

“Experimental Philosophy and Apriority”

DRAFT IN PROGRESS; DO NOT CITE UNPERMITTEDLY

15 April, 2010

ichikawa@gmail.com

Jonathan Ichikawa

Abstract

One of the more visible recent developments in philosophical methodology is the *experimental philosophy* movement. On its surface, the experimentalist challenge looks like a dramatic threat to the apriority of philosophy; ‘experimentalist’ is nearly antonymic with ‘aprioristic’. This appearance, I suggest, is misleading; the experimentalist critique is entirely unrelated to questions about the apriority of philosophical investigation. There are many reasons to resist the skeptical conclusions of negative experimental philosophers; but even if they are granted—even if the experimentalists are right to claim that we must do much more careful laboratory work in order legitimately to be confident in our philosophical judgments—the apriority of philosophy is unimpugned. The kinds of scientific investigation that experimental philosophers argue to be necessary involve merely enabling sensory experiences. Although they are not enabling in the sense of permitting concept acquisition, they are enabling in another epistemically significant way that is also consistent with the apriority of philosophy.

According to a certain kind of experimentalist critic of traditional philosophical methodology, it is a mistake to attempt to resolve philosophical questions via armchair methods. On the view in question, philosophy ought to proceed by something much more like the methodology of empirical science; this is meant to contrast with the armchair engagement prototypical of contemporary and historical philosophical argument. Our critic may have performed some relevant historical investigation himself—he comes armed with survey data, or neuroimaging studies, or some other sort of surprising evidence, that he takes to undermine the armchair methodology. There is important empirical investigation, says the critique, that must be performed before philosophical knowledge is possible. The critic in question endorses:

X-Phi: Without engaging in much more scientific investigation (perhaps in the form of surveys, neuroimaging, cognitive psychological theorizing, etc.), we cannot know many alleged pieces of philosophical knowledge.

The critique certainly represents an important challenge to traditional philosophical methodology. If X-Phi is true, then many philosophers ought to be paying much more attention to various bits of empirical research than they are. Many more of us, if the critique is sound, should even be performing experiments ourselves and analyzing the results using the statistical tools prevalent in the sciences. The critique, if sound, undermines a certain armchair conception of philosophical methodology.

It is natural, therefore, to think of the critique as undermining the apriority of philosophy: if one cannot proceed from the armchair, and must engage in empirical, scientific experiments in order to achieve philosophical knowledge, then philosophical knowledge is a posteriori. That is to say, X-Phi stands in obvious *prima facie* tension with the traditional claim:

Apriority: Many paradigmatic pieces of our philosophical knowledge (including many of those targeted by X-Phi) are a priori knowledge.

Although X-Phi and Apriority appear straightforwardly inconsistent, this admittedly natural thought is mistaken. The proper interpretation of the experimentalist critique is not one that bears on the apriority of philosophy. The project of this paper is to suggest that even if the critique is sound—even if X-Phi is true—it is plausible that much philosophical investigation is a priori. Since the critique would if sound undermine an certain kind of armchair conception of philosophy, it follows that apriority and armchair methodology are less closely connected than one might be inclined initially to think.

I won't argue for the truth of Apriority here. Although I do think it's true, establishing it is another, much larger, project.¹ My goal is to show that X-Phi, even if true, provides no reason to reject Apriority.

In §1, I will present identify and precisify the experimentalist critique with which I mean to be engaging; §2 will explore the relation between experience and apriority. In §3 I show that, thus clarified, the apriority of much philosophy is consistent with the experimentalist charge that philosophers need to be engaging with experiments.

§1. Empirical Worries about Philosophy

One of the more visible recent developments in philosophical methodology is the experimental philosophy movement. In some of its more negative paradigmatic instances, experimental philosophy collects some empirical data—perhaps some survey data from folk judgments about philosophical cases, or perhaps some neurological data about what happens in philosophers' brains when they make certain kinds of judgments—and uses it to cast doubt on traditional armchair philosophical methodology. If these experimental philosophers are right, then philosophical methodology, as traditionally practiced, is importantly misguided, and stands in need of substantial revision, with a much greater emphasis on empirical investigation. In this section, I will attempt to articulate what seems to me to be the best interpretation of the experimentalist critique. I will not here attempt to answer the critique head-on—although I do think it can and should so be answered²—my purpose is rather to explore what would follow if its central skeptical conclusion is correct.³

Preliminaries: Two Caveats

Before proceeding to offer my interpretation of the critique arising from experimental philosophy, two caveats are in order. First, the experimental philosophy movement is a wide and diverse one, and not all of its contributions are skeptical. When I say that I am

¹ In Ichikawa and Jarvis (2009), we argue that many thought-experiment judgments are a priori; Ichikawa & Jarvis (forthcoming) argues that much modal knowledge is known a priori.

² See my “The Audacity of Hopeful Philosophical Appeals to Intuition,” (manuscript).

³ As should be clear, I am not using ‘skeptical’ in any kind of pejorative sense; to be skeptical (about X) is merely to deny knowledge (concerning X or deriving from X). To argue that armchair philosophical methodology does not issue into knowledge is to provide a skeptical argument about armchair philosophical methodology. Obviously, this should be distinguished from radical skepticism, which is roughly the view that there is no knowledge at all.

characterizing ‘the experimentalist critique’, I mean to be characterizing the project of a particular significant portion of the experimental philosophy movement; I intend no suggestion that all experimental philosophy is engaged with discrediting traditional armchair methods. Knobe and Nichols (2008) aptly emphasize that much experimental philosophy is positive; it engages with empirical philosophical questions in a way continuous with more traditional methodology, emphasizing only that there are significant scientific resources worthy of exploitation in the investigation of the world. Much of Knobe’s own contribution to experimental philosophy fits this positive mode.⁴

But although not all experimental philosophy is negative, much of it—including many of the most prominent instances of it—is. So it is that, for example, Weinberg et. al. (2001) suggest that that “a sizeable group of epistemological projects—a group which includes much of what has been done in epistemology in the analytic tradition” is “seriously undermined” by survey data they have found; similarly, Machery et. al. (2004) tell us that they have uncovered evidence that shows traditional philosophical assumptions to be “spectacularly misguided,” and that as a consequence, “philosophers must radically revise their methodology.” Alexander et. al. (forthcoming) recognize the divide between experimentalist projects I have been describing, and enjoin us to ‘accentuate the negative’—the skeptical pressure generated by the negative experimentalist program problematizes armchair philosophy and positive experimentalist philosophy alike. When I speak in this paper of ‘the experimentalist critique,’ it is this negative project of experimental philosophy that is my focus.

Here are a few paradigmatic examples of negative experimental philosophy, as I am understanding it. Weinberg et. al. (2001) ran a series of surveys that apparently indicate that some standard epistemic intuitions, such as the internalist intuition about Lehrer’s TrueTemp case and the skeptical intuition about Gettier’s cases, vary according to cultural and socioeconomic background. For example, East Asian subjects seem to be more likely than Western subjects to consider Gettier cases to be instances of knowledge. In a similar study, Machery et. al. (2004) found cross-cultural variation in intuitions about reference in Kripke’s Gödel-Schmidt case. Swain et. al. (2008) found that epistemic intuitions can vary according to the order in which they’re presented. Such survey data comprise perhaps the most typical instances of the experimentalist critique, but negative arguments proceeding from other sorts of data should also be considered. For example, Greene et al (2001) present neuroscientific data demonstrating that particular brain mechanisms underlie certain kinds of moral reasoning; this is thought to undermine certain moral intuitions. The project of this section is to identify the form of the skeptical pressure this data is meant to undermine.

The second early clarification to be made is that the project of characterizing the experimentalist critique is an interpretative one. Instances of negative experimentalist philosophy that suggest the experimentalist critique are only rarely explicit about just what form the argument takes. So in characterizing ‘the’ experimentalist critique, I am attempting to distil the most plausible worry in the vicinity. Among the factors favoring a particular interpretation, then, will be considerations deriving from charity. Given the textual ambiguity of much of the skeptical argumentation in the literature, all things equal, it is uncharitable to distil an argument form that is obviously unsound. It may be that some experimentalist critics intend arguments other than the one that will be the focus of this paper—I don’t mean to be claiming that my version of the experimentalist critique is one that all experimentalist critics

⁴ See, for one influential example, Knobe (2006).

would endorse. But I do mean to claim that it is the best version of the argument; however plausible alternatives may be as experimental-philosophy exegesis, they are not plausibly sound critiques of armchair methodology. I'll begin by canvassing a couple of interpretive options that I will not endorse.

Facts about Intuition as Data

Experimentalist philosophers often emphasize that their discoveries are, from the point of view of armchair methodology, surprising. For example, Machery et. al. write that “most philosophers exploring the nature of reference assume that the Kripkean intuitions [about the Gödel-Schmidt case] are universal.”⁵ They proceed to argue that this assumption is incorrect, and take themselves, in so doing, to have undermined at least some of armchair methodology. On one interpretation of the argument, they understand traditional methodology to be relying on facts about the universality of intuitions as important data in armchair methodology. This interpretation has the advantage of implying straightforwardly that the data in question are in tension with armchair methodology: armchair methodology proceeds from premises about the universality of intuitions.

The central problem with this interpretation of the argument is that it relies on an implausible characterization of armchair philosophy. This point has been adequately made elsewhere, so I won't dwell upon it now.⁶ Max Deutsch puts the point pithily, in discussion of another case about reference, discussed in Mallon et al (2009):

Mallon et al. appear to believe that Evans, if only he had reflected on his own methodology, would have had to retract the claim that ‘Madagascar’ refers to the island, and would have had to patiently await the results of an opinion poll concerning competent speakers’ intuitions about the referent of ‘Madagascar’. But that is preposterous. A philosopher of language such as Evans, just as easily as anyone else, could have simply checked his world atlas and seen that ‘Madagascar’ refers to the island. Facts such as the fact that ‘Madagascar’ refers to the island are *data* for theories of reference. Facts such as the fact that competent speakers intuit that the island is the referent of ‘Madagascar’ are data for a psychological theory, one that does not have any clear bearing on a theory of reference.

Given my broader goal in this paper, it is worth emphasizing a second reason it is unhelpful to interpret the experimentalist critique as assuming that armchair methodology relies on facts about the distribution of intuitions as evidence. I am investigating what bearing the experimentalist critique has on the putative apriority of philosophical inquiry; if, as suggested, philosophical inquiry relies as evidence, or premises, or data, on facts about the distribution of intuitions, I take it the apriority of the relevant inquiry is off the table from the start. Facts about who has what intuitions are not, by anybody's lights, a priori. If armchair philosophy did depend on such empirical claims, we wouldn't need to turn to experimental philosophy to refute the apriority of such methodology; we'd need only point out that any access we might have to such claims would have to come through introspection (in the case of our own intuitions) and testimony (in the case of others'). So Apriority would be a non-starter. Insofar, then, as I am investigating what bearing the critique has on Apriority, this version of the critique would pretty clearly have none.

⁵ Machery et al (2004), p. B8.

⁶ See e.g. Deutsch (2009), Ichikawa (manuscript 1), or chapter 7 of Williamson (2007).

There are therefore two reasons for me not to understand the critique in the way suggested in this subsection: it interprets it as relying on an implausible assumption about philosophical methodology, and it interprets it in a way that rules out Apriority, even independent from considerations arising from X-Phi.

Disagreement

Another way we might understand the experimentalist critique is as a skeptical challenge arising from disagreement. When we discover that East Asian subjects think that Gettier cases are instances of knowledge, this challenges our own belief to the contrary; and absent some reason to prefer our judgment to theirs, we ought not to continue believing that our judgments were correct. Disagreement ought to issue into agnosticism. Sosa (2007) interprets the worry in this kind of way:

How might survey results create a problem for us? Suppose a subgroup clashes with another on some supposed truth, and suppose they all ostensibly affirm as they do based on the sheer understanding of the content affirmed. We then have a prima facie problem. Suppose half of them affirm <p> while half deny it, with everyone basing their respective attitudes on the sheer understanding of the representational content <p>. Obviously, half of them are getting it right, and half wrong. Of those who get it right, now, how plausible can it be that their beliefs constitute or derive from rational intuition, from an attraction to assent that manifests a real competence?

Not that it is logically incoherent to maintain exactly that. But how plausible can it be, absent some theory of error that will explain why so many are going wrong when we are getting it right? Unless we can cite something different in the conditions or in the constitution of the misled, doubt will surely cloud the claim to competence by those who ex hypothesi are getting it right.

...

So there will definitely be a prima facie problem for the appeal to intuitions in philosophy if surveys show that there is extensive enough disagreement on the subject matter supposedly open to intuitive access. (102)

While this interpretation represents a step in the right direction from the one given above—indeed, the treatment I will ultimately give in the following subsection, and develop in §4, will have more than a little in common with this line—there are two problems with understanding the experimentalist critique in this way in generality. First, as a number of authors have emphasized,⁷ since the relevant studies concern unreflective folk intuition, it is not clear that they ought to be afforded weight on a par with those of considered judgments of professional philosophers. We seem to have little reason to regard the subjects of the surveys in question as ‘epistemic peers,’ in the sense of the burgeoning literature on disagreement in epistemology.⁸ Second, and more to the point, understanding the experimentalist critique as one centrally involving disagreement is insufficiently general. Although some of the studies in question emphasize groups of people who dissent from standard intuitions, not all of them do; what are we to make of arguments arising from order effects or neuroimaging? These do not seem to involve disagreement in any straightforward sense. To focus too narrowly on disagreement, then, seems to be an exegetical error.

⁷ E.g. Hales (2006) p. 171, Ludwig (2007), Kauppinen (2007).

⁸ E.g. Elga (2007)

Undermining Defeaters

However, generalizing from the disagreement approach can, I think, avoid both problems given above, and provide a faithful and interesting characterization of the critique. What has disagreement in common with the other sorts of experimental phenomena used to generate skeptical pressure? In generality, such data can provide us with reason to second-guess our initial judgments. That is to say, they can provide *undermining defeaters*. In general, if I encounter someone who disagrees with me as to whether *p*, this provides me with reason to reconsider the grounds for *p*. Similarly, if my judgment that *p* derives from a source *S*, and I discover some evidence that *S* is unreliable with respect to judgments like this one, this provides me with some reason to reconsider my grounds for *p*. Notice that this effect obtains even if my initial judgment was a good one. Sometimes, upon reconsideration, I may correctly decide that my initial judgment was both correct and well-supported after all; other times, agnosticism will be the correct response. But given my cognitive limitations, the discovery of such worrying patterns provides me with reason to be circumspect.

The experimentalist concerns are easily interpretable along this model. If you discover that most members of a culture different from your own are inclined intuitively to judge that Gettier's cases are instances of knowledge, this can provide you with some reason to consider your own reasons for thinking the contrary. If you find that your belief derives from, e.g., general theoretical considerations about the epistemic value of non-accidentality, then you may well endorse it, having passed this new challenge; if, by contrast, you can identify no reason to endorse the belief, finding only that it seems like the thing to say, then you may more properly suspend judgment upon recognition that other people, with grounds apparently as strong as yours, disagree. Similarly, if you discover that epistemic intuitions are susceptible to order effects, then you know a common way in which people can make mistakes about epistemic judgments; it is natural and appropriate, upon such a discovery, to investigate whether you yourself might have been making a mistake in this way. You can do this by considering whether your judgments are stable over a variety of orders of presentation: does your judgment about TrueTemp persist regardless of which other cases you consider immediately before it?

In its broad outlines, this form of concern is not novel to experimental philosophy. Philosophers have a long-standing tradition of questioning the grounds for one another's beliefs; upon discovery that a particular judgment derives from a certain sort of process, where that process is one that is not generally trusted, we find reason to discount that judgment.⁹ The distinctively experimentalist critique adds two elements to this kind of pressure. First, it sharpens our abilities to identify the sources of particular kinds of judgments. Although it is plausible that we can know a fair amount about why we think what we do from the armchair, it should hardly be surprising that the tools of psychology can provide us with important new insights in this domain. As indeed it has done; we now know, for example, on empirical grounds, that moral judgments influence intuitions about the intentionality of actions.¹⁰ Second, experimental investigation can tell us more about the reliability of particular kinds of processes. This is, I take it, what occurs in the order effects case; we learn, via Swain et. al.'s survey results, that unreflective intuitive judgments are

⁹ See Ichikawa (2010)

¹⁰ Knobe (2006). Note that this is not to claim that one *couldn't* have noticed such a phenomenon from the armchair, had one thought to consider the question.

susceptible to a particular kind of error; upon learning this, we face, on pain of epistemic irresponsibility, an obligation to check to see whether we are making this kind of error.

I will say more about this kind of undermining evidence in §4.

It is easy to see, then, how the experimentalist critique demands philosophical investigation to move beyond the armchair. There are empirical challenges to our armchair processes. Even if these challenges can be met, they could only be met empirically, either by demonstrating that our armchair processes are not the ones impugned by the relevant data, or by defending those processes by showing that the data, or its interpretation, are incorrect. According to the critique, to provide such a response is a necessary condition for philosophical knowledge. As Alexander and Weinberg (2007) write, “empirical research into the nature of intuitions generated in response to thought-experiments, rather than supporting the use of intuitions as evidence, challenges the suitability of intuitions to function in any evidentiary role.” (66) Absent a defense to this challenge, philosophers proceed illegitimately.

There is much that one may say to dispute this challenge to armchair methodology, but this is not my present concern. I mean, for the purpose of argument, to be quite concessive to the experimentalist critique, understood as undermining philosophical judgments in the way suggested. The central contention of this paper is that, *even conceding a rather radical form* of the experimentalist critique, we have no reason to deny that philosophical investigation is often a priori. That is to say, even if X-Phi is true—even if it is a necessary condition for some bit of philosophical knowledge that one do experiments to ensure that one isn’t going wrong, it may still be that that very bit of philosophical knowledge is a priori.

This claim may strike the reader as straightforwardly contradictory. To see that it is not, we turn now to the relation between experience and apriority.

§2. Experience and the A Priori

A priori knowledge is, to a first approximation, knowledge that does not depend on experience. So approximated, the tension between Apriority and the conclusion of the experimentalist pressure is obvious and leaves little grounds for optimism for resolution. Apriority requires that knowledge be *independent* from experience, but the experimentalist pressure alleges that philosophical knowledge *requires* the experiences of engaging in scientific investigation. However, the relationship between apriority and experience is not so straightforward as this line suggests; it is widely recognized that the approximation above needs qualification.

Experience and Concept Possession

A priori knowledge can ‘depend’ *in some sense* on experience; it *is* consistent with p’s a priori knowability that a subject must have had certain experiences in order to know p. That there are at least some such roles for experience is uncontroversial; consider the standard exemption for experience that plays a role in enabling the possession of concepts; such are standardly thought to be exempt from the requirement that a priori knowledge of p be independent of experience. For example, many philosophers believe in phenomenal concepts—concepts that can only be possessed by beings who have undergone certain sorts of experiences. If there are such concepts, then tautologies involving them can only be thought—and so can only be known—by beings who have had those experiences. Suppose that RED is a phenomenal concept, only possessible by those who have had the experience of seeing red things. Then in order to think the proposition that <all red things are red>, one

must have had experience of red things; therefore, knowledge of that proposition depends essentially on visual experience. (Parallel arguments could apply the point to other putative a priori knowledge beyond logical truths: <all red things are colored>, <red things are extended>, <no surface is both red and green>, etc.) But this kind of dependence is consistent with apriority. We say that such experience is *merely enabling*; this contrasts with a *warranting* role for experience.¹¹

Even if there are no phenomenal concepts¹²—no concepts such that particular experiences are necessary for their possession—it is plausible that as a matter of fact, many of our concepts have come about via various experiences that we have had.¹³ Even if it is possible to possess RED without seeing red things, as a matter of fact, most of us came to this concept via that kind of experience; according to orthodoxy, this role for experience is still consistent with the a priori knowability of some propositions involving RED.

This orthodoxy gives rise to two important questions. First, to what uses can the experience in question—that necessary for concept possession—be put in a subject’s coming to believe and know *p*, consistent with that eventual knowledge’s apriority? It is implausible that just any use is legitimate; otherwise, for example, it might end up a priori for me that I have had the experience of seeing something red. Intuitively, this proposition is a posteriori, relying for its warrant on my past experience of seeing red things; but, on the line in question, this experience is necessary for—or at least casually involved in—my possession of the concept RED. So on what basis can this knowledge be considered a priori? Second and relatedly, what if anything is special about the experience that is necessary for concept possession? What justifies the exception to the general requirement that a priori knowledge is independent of experience? Without a theoretical motivation, one suspects the rule of adhocery; is there such motivation, or is it only included to give the intuitive verdicts about cases?

In the remainder of this section, I’ll lay out what strikes me as a plausible theoretical characterization of what is essential to a priori knowledge and justification. This characterization will answer both questions above, explaining why, and under what circumstances, important roles for experience necessary for concept possession are consistent with apriority. A central upshot, for my purposes, will be that the relevant roles for experience extend beyond experience necessary for concept possession.¹⁴

¹¹ As Timothy Williamson (2007, p. 165) writes: “Standard discussions of the a priori distinguish between two roles that experience plays in cognition, one evidential, one enabling. Experience is held to play an evidential role in my visual knowledge that this shirt is green, but a merely enabling role in my knowledge that all green things are coloured: I needed it only to acquire the concepts *green* and *coloured*, without which I could not even raise the question whether all green things are coloured.” Williamson goes on to critique this orthodoxy.

¹² Ball (2009)

¹³ Of course, a Fodorian nativist will here dissent. Fodor (1998). Concept nativism is obviously inconsistent with there being experience that is both consistent with apriority and necessary for concept acquisition. But it is consistent with the more general story I will go on to tell about non-warranting experience. As I will emphasize in this section and the next, not much actually hangs on concepts.

¹⁴ Cf. Burge (1993).

Apriority and Defeasibility

Some transitions in thought are rational, and others are not. It is rational, for example, under ordinary circumstances, to infer from <Obama went to Afghanistan> to <Obama flew on Air Force One>. The rationality of this inference is explained in part by background knowledge concerning, e.g., Obama's status as President of the United States, and the standard method of overseas transportation for the President. It is possible for a subject to have a different course of experience, such that the inference in question would not be a rational one; if one had experiences, for example, as of a perfectly convincing news report indicating that Air Force One had been dismantled last year—corroborated, let us suppose, by the experience of alternate sources giving the same news, the discussion of this surprising news by one's peers, etc.—then one would be rational in *declining* to make the inference in question.

I think that this kind of dependence on experience for rational inference underwrites the common-sense fact that the inference in question is a posteriori. However, we must take care to articulate the reasoning here properly. For an apparently-parallel argument could threaten the apriority of any inference at all. Consider: it is rational, under ordinary circumstances, to infer from <grass is green> to <grass is green or someone is hungry>. But arguably, there is a plausible course of experiences such that, if a subject had those experiences, it would not be rational for him so to infer. For example, someone might receive apparent expert testimonial advice that disjunctive elimination is not generally valid. True disjunctions, one might be misled to believe, require some kind of evidential or metaphysical connection between the disjuncts. Such advice would, I suppose, be misleading—disjunctive elimination *is* valid, and the intuitions requiring connectedness ought to be dismissed as Gricean pragmatic ones—but it might be sufficiently authoritative such that it would be irrational to ignore it. Given this possibility, it seems that the inference in question would be counted as a posteriori by lights similar to those marshaled above. And of course, the argument would generalize widely; once we see how experience can provide misleading evidence, it is plausible that some kinds of experience could undermine the rationality of any inference whatsoever.¹⁵

What this shows is that empirical defeasibility is not the correct model for the way in which a priori knowledge must be 'independent from experience.' That some possible experiences could *undermine* warrant does not show that contrary experiences are a part of the warrant. Compare: much of my knowledge is 'independent of testimony' in the sense that the rationality and knowledge-status of my belief does not derive from testimony. I now know that I am wearing a necktie, but this knowledge has nothing to do with anybody's having told me that I am doing so. My warrant for this belief is perceptual, not testimonial. But it is consistent with these obvious facts that testimonial experience could *undermine* my justification. If, perhaps, someone sufficiently authoritative told me that I had ingested a drug that was likely to confuse me about what I'm wearing, or that I am the subject of an elaborate prank involving a fake necktie, this might undermine my perceptual knowledge.

One way to attempt to draw the distinction more carefully might be to distinguish between, in Marcus Giaquinto's terminology, 'negative dependence' and 'positive dependence' on experience (or testimony, or whatever).

¹⁵ Williamson (2007), ch. 4 makes much of this fact. Williamson thinks this rules out any interesting notion of apriority along the lines of conceptual necessity. But it is possible to think that there are conceptual truths without thinking that they must be appreciated by all competent speakers; I will give an example of such a view below.

We could mark the distinction by saying that if a belief is rationally revisable in the light of future experience, its retention is *negatively dependent* on experience; and if a belief cannot have been justifiably acquired unless some experience was used as grounds in the process, its acquisition is *positively dependent* on experience.¹⁶

So developed, the suggestion would be that a priori knowledge is consistent with negative dependence on experience, but inconsistent with positive dependence. There is much that is intuitive about this distinction, and I am not sure that it is incorrect; nevertheless, I prefer a different characterization, for two reasons. First, the notion of ‘experience used as grounds,’ critical in the articulation of positive dependence, is not as straightforward to apply as might be hoped. Timothy Williamson, for example, draws attention to cases in which judgments are clearly informed—and rationally dependent on—experience, even though citing that experience as evidence would clearly be a mistake.¹⁷ This kind of feature is particularly plausibly present in various counterfactual judgments; to take one of Williamson’s examples, <if two marks had been nine inches apart, they would have been further apart than the front and back legs of an ant>.¹⁸ So unless positive dependence is characterized rather carefully, linking apriority closely to Giaquinto’s positive dependence risks potential counterexamples.¹⁹ Second, Giaquinto’s distinction does not relate in any clear way to the standard story given about experience that is necessary for concept possession.

I will attempt to characterize apriority in a way that provides a general answer to why both negative dependence reliance on concept-enabling experience is consistent with apriority, but positive dependence is not.

Propositional and Doxastic Justification

To properly characterize the a priori, I suggest, we should first clarify that apriority is most fundamentally a property of *propositional justification*. Propositional justification is to be distinguished from *doxastic justification* in that the former characterizes the abstract structure of the warrant for the proposition, where the latter describes whether a particular belief in a proposition appropriately capitalizes on a priori warrant. To say that, for some subject S, a proposition p is (propositionally) justified is to say that S has warrant for p; to say that S’s belief that p is (doxastically) justified is to say that S is warranted in believing p. A proposition might be propositionally justified for a subject without that subject’s belief in that proposition being doxastically justified; standard examples of such include cases in which a subject has good grounds in support of a proposition, but believes it on the basis of poor grounds; for example, one might ignore one’s conclusive evidence that his wife is being unfaithful, but come unreasonably to believe that she is being unfaithful on the basis of an astrological reading.

Apriority is fundamentally a property of propositional justification; we can define a sense of apriority for other states in the obvious ways. A belief is doxastically justified a priori if it is

¹⁶ Giaquinto (1998) p. 200.

¹⁷ Williamson (2007), ch. 5.

¹⁸ *Ibid.*, p. 167

¹⁹ If the central claim of Jenkins (2008) is correct, then this kind of dependence on experience, which goes beyond Giaquinto’s mere negative dependence, is ubiquitous among all a priori knowledge.

appropriately held on a priori grounds; someone knows p a priori if her a priori justified belief amounts to knowledge, etc.

Returning to the propositional sense, for p to be justified a priori is for there to be warrant for p that does not depend for its presence on any particular past or present experience. Just as the presence of propositional justification in general is neutral on the question of whether a belief in that proposition is doxastically justified, so too is the presence of a priori propositional justification consistent with a belief in that content's being justified or unjustified—or justified a priori or a posteriori. That a priori propositional and doxastic justification can so come apart explains why it is consistent with the apriority of some proposition—and indeed, the apriority of some belief in that proposition—that there are possible courses of experience that would defeat the doxastic justification for belief in that proposition.

Although a priori justification is in this sense empirically defeasible, Ichikawa & Jarvis (forthcoming) offered a principle for distinguishing a limited subset of ways in which it can be defeated. We suggested there that a priori (doxastic) justification can be defeated in only two ways: first, by a subject's rational shortcomings, and second, by evidence to the effect that one is experiencing such rational shortcomings. As we wrote:

Indeed, almost any rational commitment to infer can be defeated *due to a subject's rational limitations*. Rational limitations may arise from limitations in one's conceptual repertoire, limitations in computational capacity (e.g. in the time it takes to draw an inference), and tendencies to make performance errors in drawing inferences, or, for that matter, any other sorts of proclivities the subject has to make or exhibit confusions in attempting to execute in accordance with what he has reason to think. (This list may not be exhaustive, but it indicates that the term 'rational limitations' is not a catchall; rational limitations contrasts with limitations in experience that result specifically in a paucity of *evidence*, which might limit the subject's ability to competently infer in a wholly different way. Rational limitations are limitations concerning the processing of evidence.) Testimony from a panel of expert logicians can defeat John's rational commitment to infer in accordance with *modus ponens*, but only because of John's rational limitations vis-à-vis logic. If John were an acknowledged über-logician, he would have no reason to kowtow to the panel of "expert" logicians any more than we have reason to defer to elementary school children on matters of basic arithmetic.

Similarly, almost any rational commitment to infer can be defeated *due to evidence regarding a subject's (current) rational limitations*. Evidence to the effect that Jane has taken a pill that inhibits rational capacities can defeat Jane's rational commitment to infer in accordance with *modus ponens* as can evidence to the effect that Jane is crazy even if in fact Jane has not taken such a pill and is not crazy.

These two forms of defeat, we said, are the only ways in which a priori justification can be defeated.²⁰

This framework suggests a natural explanation for why experience involved in concept acquisition is consistent with the apriority of propositions containing those concepts: the relevant role of such experience is merely to reduce the relevant cognitive limitations. One way to fail to be in a position to take advantage of a priori warrant is to be unable, due to limitations in one's conceptual repertoire, to entertain the thoughts needed. Experience that plays a role in enabling such thoughts is experience that helps a subject to capitalize on the

²⁰ Our terminology then was somewhat different; we wrote of conceptual entailment, not apriority; our paper did not attempt to characterize apriority precisely. What we said then for conceptual entailment, I say now for apriority.

propositional justification he already has. So construed, this experience could be necessary for doxastic justification, but not for propositional justification.

§3. Enabling Beyond Concept Possession

In §2, I offered a theoretical motivation for the commonplace that experiences that are necessary for concept possession are consistent with the apriority of knowledge of propositions constituted by those concepts. Such roles for experience are theoretically interesting, but are of limited use for my present project of reconciling X-Phi and Apriority; according to the former, epistemologists cannot know, e.g., Gettier cases to be counterexamples to the JTB theory of knowledge unless they first engage in some folk survey methodology—but it is utterly implausible that such surveys are necessary for possession of KNOWS or any other of the concepts constituent in the conclusion of the Gettier argument. Experience necessary for concept possession is pretty clearly a non-starter with respect to my present project.

However, once we think of concept-enabling experiences in the way articulated in §2—as consistent with apriority because they merely help a subject to capitalize on propositional warrant that is already in place—it is not difficult to see that this role for experience will generalize beyond enabling concept possession. There are other ways to be merely enabling than merely to enable concept possession.

Consider the experiences typical students undergo when being trained in a priori disciplines like math, logic, and philosophy. Sometimes, students learn contingent facts via sensory experience—as, for example, when a student accepts an instructor’s testimony that the symbol that looks like ‘ \supset ’ symbolizes a truth-functional connective, with such-and-such truth conditions, etc. Sometimes, experiences play a role in establishing propositional justification. Prior to the relevant experiences, the students simply suffered no rational pressure to believe the relevant facts. Such knowledge, then, is a posteriori.

But often, the way students learn in such environments is quite different; often, students are trained to recognize rational commitments they already have. It is a confusion to affirm the consequent; doing so reflects a rational limitation, whether before the student is told so by the instructor or after. (Of course, we blame subjects for such failings more after they receive some training; we legitimately expect a higher degree of rationality of some than of others.) Sometimes, experiences play important roles in developing a priori cognitive abilities, which enables subjects better to capitalize on their extant a priori warrant; this represents another way in which experience can play an important role in eventual a priori knowledge. Although the experience plays no role in establishing the propositional a priori justification, it helps subjects to be able to capitalize on it, which ultimately can issue into a priori knowledge.

Plausibly, the vast majority of actual human a priori knowledge relies on experience in this way. We are creatures who develop our cognitive abilities by relying on ubiquitous experience; it is a gross understatement to say that a subject without the rich sorts of educational experiences with which we have been privileged would not often think particularly clearly. Without experience, we would not be able to have much knowledge at all—and this goes for the special case of philosophical knowledge, too. But this is not to say that philosophical knowledge cannot be a priori in the sense given above.

So there is no contradiction in the idea of experience—even experience unnecessary for concept formation—being necessary for philosophical knowledge, even though philosophical

knowledge is a priori. In §4, I will argue that the sorts of experiences that are, according to the experimentalist pressure discussed in §1, necessary for philosophical knowledge is indeed consistent with the apriority of philosophy. If the experimentalist critics are right, we must do experiments in order to get philosophical knowledge. Even if we must, this provides no reason to deny that the relevant knowledge, once achieved, is a priori.

§4. Experimental Philosophy and Apriority

Experimental Philosophy as Rational Improvement?

The case of rational training shows there are possible cases where subjects couldn't know that p without having some particular experiences first, even though p is a priori. This can happen when those experiences serve merely to decrease one's rational limitations, thereby making them better able to capitalize on a priori warrant. Can we understand the experimental work we suppose to be needed for philosophical knowledge in this way? To do so would be to suggest that engaging in experimental work is necessary for philosophical knowledge because it is the means we need to use in order sufficiently to develop our a priori capacities. This would be in effect to treat the relevant sorts of experiments as providing something roughly akin to the training students receive in logic classes.

The prospects for reconciling X-Phi with Apriority in this way do not seem especially strong. For one thing, the empirical facts to be discovered—the fact that epistemic intuitions vary according to cultural background, or that the messiness of a desk influences moral judgments—are deeply contingent. We could have been beings with more universal epistemic intuitions, or who processed moral judgments in a way independent from environmental features; the fact that we turn out not to be is something that could not, even in principle, be discovered a priori. By contrast, when I undergo the experiences that constitute my training in logic or mathematics, although it is of course contingent that I have the particular sorts of experiences I have, the new facts which I come to appreciate—that, for example, $\sim(A \vee B)$ is equivalent to $\sim A \ \& \ \sim B$ —are necessary. And even if I could not have come to appreciate this fact without having something like the experiences that I had—even if I needed to have the experiences of being taught logic in order to come to appreciate these truths—this is so only because of my rational limitations. In principle, someone clever enough could come to this sort of knowledge without having to be taught it. In general, when experience is necessary to enable a subject to capitalize on genuinely a priori warrant, it is at least in principle possible for a sufficiently rational subject to capitalize on it without that experience. But this is not the case with the relevant experimental results. Even a perfectly ideal rational agent with no cognitive limitations couldn't know, without having the appropriate experiences, that epistemic intuitions vary culturally; so if, as X-Phi says, such knowledge is necessary for philosophical knowledge, it cannot be so merely by being necessary for rational improvement.

Defeaters via Evidence of Rational Limitation

According to Apriority, there is a priori warrant for many philosophical claims. According to X-Phi, we can have no knowledge of these claims without experimental experiences. So the conjunction of these views entails that the absence of the relevant experiences defeats a priori warrant. In §3, I suggested that there are only two ways for a priori warrant to be defeated: via limitations in a subject's rational capacities, and via evidence that such limitations obtain. In the previous subsection, I rejected the idea that the relevant experiences are necessary because they are necessary for the achievement of sufficiently high rational capacities. We are left, then, with only one choice for modeling the experimentalist pressure in a way

consistent with Apriority. If we do not engage in experimental investigation, finding sufficiently promising results, and have the experiences that such investigation would involve, then our a priori warrant in philosophical propositions can be defeated by evidence that we are relevantly cognitively limited.

What does this sort of defeat look like? Here again is the example given in Jarvis & Ichikawa (forthcoming):

Evidence to the effect that Jane has taken a pill that inhibits rational capacities can defeat Jane's rational commitment to infer in accordance with *modus ponens* as can evidence to the effect that Jane is crazy even if in fact Jane has not taken such a pill and is not crazy.

This sort of effect occurs in more mundane contexts, too. Suppose I am evaluating two philosophers—in order, perhaps, to decide which one should be hired for a particular job. We can suppose that it's not a trivial decision—each has different strengths, and neither is *far* better than the other—but that in fact, Emily is better than Lauren, and I have an evidential base that makes this clear. This evidential base, plausibly, includes familiarity with their respective writings, research interests, motivation, teaching skills, administrative abilities, etc. It provides propositional justification for me for the proposition that Emily is better than Lauren. All things equal, it would suffice for knowledge. But now let us also suppose this complication: Emily is a dear friend, and I want very much, for reasons external to evaluating each as a philosopher, for her to get the job. Under such circumstances, it may well be that my warrant for the relevant proposition is defeated, and I won't be able to know which is better. This is because I, like most people, am susceptible to biases favoring people I like and outcomes I desire. I have a natural tendency to emphasize the positive features and ignore the negative ones, when it comes to people that I like. And I know this about myself. Under the circumstances, then, it may not be reasonable for me to trust myself in this tricky judgment; it would be too easy for me to think Emily the better candidate, even if Lauren really were preferable. This psychological data concerning my unreliability undermines my evidential warrant for Emily's preferability. Note also that it is inessential that I am actually susceptible to this widespread bias; if I know that almost everybody is susceptible to this bias, but I happen to be exceptionally immune to it, I still have strong evidence that I am susceptible, and this evidence still defeats my warrant. Of course, if I engaged in some careful self-experimentation and came to *know* that I was in this way exceptional, that would be a different story; such could put me in a position to capitalize on the a priori warrant after all. So there are two roles for psychological investigation that bear on knowledge of the non-psychological of identifying the better philosopher: it can provide evidence of rational limitations, thus preventing the use of propositional justification, and it can provide evidence restoring reasonable confidence in one's rational abilities, thus defeating the defeater and licensing the use of the non-psychological warrant.

That Emily is a better philosopher than Lauren is, of course, not a priori. The point of the example is rather to illustrate how experience can be relevant to establishing and undermining defeat in the form of evidence of rational failure. But the phenomenon is familiar in a priori reasoning as well. Suppose I wonder whether a certain sentence is a theorem of classical logic; after some moderate effort, I produce a proof of it, and come thereby to believe that it is indeed a theorem. My belief may not amount to knowledge, for it may be that, even if I have properly produced the relevant proof, it is unreasonable for me to rest confident on having done so. For, like the rest of us, I have strong empirical evidence that I sometimes engage in fallacious reasoning; sometimes, I think I've produced a proof, when in reality, I haven't. It is too risky for me simply to assume that I've been operating with sufficient rational capacities; it would be too easy for me to have gone wrong. In order

to come to knowledge of the proof's conclusion, therefore, I may need to double-check it. The sort of double-checking that can make it reasonable for me to be confident in my abilities involves, of course, experiences—the visual experiences involved in examining the marks on my page or blackboard, for example, that provide a record of my performance. These experiences are what it takes in order for me legitimately to be confident in my logical performance.

Here's another example: suppose I perform a moderately complex mathematical calculation in my head; I judge, say, that 26×31 is 806. We'll stipulate that I came to this answer competently in a way that, under ordinary circumstances, would lead me to knowledge that it is correct. But I go on to encounter an apparent epistemic peer who tells me that he has done the calculation in his head, and come to a different answer: 628. We can suppose that my rational confidence in my own mathematical ability is comparable with that in his, such that it would be unreasonably dogmatic simply to discount his answer on the basis of its conflict with the correct one; since it would be unreasonable, at this moment, to continue to believe that the answer is 806, I no longer know that fact. I restore my knowledge by having some experiences: the experiences involved in, say, talking over how we came to the answer, or checking to see whether he's any good at math, or using a calculator. My initial belief was a priori knowledge; it was defeated by evidence of rational limitations, and that defeater was overcome by experience. My experience restored my a priori knowledge.

I don't know of any simple way to explain in generality under what circumstances this kind of experiential validation is necessary for reasonable belief. I certainly don't think that it's *always* necessary—we don't always need to double-check, and sometimes we can ignore evidence that is misleading; the point is merely that it sometimes such experiences are necessary to restore reasonable confidence. And that when they are, the experiences involved are not properly thought of as part of the warrant—as part of the explanation for the propositional justification in question—instead, they are merely enabling. They do not enable concept possession, but they enable subjects to capitalize on their extant justification by defeating evidence that would otherwise be suggestive of disruptive cognitive limitations. They are therefore consistent with Apriority.

X-Phi and Reasonable Self-Confidence

The model of the previous subsection is the one to which we should apply the case of experimental pressure against armchair philosophy; indeed, it matches the interpretation of the experimentalist critique I argued for in §1. According to the experimentalist pressure, we now have substantial empirical evidence impugning our rational capacities, such that, to continue to rely on them unreflectively would be epistemically irresponsible. Learning that some cultural groups judge Gettier cases to be cases of knowledge impugns our rational abilities, since this shows that either we or they are wrong. Without some reason to privilege our own judgments over theirs, it would be unreasonably dogmatic to proceed on the assumption that it is we, and not they, who are right.

I should emphasize that I do not mean to be endorsing this line. This is how I think the experimentalist pressure ought to be understood. Obviously, there are many possible responses one could make at this stage in order to resist X-Phi. One could argue that we ought to interpret the relevant subjects as engaging with different scenarios, or using different concepts, and so do not disagree at all; or one could suggest that our judgments are more likely to be right, given our expertise, or our interested engagement. There is much one can say here, and most of it has been said. Again, my point in this paper is that *even if X-Phi is*

right, the experience necessary for the relevant philosophical knowledge is consistent with Apriority.

Other experimentalist challenges are also easily thought of in this way. We learn that order effects influence our epistemic intuitions; this shows that epistemic intuitions are easily influenced and so not particularly reliable; this should undermine our confidence that our judgments are not similarly unreliable. We could restore our initial reasonable confidence if we engaged in experimental investigation that showed that we philosophers are not similarly susceptible to order effects, or that our arguments did not rely on the kinds of intuitions that are susceptible to order effects.

Of course, in saying that we could restore rational confidence and knowledge in this sort of way, I assume that our initial judgments were indeed correct and a priori. I therefore assume that the experimentalist challenge can ultimately be met. But nothing much hangs on any particular case; if I am wrong about the Gettier intuition—if our ultimate investigation indicates that everyone, including philosophers, is unreliable about knowledge attributions in a way that undermines any claim to know Gettier cases to be non-knowledge, that shows that this particular case was not known a priori, since it was not known. We evaluate judgments on a case-by-case basis.

There is much room to consider and debate how far the skeptical pressure unearthed by the experimentalists extends. A conservative interpretation would have it that only particular patterns of rational failure need be investigated; one need only check to make sure there aren't order effects operative to restore legitimate confidence. A more radical line would suggest that the presence of order effects raises much broader worries: we've discovered that order effects influence intuitions; *who knows* what other epistemic dangers lurk? On the radical line, we'd need to do much more checking for potential sources of error in order to be confident in our rational abilities. Messy desks influence moral judgments; what other disruptive factors might be around? Maybe the smell of coffee disrupts one's ontological intuitions? We can't tell unless we do some experiments and check! The radical line demands quite a lot of experimental experience for philosophical knowledge; even still, the knowledge in question is a priori. The truth is surely somewhere between the conservative and radical lines here, but these questions are simply orthogonal to the apriority of philosophy.

Conclusion

Whether some knowledge is a priori depends on the source of its warrant. What I have demonstrated is that this question is separable from the question whether it is achievable from the armchair; that indeed, the questions may in principle come apart rather dramatically. Even though I have not attempted to establish Apriority or to judge definitively on X-Phi, exploring the connection—or, as I say, independence—between the two has, I hope, clarified both notions.

One takeaway moral is that epistemological questions about the nature and structure of warrant must be distinguished from methodological questions about the appropriate engagement with and processing of evidence. The methodological questions are epistemological ones too: they bear, as I have emphasized, on what is known. But they do so

in a way importantly distinct from questions about the source of warrant, and so from questions about apriority.²¹

References

- Alexander, J. and Weinberg, J. (2007), Analytic Epistemology and Experimental Philosophy, *Philosophy Compass* 2(1): 56-80
- Alexander, J., Mallon, R., and Weinberg, J. (forthcoming), Accentuate the Negative, *Review of Philosophy and Psychology*
- Ball, D. (2009), There Are No Phenomenal Concepts, *Mind* 118(472): 935
- Burge, T. (1993), Content Preservation, *The Philosophical Review* 102(4): 457-488
- Deutsch, M. (2009), Experimental Philosophy and the Theory of Reference, *Mind and Language* 24(4): 445-466
- Elga, A. (2007), Reflection and Disagreement. *Noûs* 41(3): 478-502
- Fodor, J. (1998), *Concepts: Where Cognitive Science Went Wrong*, Oxford University Press
- Giaquinto, M. (1998), Epistemology of the Obvious: A Geometrical Case, *Philosophical Studies*, 92(1): 181-204
- Greene, J., Sommerville, R., Nystrom, L., Darley, J., and Cohen, J. (2001), An fMRI Investigation of Emotional Engagement in Moral Judgment. *Science* 293: 2105–2108
- Hales, S. D. (2006). *Relativism and the foundations of philosophy*, Cambridge, MA: MIT Press.
- Ichikawa, J. (manuscript 1), Who Needs Intuitions? Available online: <http://jonathanichikawa.net/papers/wni.pdf>
- Ichikawa, J. (manuscript 2), The Audacity of Hopeful Philosophical Appeals to Intuition. Available online: <http://jonathanichikawa.net/papers/ahpai.pdf>
- Ichikawa, J. and Jarvis, B. (2009), Thought-Experiment Intuitions and Truth in Fiction, *Philosophical Studies*, 142 (2): 221-246.
- Ichikawa, J. (2010), Explaining Away Intuitions, *Studia Estonica Philosophica*.
- Ichikawa, J. and Jarvis, B. (forthcoming), Rational Imagination and Modal Knowledge, *Nous*. Available online: <http://jonathanichikawa.net/papers/rimk.pdf>
- Jenkins, C. (2008), *Grounding Concepts*, Oxford University Press
- Kauppinen, A. (2007). “The Rise and Fall of Experimental Philosophy”, *Philosophical Explorations* 10(2): 95-118.
- Knobe, J. (2006). The Concept of Intentional Action: A Case Study in the Uses of Folk Psychology, *Philosophical Studies* 130 (2):203-231.
- Knobe, J. and Nichols, S. (2008) An Experimental Philosophy Manifesto. In Knobe, J. and Nichols, S., eds, *Experimental Philosophy*, Oxford University Press: 3-16

²¹ I'm grateful for invaluable discussions about this material to Derek Ball, Yuri Cath, Benjamin Jarvis, Carrie Jenkins, and Jonathan Weinberg.

- Ludwig, K. (2007). The epistemology of thought experiments: First person versus third person approaches. *Midwest Studies in Philosophy* 31: 128-159
- Machery, E., Mallon, R., Nichols, S. and Stich, S. (2004) “Semantics, Cross-Cultural Style”, *Cognition* 92: B1-B12.
- Mallon, R., Machery, E., Nichols, S. and Stich, S. (2009) Against Arguments From Reference. *Philosophy and Phenomenological Research* 79(2): 332-356.
- Sosa, E. (2007). Experimental philosophy and philosophical intuition. *Philosophical Studies*, 132(1): 99-107.
- Swain, S., Alexander, J., and Weinberg, J. (2008) “The Instability of Philosophical Intuitions: Running Hot and Cold on Truetemp.” *Philosophy and Phenomenological Research* 76(1): 138-55.
- Weinberg J., Nichols, S., and Stich, S. (2001) “Normativity and Epistemic Intuitions”, *Philosophical Topics* 29: 429-60.
- Williamson, T. (2007), *The Philosophy of Philosophy*, Malden, MA: Blackwell.