that it is not an analogy to Bengson's argument. Sabbaticism is silly; Bengson's proposal is most definitely not. In fact, it seems to us to be a legitimate scientific hypothesis, one that is absolutely worthy of further investigation, and we hope that he will pursue such an investigation. There may well be yet-unexplored divisions within System 1, such that some particular sorts of effects like the ones proposed by restrictionists do not, as a matter of fact, manifest in seeming-ful cognition. The current state of the literature does not, we think, rule this hypothesis out. But Bengson has failed, we fear, to recognize that it is a hypothesis that cannot do the dialectical work he wants it to do, unless and until there is some substantial evidence that it is at least somewhat likely to be true. As we noted above, Bengson tends in his paper to frame matters in terms of possibilities. But once we start actually looking at probabilities, his version of the thickness defense looks rather implausible, although, again, eminently worthy of further investigation. It is, we think, an open, although somewhat improbable hypothesis.

B. *Etiological Approaches*

Etiological approaches would distinguish genuine from ersatz intuition by their origins. There's really only one extant version of this kind of approach at this time, and it is the view like Kirk Ludwig's described earlier, that genuine philosophical intuitions are intuitions made only on the basis of conceptual competence. To see how conceptual competence versions of the thickness defense fair, let's start with the immunity condition. It turns out that conceptual competence approaches fair somewhat better vis-à-vis the immunity condition than do phenomenological approaches. After all, it is plausible to think (if not yet empirically demonstrated) that issues of conceptual *performance* might be responsible for the kinds of problematic effects (e.g., order and framing effects) that form the basis of the restrictionist challenge. And, nothing that we know of from psychology or cognitive science seems to *directly* undermine the plausibility of this suggestion. But, it is important to keep in mind that whether or not conceptual competence approaches will sufficiently immunize genuine philosophical intuitions depends crucially on what theory of concepts we are working with. Many theories attribute structures to the mind that are themselves contextually sensitive and widely variable across both individuals and populations – for example, theories according to which our concepts contain representations of exemplars (see, e.g., Smith 1999) or theories according to which sensitivity to contextual factors is practically a defining characteristic (see, e.g., the discussion of proxytypes in Prinz 2002). Of course, philosophers might have an alternative theory of concepts in mind, perhaps one that would impose context-free sets of necessary and sufficient conditions on category membership.

¹⁹ Oddly enough, it would perhaps also do fine in terms of hippocraticity: there's no reason to think that there's any aspects of our cognition that are *worse* on Sundays than on other days, either.

But, it is important to point out both that these kinds of theories haven't proved particularly popular to folks working on concepts in psychology or cognitive science and that the truth of such theories will depend on how the literature on concepts plays out. As such, if philosophers want to defend their armchair practices by appealing to this kind of theory of concepts, they need to recognize that doing so involves a *substantial* empirical commitment about the nature and structure of concepts, one that doesn't have much empirical support from psychology or cognitive science.

Along similar lines, there is reason to worry that conceptual competence approaches aren't going to be particularly hippocratic, either. To see why, it is important to note that in order for these approaches to even get off of the ground as a defense of our intuition-deploying practices, it must be that conceptual competences have some sort of rich structure; there is no way for a conceptual atomist to even begin to run this line of argument. Proponents of this view will require concepts to have theory-like structures. But if some sort of internalized theory determines conceptual content, then we have good reason to worry about the effects of theory contamination on our conceptual judgments. If the conceptual theories are acquired or replaced by tacit learning processes (cf. Reber 1993) or by the modularization of explicitly formed theories (cf. Karmiloff-Smith 1993), then we would face a significant problem for philosophy's intuition-deploying practices. We advance philosophical theories on the basis of their ability to explain our philosophical intuitions, defend their truth on the basis of their overall agreement with our philosophical intuitions, and justify our philosophical beliefs on the basis of their accordance with our philosophical intuitions. If those same theories play a role in shaping our conceptual judgments – if we have our intuitions because of our theories, and not the other way around – we risk a most vicious kind of epistemic circularity (for discussion, see, e.g., Cummins 1998). Because this risk is rather more hypothetical, though, given that these philosophers haven't even put a real conceptual competence account of the right sort of the table, it is probably appropriate to err on the side of charity in our evaluation, and so we are willing to say that it is an open hypothesis whether conceptual competence approaches are going to be particularly hippocratic.

Things get worse. Even if genuine philosophical intuitions were to come out as both immune and hippocratic on conceptual competence, there would still be reason to worry that the conceptual competence approach – and indeed etiological approaches more generally – won't be able to render genuine philosophical intuitions particularly manifest. It is going to be quite hard, indeed maybe even impossible, to introspectively determine when our own philosophical intuitions are the product of our conceptual competence, and when instead they have been influenced by other factors. For *some* other factors, especially pragmatics, we do have at least some tests that can perhaps be performed in the armchair. But for many others we don't, and indeed many philosophers (e.g., Lynch 2006, and Henderson & Horgan 2001) have particularly emphasized just how *opaque* to us the origins of our intuitions are, in a manner well consistent with general cognitive psychological doubts (e.g., Nisbett & Wilson 1977, Wilson 2002) about the reliability of introspection to get at origins of or processes underlying our actions and thoughts. And we also noted above how research on System 1 cognition generally concurs with this view about the unavailability of the underlying processes to consciousness. In short, we are just not at all likely to be well suited to tell – *especially* with only the limited resources of the armchair – which of our philosophical intuitions are genuine. All of this makes it rather unlikely that genuine philosophical intuitions are going to be manifest on the conceptual competence approach. We

would also caution that these same reasons would likely apply to any other etiological approach to the thickness defense, were someone to try to formulate one.

C. Methodological Approaches

According to methodological approaches, genuine philosophical intuitions are the products of careful philosophical reflection. (In principle there could be other methodological proposals, such as the role of peer review, or the applications of formal tools, and so on. But the method of reflection is the only current contender in the literature.) To see how this approach fares, let's start with manifestability. We take it that the condition of manifestability is in fairly good shape here, in that reflection is an armchair-pursuable activity *par excellence*. Nonetheless, we will note two worries, without pressing them further. First, proponents of any reflection-based approach owe us some sort of story about how much reflection is enough, and how we can tell. Second, although one can engage in reflection from the armchair, it may sometimes be the case that correction can only come by reflecting on information gathered from outside the armchair. For example, no amount of reflecting on our intuitions about space and time was ever going to teach us that space-time is not Euclidean! It might be that reflection can indeed help improve people's intuitive judgments, without it being the case that *armchair* reflection is what is needed. Nonetheless, we are willing to grant that it is fairly likely that genuine philosophical intuitions will be manifest on this approach.

The problem with reflection-based approaches is that they score rather more poorly on both immunity and hippocentricity.²⁰ Although it is at least plausible that the greater reflection would help with at least *some* of the effects (see, e.g., Smith & Levin 1996, on need for cognition and framing), there is no reason at this time to think that it will help with *all* of them. Depending on the exact source of the demographic variations, reflection *might* help; in particular, if it turns out that the demographic differences are due to differences in motivation to reflect in the first place, then making everyone reflect more might bring everyone to a convergent solution (for discussion of this point, see Nagel (forthcoming)). But, similar to what we saw with Bengson's proposal above, all we have are a string of possibilities, currently with at best very little evidence on their behalf. Moreover, there is also some recent evidence suggesting that, in at least some cases, further reflection does not make some effects go away (see, e.g., Weinberg et al. forthcoming). At best, then, the reflection version of the thickness defense makes it rather unlikely that genuine philosophical intuitions will be appropriately immune.

Furthermore, work in psychology and cognitive science suggests that for some influences and under some conditions, reflection might intensify their effects. For example, Petty et al. (2001) found that more reflective subjects, when evaluating evidence in a "chunked" setting (that is, where the evidence comes in distinct packages), although they were *less* likely to show recency effects, were *more* likely to show primacy effects. And LeBoeuf & Shafir (2003) report that:

Interestingly, research on persuasion suggests that high-NC participants may sometimes be more affected by framing than their low-NC counterparts. For example, high-NC, but not low-NC, participants find arguments framed to match their moods more persuasive than arguments not so framed (Wegener et al., 1994). Similarly, negative frames are more

²⁰ Hilary Kornblith (2002, 2010) nicely describes further problems associated with reflection.

persuasive than positive frames in encouraging detection-related health behaviors among those high, but not those low, in NC (Rothman et al., 1999). Note also that verbalization during decision making can hurt performance (Schooler & Melcher, 1995; Schooler et al., 1993), and that justification provision can lead to the deterioration of decisions relative to experts' standards (Wilson & Schooler, 1991). Taken together, these studies paint a fairly pessimistic view of the power of extra thought to improve the quality of many types of decisions" (89).

Also relevant to the evaluation of immunity is their summary claim that, while "deeper thinking is likely to prove critical in avoiding some errors, it is unlikely to play a significant role in avoiding others" (90). This is entirely consistent with the restrictionists' general argument: determining just where intuitions are susceptible to what sorts of effects, and where that susceptibility can be reduced or even eliminated – and where it may instead be exacerbated – is a complicated question for empirical science. The armchair simply lacks the resources to sort the signal from the noise in the tangled instrument that is intuitive human judgment. Given a fair amount of evidence indicating that reflection can occasionally induce epistemic harms, it is unlikely that genuine philosophical intuitions are hippocratic on the reflection version of the methodological version of the thickness defense.

5. Conclusion:

Let's take stock. Philosophers need to recognize that that this debate is one that has to play out not in terms of mere possibilities, even legitimately real empirical possibilities, but in terms of probabilities, and what the current state of scientific evidence does – and largely doesn't – make likely. Every move, both for and against the armchair, has to be measured against what commitments it incurs, and how well those commitments actually seem to fare. And it turns out that none of the most popular ways of setting up the thickness defense end up satisfying all three empirical conditions of adequacy set out for a successful response to the restrictionist challenge. Among the phenomenological proposals, Bealer's approach fails most clearly because of its poor score on manifestability, whereas Bengson's succeeds better in that regard, while scoring as at best open-but-implausible on the veritist commitments. The conceptualist version of the etiological approach scores middling at best on the veritist conditions, while completely failing the manifestability condition. And, the reflectionist version of the methodological approaches also scores poorly on both veritist commitments.

Now, it is surely possible that some other specific version of the thickness defense could fare better. Nonetheless, there is a discernible pattern in the distinct shortcomings and limited-at-best successes of the proposals that are on the table. We would close with some speculation as to a good reason why there might in fact be a more *general* problem that will beset *any* thickist attempt to meet the restrictionist challenge from the armchair. It turns out that there is a tension between the condition of immunity and the conditions of hippocraticity and manifestability that might explain why the three conditions can't be simultaneously realized. In order to satisfy immunity, there needs to be a stark difference between genuine philosophical intuitions on the one hand, and ersatz philosophical intuitions on the other – a difference that explains both why faux philosophical intuitions are susceptible to the kinds of problematic effects that form the basis of the restrictionist challenge, and why genuine philosophical intuitions are not. But the

more that genuine philosophical intuitions are a different cognitive creature altogether, the more likely it is that they will have their own peculiar epistemic pathologies, and thus violate hipocraticity. Additionally, in order to satisfy immunity, the difference between genuine philosophical intuitions and faux philosophical intuitions should not be one that folks can easily identify. After all, if folks could easily identify the difference between genuine and ersatz philosophical intuitions, then there is less reason to suppose that faux philosophical intuitions are the ones showing up in the experimental studies that form the basis of the restrictionist challenge. But, the more opaque we make the distinction between genuine philosophical intuitions and faux philosophical intuitions, the harder it will be to satisfy manifestability, as we saw with the appeal to conceptual competence above.

Where does this leave us? It seems that philosophers face a dilemma on a sliding scale: the thinner their conception of philosophical intuition, the more clearly they are challenged by recent work in experimental philosophy; the thicker their conception of philosophical intuition, the more likely they are to incur different kinds of methodological problems and the harder it will be to tell when anyone is doing philosophy correctly. It is very unlikely that a course through these contesting forces could be navigable from the armchair. If there's a useful distinction to be deployed between genuine and ersatz intuitions, it's going to turn out to be one that we will need not just philosophical but also substantive scientific work to chart out.

References

- Alexander, J. 2010: Is Experimental Philosophy Philosophically Significant? *Philosophical Psychology* 23, 377-389.
- Alexander, J. Forthcoming: *Experimental Philosophy: An Introduction*. Cambridge: Polity Press.
- Alexander, J., and Weinberg, J. 2007: Analytic Epistemology and Experimental Philosophy. *Philosophy Compass* 2, 56-80.
- Bealer, G. 1998: Intuition and the Autonomy of Philosophy. In M. DePaul and W. Ramsey (eds), *Rethinking Intuition*. Lanham, MD: Rowman and Littlefield, 201-240.
- Bengson, J. Forthcoming: Experimental Attacks on Intuitions and Answers. *Philosophy and Phenomenological Research*.
- BonJour, L. 1998: *In Defense of Pure Reason: A Rationalist Account of A Priori Justification*. Cambridge: Cambridge University Press.
- Buckwalter, W. and Stich, S. Manuscript: Gender and Philosophical Intuitions.
- Cappelen, H. Forthcoming: *Philosophy Without Intuitions*. Oxford: Oxford University Press.
- Cullen, S. 2010. Survey-Driven Romanticism. *Review of Philosophy and Psychology* 1, 275-296.

- Cummins, R. 1998: Reflection on Reflective Equilibrium. In M. DePaul and W. Ramsey (eds), *Rethinking Intuition*. Lanham, MD: Rowman and Littlefield Press, 113-128.
- Devitt, M. 1994: The Methodology of Naturalistic Semantics. *Journal of Philosophy* 91, 545-572.
- Elgin, C. 1996: *Considered Judgment*. Princeton: Princeton University Press.
- Evans, J. 2003: In Two Minds: Dual-Process Accounts of Reasoning. *Trends in Cognitive Sciences* 7, 454-459.
- Gettier, E. 1963: Is Justified True Belief Knowledge? *Analysis* 23, 121-123.
- Goldman, A. 2007: Philosophical Intuitions: Their Target, Their Source, and Their Epistemic Status. *Grazer Philosophische Studien* 74, 1-26.
- Gutting, G. 1998: Rethinking Intuition: A Historical and Metaphilosophical Account. In M. DePaul & W. Ramsey (eds), *Rethinking Intuition*. Lanham, MD: Rowman and Littlefield Press, 3-16.
- Hales, S. 2006: *Relativism and the Foundations of Philosophy*. Cambridge, MA: MIT Press.
- Henderson, D., and Horgan, T. 2000: What is the A Priori and What Is It Good For? *Southern Journal of Philosophy* 38, 51-86.
- Henderson, D., and Horgan, T. 2001: The A Priori Isn't All It Is Cracked Up To Be, But It Is Something. *Philosophical Topics* 29, 219-250.
- Ichikawa, J. Manuscript: Intuitions and Begging the Question.
- Kahneman, D. 2003: A Perspective on Judgment and Choice. *American Psychologist* 58, 697-720.
- Karmiloff-Smith, A. 1993: Beyond Modularity: A Developmental Perspective on Cognitive Science. *European Journal of Disorders of Communication* 28, 95-105.
- Kauppinen, A. 2007: The Rise and Fall of Experimental Philosophy. *Philosophical Explorations* 10, 95-118.
- Kornblith, H. 1998: The Role of Intuitions in Philosophical Inquiry: An Account with No Unnatural Ingredients. In M. DePaul & W. Ramsey (eds), *Rethinking Intuition*. Lanham, MD: Rowman and Littlefield Press, 129-142.
- Kornblith, H. 2002: *Knowledge and It's Place in Nature*. Oxford: Oxford University Press.

- Kornblith, H. 2010: What Reflective Endorsement Cannot Do. *Philosophy and Phenomenological Research* 80, 1-19.
- LeBoeuf, R., and Shafir, E. 2003: Deep Thoughts and Shallow Frames: On the Susceptibility to Framing Effects. *Journal of Behavioral Decision Making* 16, 77-92.
- Levin, J. 2005: The Evidential Status of Philosophical Intuition. *Philosophical Studies* 121, 193-224.
- Lewis, D. 1983: *Philosophical Papers: Volume I*. New York: Oxford University Press.
- Ludwig, K. 2007: The Epistemology of Thought Experiments: First Person versus Third Person Approaches. *Midwest Studies in Philosophy* 31, 128-159.
- Lynch, M. 2006: Trusting Intuitions. In P. Greenough and M. Lynch (eds) *Truth and Realism*. Oxford: Oxford University Press, 227-238.
- Morewedge, C., and Kahneman, D. 2010: Associative Processes in Intuitive Judgment. *Trends in Cognitive Sciences* 14, 435-440.
- Nagel, J. Forthcoming: Intuitions and Experiments. *Philosophy and Phenomenological Research*.
- Prinz, J. 2002: *Furnishing the Mind: Concepts and their perceptual basis*. Cambridge, MA: MIT Press.
- Pust, J. Forthcoming: Intuitions. *Stanford Encyclopedia of Philosophy*.
- Reber, A. 1993: Implicit Learning and Tacit Knowledge: An Essay on the Cognitive Unconscious. Oxford: Oxford University Press.
- Rothman, A., Martino, S., Bedell, B., Detweiler, J., and Salovey, P. 1999: The systematic influence of gain- and loss-framed messages on interest in and use of different types of health behavior. *Personality and Social Psychology Bulletin* 25, 1355-1369.
- Samuels, R. Forthcoming: The Magic Number Two, Plus or Minus: Dual Process Theory as a Theory of Cognitive Kinds. In K. Frankish and J. Evans (eds.) *In Two Minds: Dual Process and Beyond*. Oxford: Oxford University Press.
- Schooler, J., and Melcher, J. 1995: The ineffability of insight. In S. Smith and T. Ward (eds) *The Creative Cognition Approach*. Cambridge, MA: The MIT Press, pp. 97-133.
- Schooler, J., Ohlsson, S., and Brooks, K. 1993: Thoughts beyond words: when language overshadows insight. *Journal of Experimental Psychology: General* 122, 166-183.

- Slovic, P., Finucane, M., Peters, E., and MacGregor, D. 2002: The Affect Heuristic. In T. Gilovich, D. Giffin, and D. Kahneman (eds.) *Heuristics and Biases: The Psychology of Intuitive Judgment*. Cambridge: Cambridge University Press.
- Smith, E. 1999: The Exemplar View. In E. Margolis and S. Laurence (eds) *Concepts: Core Readings*. Cambridge, MA: MIT Press, pp.
- Smith, S., and Levin, I. 1996: Need For Cognition and Choice Framing Effects. *Journal of Behavioral Decision Making* 9, 283-290.
- Sosa, E. 1998: Minimal Intuition. In M. DePaul and W. Ramsey (eds) *Rethinking Intuition*. Lanham, MD: Rowman and Littlefield Press, 257-270.
- Nisbett, R., and Wilson, T. 1977: Telling More than we can Know: Verbal reports on mental processes. *Psychological Review* 84, 231-259.
- van Inwagen, P. 1997: Materialism and the Psychological-Continuity Account of Personal Identity. *Nous* 31, 305-19.
- Wegener, D., Petty, R., and Klein, D. 1994: Effects of mood on high elaboration attitude change: the mediating role of likelihood judgments. *European Journal of Social Psychology* 24, 25-43.
- Weinberg, J. 2007: How To Challenge Intuitions Empirically Without Risking Skepticism. *Midwest Studies in Philosophy* 31, 318-343.
- Weinberg, J., Gonnerman, C., Buckner, C., and Alexander, J. 2010: Are Philosophers Expert Intuiters? *Philosophical Psychology* 23, 331-355.
- Weinberg, J., Alexander, J., Gonnerman, C., and Reuter, S. Forthcoming: Restriction & Reflection: Challenge Deflected, or Simply Redirected. *The Monist*.
- Williamson, T. 2007: *The Philosophy of Philosophy*. Oxford: Blackwell Publishing.
- Williamson, T. 2011: Philosophical Expertise and the Burden of Proof. *Metaphilosophy*, 42, 215-229.
- Wilson, T. 2002: *Strangers to Ourselves: Discovering the Adaptive Unconscious*. Cambridge, MA: Harvard University Press.
- Wilson, T., and Schooler, J. 1991: Thinking Too Much: Introspection Can Reduce the Quality of Preferences and Decisions. *Journal of Personality and Social Psychology* 60, 181-192.