Intuition Fail: Philosophical Activity and the Limits of Expertise

Intuition Fail: Philosophical Activity and the Limits of Expertise

Wesley Buckwalter

University of Waterloo

wesleybuckwalter@gmail.com

Abstract

Experimental philosophers have empirically challenged the connection between intuition and philosophical expertise. This paper reviews these challenges alongside other research findings in cognitive science on expert performance and argues for three claims. First, evidence taken to challenge philosophical expertise may also be explained by the well-researched failures and limitations of genuine expertise. Second, studying the failures and limitations of experts across many fields provides a promising research program upon which to base a new model of philosophical expertise. Third, a model of philosophical expertise based on the limitations of genuine experts may suggest a series of constraints on the reliability of professional philosophical intuition.

Keywords: expertise, intuition, expert performance, evidence, philosophical activity

Intuition Fail: Philosophical Activity and the Limits of Expertise

> *Even when the experts all agree, they may well be mistaken.*
> — Bertrand Russell, *On the Value of Scepticism*

1   Introduction

In 2008, the American College of Surgeons released some astonishing results from a consumer survey they had commissioned. They found that patients spend one hour on average researching a surgeon's credentials prior to scheduling a procedure. But comparatively, they spend roughly 5 times as long researching the purchase of a major appliance or piece of home furniture. They spent 8 times as long reading customer reviews before buying or leasing a new car, and 10 times as long when debating a change in jobs. In fact, their results showed that 36%, or roughly one in three Americans surveyed between the years of 2003-2008 did not bother to review the credentials of these expert medical professionals before going under the knife.[1] After all, they are the experts.

Let this illustrate the powerful allure of expertise. Experts possess very high levels of skill and technical proficiency. Due to their advanced training and experience, we typically trust experts, all else equal, over those lacking such proficiency. The wisdom and limits of placing this level of trust in our expert professionals is a question we will revisit at the end of this essay. The question we will consider presently, and one that has dominated recent metaphilosophical discussions in philosophy, is whether we should

---

[1] See "Survey Says: Patients Prep Harder for Vacation than for an Operation" <http://www.facs.org/news/russell.html>

extend an analogous level of trust to professional philosophers when it comes to philosophical intuition.

According to many challenges to intuition and philosophical expertise offered by experimental philosophers, the answer is "no". This paper reviews these challenges alongside other research findings in cognitive science on expert performance in several fields and argues for three claims. First, evidence taken to challenge philosophical expertise may, paradoxically, be well explained by the failures and limitations of genuine expertise. Second, studying the failures and limitations of genuine experts across many fields provides a promising research program upon which to base a new model of philosophical expertise. Third, a model of philosophical expertise based on the limitations of genuine experts also suggests a series of constraints on the reliability of professional philosophical intuition.

2   Philosophical Expertise

Philosophers have long been concerned with the nature and value of expert human performance. Socrates noted the importance of differentiating between experts and novices "in all matters which are considered learnable and teachable" when he observed in Plato's *Protagoras* that "we send for builders to advise us on what is proposed to be built," and that "if anyone else whom the people do not regard as a craftsman attempts to advise them" they will "merely laugh him to scorn and shout him down" (319b-c).

Today many philosophers continue to distinguish between experts and novices in matters of philosophical activity, and particularly, philosophical intuition (for variations see Deutsch 2009; Hales 2006; Ludwig 2007; and Williamson 2005, 2011). The same basic Socratic spirit above is closely embodied in the work of Steven Hales, for instance,

when he writes "Intuitions are and should be sensitive to education and training in the relevant domain" and subsequently that the "intuitions of professional philosophers are much more reliable than either those of inexperienced students or the 'folk'" (Hales 2006, p. 171-2).

The reason why Socrates and his fellow Athenians send for craftsmen in their deliberations is presumably the same reason why Hales and others have called for professionally trained philosophers in our own. By taking the advice of expert builders over non-builders, for example, you all but guarantee the creation of a better building. Likewise, it might be thought, philosophers are the experts of evaluating philosophical thought experiments.[2] Consequently, one should rely on the intuitions of trained philosophers over those of non-philosophers when building or evaluating philosophical theories.

The attempt to specify, develop and defend this claim of philosophical expertise in evaluating thought experiments is commonly referred to in philosophy as "the

---

[2] Some philosophers have argued that intuitions are hardly ever used as evidence in philosophical methods (Cappelen 2012; Kuntz & Kuntz 2011). This paper remains neutral on this point, focusing instead on the question of philosophical expertise (though also see responses by Chalmers forthcoming; Bengson forthcoming; Buckwalter 2012b). For a review of the majority of work in experimental philosophy that does not focus on expert intuitions and related methodological concerns, see Knobe et al. (2012).

expertise defense." [3] Perhaps the most straightforward articulation of this defense comes

from Kirk Ludwig in his paper "The Epistemology of Thought Experiments: First Person

Versus Third Person Approaches":

> What is called for is the development of a discipline in which general expertise in
> the conduct of thought experiments is inculcated and in which expertise in
> different fields of conceptual inquiry is developed and refined. There is such a
> discipline. It is called philosophy. Philosophers are best suited by training and
> expertise to conduct thought experiments in their areas of expertise and to sort out
> the methodological and conceptual issues that arise in trying to get clear about the
> complex structure of concepts with which we confront the world. (2007, p. 150-1)

While a number of articulations are available (Williamson 2005, 2011; Hales 2006;

Ludwig 2007; Horvath 2010; Grundmann 2010) such defenses typically turn on a simple

analogy. We give special evidentiary weight to the performances of experts in their

respective fields. We do so on the basis of the assumption that their performances are

more accurate or reliable than those lacking expertise in the target discipline. And

expertise in philosophy is no different. By way of this analogy, philosophers are best

suited to construct and evaluate thought experiments and reach intuitive verdicts in virtue

of their expertise in the discipline of philosophy. They are the expert intuiters.

3   Questioning the Analogy

Timothy Williamson, perhaps the most prominent advocate of the expertise defense,

writes that we should continue to uphold the analogy until evidence is provided to

question it. He observes that "from a sociological perspective, philosophy is a fairly

---

[3] Professional philosophers might be experts in many senses. The crucial question

presently at issue is whether or not they are 'expert intuiters', a phrase of art from

Weinberg et al. (2010).

normal academic discipline," and consequently, "since thought experimentation is a cognitive task distinctive of contemporary analytic philosophy, the initial presumption should be that professional analytic philosophers tend to display substantially higher levels of skill in thought experimentation than laypeople do" (2011, p. 221). Many philosophers find this initial presumption plausible and continue to endorse the general analogy advocated by Williamson and others.

By contrast, some experimental philosophers find this initial presumption implausible and have offered arguments and evidence suggesting that the link between intuition and philosophical expertise is disanalogous to expert performances in other disciplines. They have done so either by casting doubt on the claim that philosophers meet some of the basic empirical conditions required of experts as a class (Weinberg et al. 2010), or by demonstrating that philosophers fail to have intuitions that are any more reliable than those of non-philosophers (Machery 2011; Schwitzgebel & Cushman 2012; Tobia, Buckwalter & Stich 2013; Tobia, Chapman, & Stich 2013).

The most comprehensive challenge to date is given by Weinberg et al. 2010.[4] They discuss two main factors questioning the analogy between expertise in philosophy and other disciplines. The first factor is that philosophers lack robust sources of feedback regarding the correctness of their intuitions. The second and related factor, Weinberg et al. argue, is that philosophers lack more objective sources of evidence available to other experts upon which to calibrate and hone intuition. For example, a clinician may receive

---

[4] Also see Weinberg et al. 2013 for the related issue of whether philosophical intuitions may be improved by reflection.

direct feedback after making an incorrect diagnosis or misreading an x-ray. There are also other objective sources to evaluate her performance. But we typically lack the same kind of direct feedback concerning expert philosophical intuition into the nature of causation or justice. And compared to reading an x-ray, it is initially unclear, at least, whether there are similarly objective sources of evidence besides intuition with which to calibrate these judgments.

On this basis, the connection between philosophical expertise and intuition appears less straightforward than common examples of expertise in other professions. When facing many of the deep questions philosophers are interested in asking, we typically do not know which intuitions correctly capture the philosophical explanandum they were provided to explain. Neither are there clearly objective evidential sources with which to calibrate philosophical intuition. These observations lead Weinberg et al. to conclude that advocates of the expertise defense "can no longer offer quick armchair analogies between philosophers' judgments about thought-experiments, and those of the practitioners of other fields in their own home domains" (2010: 67). In their view, the analogy supporting typical deployments of the expertise defense itself requires empirical defense. There are even reasons to doubt it.

The comparison between philosophers and experts of other fields is further complicated by the fact that no one is quite sure which factors mediate superior performance in evaluating thought experiments. In other words, we do not understand the acquired skills, mechanisms or adaptations underlying expert intuitions. There are several proposals on offer. One study has shown that professional philosophers have a more reflective or inquisitive temperament than the general public, even after controlling for

overall level of education (see Livengood et al. 2010).[5] Alternatively, others have argued that philosophers possess a deeper understanding of complicated linguist or semantic theories relative to the content of their intuitions, and can apply these considerations more efficaciously in unfamiliar or novel cases (Devitt 2006). Still others have pointed to a cluster of skills involving the ability to focus on key philosophical details of thought experiments, while remaining minimally distracted by potentially irrelevant information (Ludwig 2007; Williamson 2007). These are all interesting hypotheses. But without a deeper understanding of the particular skills that actually mediate exceptional performance in evaluating thought experiments, we lack an adequate account of philosophical expertise in intuition.

Lastly, a growing body of empirical work has challenged the link between intuition and philosophical expertise directly. As mentioned above, a number of experimental philosophers have begun investigating the actual intuitions that professional philosophers report when they evaluate thought experiments. Results are typically taken to show that the various intuitions reported by philosophers are not always any more accurate or reliable than non-professional. Though this work is ongoing, the latest results continue to put pressure on the notion that philosophical expertise is best understood in terms of evaluating thought experiments and delivering superior intuitions (see Machery 2011; Machery 2012; Schwitzgebel & Cushman 2012; Schultz et al. 2011; Tobia, Buckwalter & Stich 2013; Tobia, Chapman & Stich 2013).

---

[5] Livengood et al. 2010 do not endorse the expertise defense, though it's possible their results could be recruited by others for this purpose.

4   A New Approach

We have reviewed four challenges for both the expertise defense of philosophical intuition, as well as the study of philosophical expertise and philosophical activity more generally. These points are made by way of the comparison to expertise and an analogy to professionals in other disciplines. First, when philosophers conduct thought experiments they often lack adequate indications of whether their intuitions are correct. Second, philosophers lack other objective sources with which to adequately calibrate intuition. Third, we do not know which factors actually mediate philosophical expertise in evaluating thought experiments. Fourth, there is some direct evidence regarding the performance of philosophers counting against the claim that philosophers are expert intuiters compared to other fields.

When taken together, these observations seem like an almost insurmountable challenge for linking intuition and expertise. There is direct evidence that may tell against it. The lack of evidence for the underlying factors to potentially support such an account of philosophical expertise is troubling. Moreover, if the points made about correctness or calibration are true, it is difficult to see how we could ever adequately understand the factors that mediate expert intuitional verdicts. Perhaps there are potential responses one might develop to any one of these challenges (e.g. Williamson 2011: 224). But taken at face value, they provide a daunting challenge to studying professional philosophical intuition, and to the extent that it exists, philosophical expertise.

While important research to date has challenged the existence of expert philosophical intuition, the challenge is not yet definitive and the question requires further study. Four observations motivate the further study of intuitional expertise. First, prior research comparing philosophical performance to expertise in other fields has only

focused on some features of expertise across those fields. As we will see, a large body of research in cognitive science on expert performance has yet to be explored and may provide evidence for similarities between philosophizers and other expert professionals that have gone unnoticed. Second, and with respect to philosophical activity specifically, the factors mediating philosophical performances in thought experiments are not understood, which renders claims for or against philosophical expertise largely undecided. Third, prior metaphilosophical discussions attempting to locate philosophical expertise have mostly been restricted to philosophical education, and typically involve studying the effect that acquiring a PhD in a branch of philosophy has on philosophical activity. However there is an array of other factors beyond receiving a PhD in philosophy through which one might locate philosophical expertise, such as through individual differences in experiences, talent, preferences, opportunities, and habits (Ackerman 1990, though also see Howe et al. 1999). These factors have not been studied by critics and any of them might contribute significantly to locating philosophical expertise. Lastly, some empirical evidence produced by experimental philosophers gives us some positive reason to expect intuitional expertise (Buckwalter & Phelan 2013; Buckwalter, Rose & Turri 2013; Turri 2013b). Such results suggest that in some cases, professional philosophical intuition did gleam crucial and accurate insights into the nature of important philosophical phenomena.

Though prior challenges do not yet forestall the existence of genuine philosophical expertise, they do encourage the development of a new approach to studying it. Such an approach would avoid difficult questions regarding intuitional correctness and feedback. It would address empirical evidence telling against

philosophical expertise to date. And it would provide the basis for broader, empirically informed comparisons between philosophical expertise and the nature of expertise seen within and across other disciplines. This is what this paper attempts to provide.

The inspiration for this new approach comes from Henry Wadsworth Longfellow, who writes, "Sometimes we may learn more from a person's errors, than from his virtues" (1848, p. 313). While it is important to understand philosophical progress and success, as Longfellow suggests, it may be just as important to study its failure. In the spirit of this approach, let's retain the basic analogy between philosophy and expertise in other fields endorsed by Williamson, Hales and others. And following this insight, let's continue to look for comparisons between philosophical and non-philosophical experts across other domains—this time with a twist. Instead of engaging in a direct comparison between the practices of professional philosophers and the various ways in which experts in other domains succeed, let us take the opposite approach and investigate the well-documented and domain-general ways in which highly trained *experts regularly fail*.

In the remainder of this paper, I survey the literature in cognitive science on the failures and limits of expert performance. Following past work I highlight four specific areas—including domain limitation, poor prediction, glossing, and bias— in which experts fail to outperform, and even underperform, novices. I then sketch the contours of existing research in experimental philosophy on philosophical activity and compare these results to the failures and limitations of experts more generally. The results of this comparison circumvent the challenges to studying philosophical expertise above (see section 5) and serve as the basis for three arguments. First, some evidence generated by experimental philosophers challenging expertise is consistent with commonly accepted

accounts of expertise developed across the social sciences. In many cases, the performances of expert philosophers are also well explained by some of the empirically well-researched failures and limitations of genuine expertise as they are by the absence of expertise. If true, much of the extant direct and indirect evidence taken to count against the existence of genuine philosophical expertise in thought experiments may need to be reevaluated. Second, evidence from cognitive science and experimental philosophy consistent with cross-domain failure can be used to ground a promising new approach to studying philosophical expertise and activity. If philosophers are prone to fail in similar ways that genuine experts predictably fail across a series of different professional domains, then these behaviors may be used to begin to build a model of professional philosophical expertise. In other words, such results can potentially be used to reverse engineer the beginnings of the first empirically supported account of philosophical expertise. As this research progresses, we can begin to gain better insight into the nature of expert philosophical activity, both its strengths and weaknesses. Third, new appreciation of these strengths and weaknesses can help us do philosophy even better in the future. Clearer insight into the nature of expert philosophical activity helps to identify when expert intuitions reliably inform philosophical argumentation, and when they might not.

5   How Experts Fail

As it turns out, we know a lot about experts. The literature across psychology, cognitive neuroscience, sociology, and computer science attempting to explain the factors mediating expert human performance is vast and growing. One general approach researchers have used to study these questions is to closely compare individual experts

and non-experts within and across domain. Under this kind of comparative method, human expertise is studied relative to the performance of agents of different levels of proficiency (see Chi 2006; Hoffman 1998). The basic idea of this approach is to compare the relative expertise between experts and non-experts in the hopes of understanding what abilities and underlying processes differentiate levels of expert performance, as well as how a novice rises to the level of expert. So, for instance, one research question would be to ask, what is it about an expert or master that distinguishes her from a novice?

This comparative approach to studying expertise is also the one that is implicated in the philosophical literature on the use of intuitions in philosophy. Simply stated, a crucial claim made by the expertise defense is that there are key differences between philosophers and non-philosophers (akin to the differences between the expert and the novice more generally) in virtue of which they are experts. A corollary to this claim is that professional philosophical intuition is more reliable as a result.

But one important thing we have learned from this research on experts in psychology is that the most accurate comparison between experts and non-experts as a class includes not only the general ways in which experts typically outperform novices. It also includes studying the general ways that highly trained experts fail to outperform and sometimes even underperform novices. While it is important to understand the factors that mediate exceptional human performance, it is also equally important to understanding the ways that experts typically fall short. This observation was notably made in Chi (2006)'s seminal review of the categories and research classifying the strengths and limitations of general expertise, which I rely on when comparing the ways that failures of expertise associated with domain specificity, prediction, glossing, and bias

may relate to professional philosophical activity (for discussion of these categories and research see Chi 2006, pp. 24-27).

Can basic lessons regarding how experts are prone to error rather than excel advance our understanding of philosophical expertise? If philosophers are experts in the traditional sense, then we should not only expect them to succeed like other experts, but also to fail like them. This observation clears the way for a new approach to studying philosophical expertise. Instead of looking to cases where philosophers get it right, let's opt for the opposite: spotting when and where they go wrong.

This section reviews the psychological research on the general conditions under which genuine expertise fails. Experimental philosophers have provided evidence for when philosophical intuition tends to be unreliable. If a general pattern emerges between these conditions and the failures of expertise more generally, then evidence to date may not challenge philosophical expertise. Rather, such evidence may serve as a new means for studying philosophical expertise and potentially reveal a better understanding of philosophical activity. To paraphrase Robert F. Kennedy, let's look to the experts, dare to fail, and see what we can achieve.

Before proceeding it is important to note that the following discussion compares the findings of two large research programs in cognitive science and experimental philosophy. The results of this comparison are suggestive. However the comparison is not a substitute for controlled empirical research. Only research dedicated to studying professional judgments directly can demonstrate that philosophical expert intuiters share the typical failures and shortcomings of genuine experts. Instead, the results of this

comparison are perhaps best viewed as a sketch for a promising theoretical framework guiding future understanding and study of philosophical expertise.

5.1 Domain Specificity

The literature on skills and expert performance suggests that genuine expertise is domain-specific. Professionals typically do not outperform, and can even underperform novices when completing tasks outside of their relevant domains of expertise. This tendency is perhaps best illustrated in famous studies by Gobet & Simon on chess grandmasters' ability to recall certain kinds of chessboard positions (1996a, b, c; see Chi 2006, pp. 24-26). In many of their studies, they presented chess players of various abilities with randomly selected textbook chessboard positions of about 25 pieces each. These positions were selected using a number of different criteria with respect to the domain of expertise. For instance, some criteria involved board positions reached after a certain number of moves during a match, others where the match was played by grandmasters, or where the gameplay was sufficiently obscure.

In one particular study, Gobet & Simon asked whether impairment of position recall interacted with player ability when positions were switched with their mirror image reflections. So once they had chosen the positions they wanted to administer, they altered the positions by taking their horizontal and vertical mirror images. As a control task, and to check and see if experts simply excelled at memory tasks outside of the domain more generally, Gobet & Simon also included some very unusual chessboard positions in their study. These positions had the same basic structure as the mirror image boards from before, but the individual pieces were rearranged and randomly assigned by computer to each of the previous spots. These five sorts of chessboard positions: the normal ones, the

mirror images (horizontal, vertical, and center), and the randomized pieces, were presented to players of various different skill levels (Masters, Experts, and Class A ranging from 1680 to 2540 ELO).

Gobet & Simon found that when the players were asked to recall the positions from memory, there was a large effect of expertise on the percentage of pieces that were recalled correctly, and that mirror image distortions did have a deleterious influence on the ability of players to report piece locations. These results are represented graphically below:
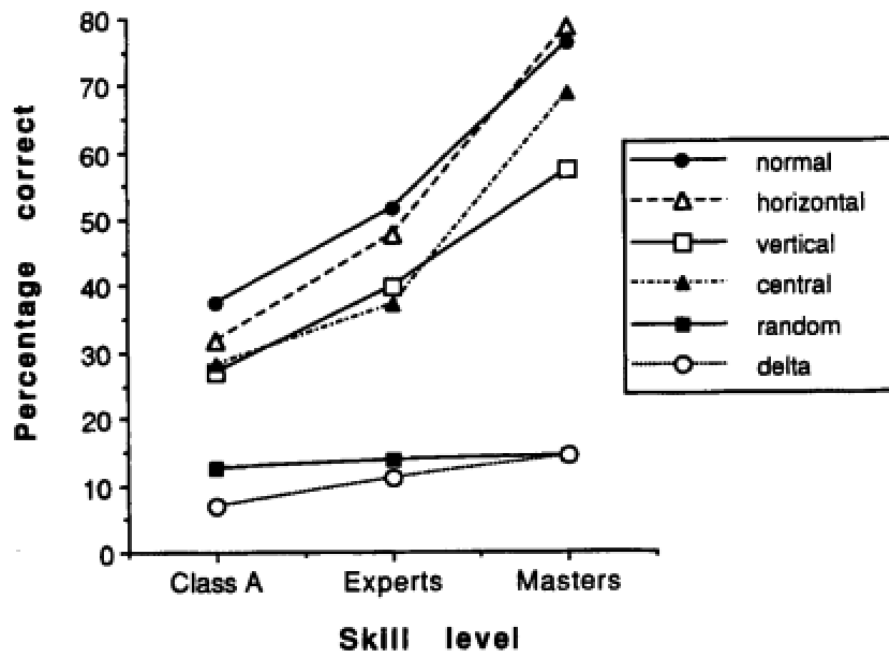

Figure 1. Results from Gobet & Simon (1996a) Experiment 2. Mean chess recall for different board positions grouped by expertise.

Gobet & Simon take these data as evidence supporting the need to modify certain aspects of the chunking theory of memory (as proposed by Richman, Staszewski, & Simon 1995) to include these more complicated aspects of various recall tasks. For our present

purposes however, this research begins to demonstrate how expert performance is often constrained by domain. While there were large effects for superior expert performance throughout the study, stronger players were not able to recall the random board positions (solid squared points, Fig. 1.) significantly better than weaker players. In other words, the superior performance of the expert chess player is restricted because one of the skills thought to mediate higher performance only operates within certain situations. That is, superior position recall only seems to lead to superior results when the players were presented with board positions that have been dubbed relevant and meaningful within the game of chess. This same basic result has been found again and again, not just for chess players, but also for a series of other domains (Voss et al. 1983, Ericsson & Lehmann 1996; Vicente & Wang 1998) and especially in understanding physical systems (Vicente 1992).[6]

Return now to the analogy between expertise in philosophy and other disciplines. If philosophers are experts with respect to delivering intuitions in philosophically meaningful thought experiments, then we should also expect superior performance to be constrained by domain of expertise. In other words, under this picture professional intuitions should analogously, only excel in those cases that have been dubbed relevant and meaningful with regard to various areas of philosophical study. And of course, within philosophy there are a number of diverse fields and foci of inquiry.

---

[6] Also see Roediger & McDermott (1995); Smith et al. (2000) for how increased knowledge in a domain can directly limit performance in other domains.

According to a domain-general model of expertise, it would be surprising if an expert ethicist, for instance, performed at the same level as an expert epistemologist or philosophers of language within the other's specialty. If the general analogy to other disciplines holds, then the general skills that these professionals bring to bear on responding to thought experiments will probably not result in superior performance outside of the domain of expertise. Philosophical expertise, like expertise across a number of disciplines, is not a blank check.

There may be some empirical precedent for this finding among philosophers. For instance, researchers have recently found that both non-philosophers and professional philosophers are subject to a certain framing effect when making judgments of moral permissibility and moral obligation: the actor-observer bias (Tobia, Buckwalter & Stich 2013). Specifically, researchers found that when both groups were presented with the famous "Jim and the Natives" case originally proposed by Smart and Williams (1978) or standard Trolley problems, the intuitive judgments both groups gave when they are the actor (or in response to cases framed in the first-person) differ from the intuitive judgments given when they are the observer (or cases framed in the third-person). If one accepts that whether an action in a moral scenario is framed in the first or third person is irrelevant to the truth of the moral judgment, then the intuitions of professional philosophers are also influenced by factors that are irrelevant to the truth of the intuition. If professional intuitions are influenced by such factors, then they may not be more reliable than those of non-professionals.

But interpreting these results requires caution. As these researchers themselves note:

One objection to our conclusion may be to argue that because only ethicists have expert moral intuitions, the data here do not pose a challenge to the expertise defense, since they represent the intuitions of a wide range of philosophers, including metaphysicians, philosophers of language, epistemologists, etc., rather than only those of ethicists. The objection is well taken…Our findings do not pose a significant challenge to the advocate of the expertise defense who maintains that only professional ethicists are experts in moral intuition. (p. 7)

Another way to interpret these findings is that just as in other fields like chess or engineering, expertise in philosophy is domain specific. In this case, intuitions in these central thought experiments of moral philosophy reported by philosophers who specialize in some other area of philosophy outside the relevant domain will also likely suffer. Superior performance is specialized to their domain of specialization.[7]

These findings do not demonstrate that professional philosophers are expert intuiters, but they do suggest an alternative explanation of the evidence available. Philosophical expertise, like genuine expertise across many fields, may be very domain specific. In other words, such findings need not threaten philosophical expertise understood above. There's another possible explanation right in line with what a domain

---

[7] See also recent work by Nado (2012) in favor of the claim that philosophical intuitions are the result of a number of disparate mental mechanisms and by Machery (2012) for direct evidence of a relationship between areas of professional specialization and confidence in Kripkean intuitions about reference. By contrast, see recent experimental findings by Schwitzgebel & Cushman (2012) suggesting that philosophers specializing in ethics display order effects in their judgments about some ethical scenarios but not in others, and by Schulz et al. (2011) for an effect of the heritable personality trait extraversion on free will intuitions after controlling for domain knowledge.

general model of expertise would predict. They may reflect domain specificity of expert philosophical intuition.

Moreover, such findings may constitute a first step toward providing a model of philosophical expertise grounded by the limitations of genuine experts. If philosophical expertise fits the model of expertise in other fields, we should expect the link between philosophical expertise and intuitions to hold strictly within the domain of expertise. To that end, future empirical work supporting expert performance in philosophy might profitably explore the effects of philosophical specialization on philosophical judgments in greater detail. In the meantime, this research suggests that if philosophers are expert intuiters, their expertise is most likely restricted to the area of specialty and their domain of expertise.

5.2   Inaccurate Prediction

Philosophers have long argued that ordinary behavior, and in some cases the ordinary application of certain concepts, can be very useful in guiding and building certain sorts of philosophical theories (Aristotle [1984]; Reid 1785 [2002]; Moore 1942; Austin 1956). The basic idea is that these judgments, in certain circumstances, can serve as valuable evidence in philosophical arguments. We see many examples of such arguments involving appeals to ordinary behavior within recent debates in epistemology and in metaphysics.[8]

---

[8] For reviews of the former see Pinillos (2011, 2012) and Buckwalter (2012). For an excellent discussion of the role of ordinary intuitions in research on the metaphysics of causation see Collins, Hall, and Paul (2004, 30-39).

Perhaps nowhere is this more evident than recent work in epistemology on ascriptions of knowledge. For instance, Jason Stanley motivates his defense of interest relativity in part based on "the intuitive reactions we have," to pairs of cases that manipulate a subject's practical interests. His claim is that our intuitions regarding knowledge ascription are stakes sensitive, and that his "central interest is to evaluate accounts that make as much sense of these intuitions as possible" (2005). Likewise, John Hawthorne touts that his view offers "the best hope yet for respecting the intuitive links between knowledge, assertion, and practical reasoning" (2004). In related discussions concerning epistemic contextualism, Keith DeRose claims that some of "the best evidence" for this view comes from "what ordinary speakers will count as 'knowledge" (2005).

Before drawing on ordinary behavior in philosophical argumentation, one must first understand what ordinary behavior is like. To find out what ordinary behavior is like, philosophers often turn to intuition and thought experimentation. In the epistemic literature for instance, claims of ordinary practice are often based on the familiar appeal to intuition about knowledge attribution, together with armchair claims about "what people will say" in response to philosophical thought experiments. If philosophers are expert intuiters however, we should expect these performances to resemble the performances of other experts when predicting ordinary responses to their respective areas of inquiry. That is, we would expect them to exceptionally bad at it.

The psychological research on expert performance suggests that experts consistently fail in making predictions about the behavior of novices or non- experts. In a series of studies, researchers have shown that one's level of expertise is negatively

correlated with making correct predictions regarding novice performance (see, for instance, Hinds 1999). They have also shown that experts typically fail to make better predictions than non-experts regarding the future behavior of others under conditions of uncertainty—such as predicting which prospective students will be successful in graduate school (Dawes 1971) or in evaluating medical internship candidates (Johnson 1988).

For instance, consider a famous study by Hinds (1999). Hinds compared the predictions that expert, intermediate, and novice participants made about the performances of other novices. In these studies, members of each expert level were asked to make predictions about the novice completion of a task, such as "using advanced cellular telephone technology" (though recall the year was 1999), or the assembling of a set of LEGOs (the Star Wars V-Wing Fighter). In both cases, Hinds found that experts were significantly worse than novices at successfully predicting novice performances in these tasks.

Hinds also examined how certain procedures might aid experts in making better predictions about novice behavior. One such technique was to "prompt people to recall their own experiences and instruct them to use this experience to construct a scenario of how the task being estimated may progress" (p. 207). Another technique was to present participants with a list of potential problems faced by novices prior to collecting predictions about their behavior in the tasks. However Hinds found that experts were significantly less responsive to both the list or recall procedures. They continued to be

significantly worse predictors of novice behavior than the intermediate or novice participants across the studies.[9]

One hypothesis offered to explain these findings is that experts struggle when taking the perspective of non-experts. Hinds speculates of this case that, "Experts may anchor on their own performance and fail to adjust adequately for the differences in skills between themselves and novices. People often anchor on their own attitudes, beliefs, and knowledge and use this anchor as a basis for predicting what others believe, feel, and know (Davis, Hoch, & Ragsdale, 1986; Nickerson, Baddeley, & Freeman, 1987)." (Hinds, 1999, p. 206). An expert's experience and knowledge can interfere with perspective-taking and, consequently, cause them to poorly predict novice behavior.

Are professional philosophers just as bad as other experts at predicting novice behavior? Countless results accumulated by experimental philosophers suggest that the answer to this question is an emphatic "yes". These data suggest that philosophers are also strikingly bad at making predictions about ordinary behavior, or how ordinary people will respond to philosophically relevant questions or thought experiments across several areas of philosophical inquiry (Arico & Fallis 2013; Batson 2008; Bengson et al. 2009; Braddock 2010; Church et al. 2005; Faraci and Shoemaker 2010; Feltz et al. 2009; Livengood & Machery 2010; Liao et al. 2012; Murray et al. 2013; Myers-Schulz & Schwitzgebel 2013; Reuter 2011; Starmans & Friedman 2012; Swain et al. 2008; Sytsma & Machery 2010).

---

[9] Interestingly, despite the effect for poor expert prediction, Hinds also found that intermediates generally tended to be among the best predictors of novice behavior.

Returning to the example of practical interests and knowledge ascription, the questions of stakes and interest relativity serve as a case in point. While there has not been a systematic study of professional philosophical judgments on the matter, a dominant assumption in contemporary epistemology is that ordinary ascriptions of knowledge are stakes sensitive (see Buckwalter 2014 for discussion). The results produced by several independent researchers suggest that professional predictions of ordinary practice are wrong (Feltz & Zarpentine 2010; May et al. 2010; Buckwalter 2010, and Schaffer & Knobe 2012). Using a series of sophisticated empirical techniques, the latest round of studies strongly suggesting that there is no distinctive unmediated epistemic role for stakes on knowledge judgments (Buckwalter & Schaffer 2013; Turri & Buckwalter under review).[10]

Again, these findings do not demonstrate that professional philosophers are expert intuiters. But the results are suggestive. If philosophers are notoriously bad at predicting the behavior of non-experts, then this just as well supports the link from intuition to philosophical expertise. These findings are right in line with what a general model of expertise would predict. Such findings constitute another step toward providing a model of philosophical expertise grounded by the limitations of genuine experts.[11] If

---

[10] See Buckwalter & Schaffer for responses to experimental challenges to stakes insensitivity from Pinillos (2012) and Sripada & Stanley (2012). They provide a competing model for the role of error salience in ordinary epistemic judgment.

[11] Presumably we might expect laypeople to make more accurate predictions of lay intuitions than professional philosophers simply in virtue of reporting their own intuitive

philosophers are experts in the traditional sense, then we should expect them to err when predicting ordinary behavior. And despite the philosophical value of ordinary behavior, poor prediction of ordinary intuition appears to be one of the most consistent findings in all of experimental philosophy. To that end, future empirical work supporting expert performance in philosophy might profitably explore the overall accuracy of philosophical prediction. For instance, future work might conduct a representative survey of failed expert predictions in philosophy, and experimental evidence of such failed predictions of the kind that one finds in the psychological literature on expertise. In the meantime, this research suggests that if philosophers are expert intuiters, their expertise is most likely at the cost of expert prediction of novice behaviors.

5.3 Glossing Over (In)significant Details

One virtue often attributed to professional philosophers is their keen ability to pay attention to key philosophical details when evaluating thought experiments. As Williamson writes, "The expertise defence does not imply that a good philosophical education involves the cultivation of a mysterious sui generis faculty of rational intuition, or anything of the kind. Rather, it is supposed to improve far more mundane skills, such

---

reactions. While dedicated empirical work is necessary to fully test this claim, some evidence for it is available. For instance, Turri, Buckwalter & Blouw (in press) has shown that patterns in ordinary knowledge attribution are more attuned to certain variations with respect to types of Gettier case constructions than the professional philosophical literature has been, suggesting that predictions of ordinary behavior made on this basis will likely be more accurate than those of professionals.

as careful attention to details in the description of the scenario and their potential relevance to the questions at issue" (2011, p. 216).

The basic idea seems to be that due to their advanced training, philosophers are better than novices at holding fixed in their minds certain information of philosophical importance, while dismissing other features of thought experiment that are irrelevant to the philosophical question at issue. This is indeed an important skill. If philosophers possess it, then they would also share this important skill with experts from many different fields. We know that generally speaking, experts possess the remarkable ability of being able to overlook visual or surface details, and understand the deep or underlying structure of problems relevant to their intended focus of inquiry (see Chi et al. 1981; compare: Voss 1980; Schmidt & Boshuizen 1993a; and discussion by Chi 2006). But although this ability is often considered a strength of genuine expertise, it can also be a weakness.

Researchers have found that although experts excel at understanding deep structures of problems, they often fail to process important surface details. For instance, Voss et al. (1980) found that when participants of different levels of expertise were presented with a narrative account of a fictitious baseball game, "high knowledge" individuals recalled more information relevant to the structure of the game than "low knowledge" participants. But they also found that low knowledge individuals were able to recall more information than the high knowledge participants regarding auxiliary and non-game action. Specifically, they found that participants with greater domain knowledge generated more complete descriptions of the mechanics of plays, in contrast to participants with less domain knowledge, who were more likely to provide more

details concerning "fans' reactions and thoughts, the pressure of the game, and so on" (p. 659).[12]

Researchers have also obtained similar results in the medical domain by comparing the performances of medical students to those of general practitioners (Boshuizen, Schmidt & Coughlin 1987; Schmidt & Boshuizen 1993a,b; Patel et al. 1989). For example, in a famous study by Schmidt & Boshuizen, medical students and internists were presented with vignettes describing a hypothetical scenario in which a patient was brought into the ER displaying certain symptoms (acute bacterial endocarditis). After different lengths of time examining the case, the participants were then asked to recall as many propositions as they could about the patient's condition, as well as make a diagnosis. What Schmidt & Boshuizen found was that advanced students were able to recall significantly more propositions about the endocarditic patient than the novices and the expert internists.

_____
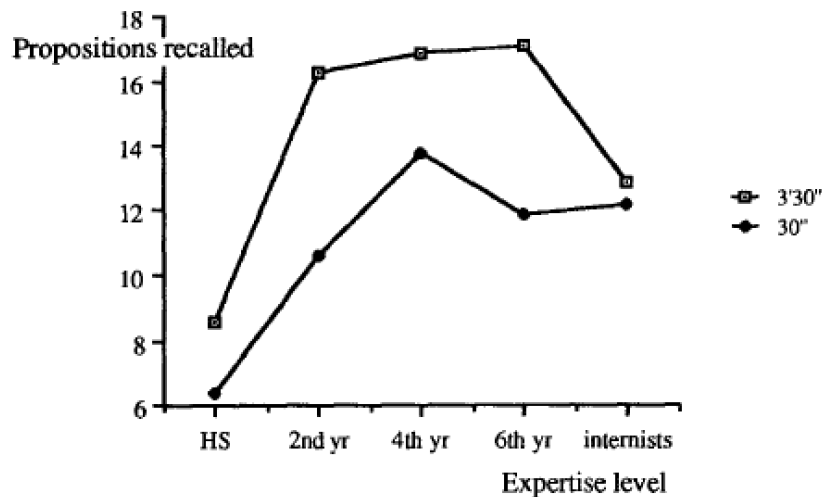
[12] Also see Adelson (1984) for similar results in the domain of computer science.

Figure 2. Results from Schmidt & Boshuizen (1993a). Propositions recalled based on time with case and level of expertise.

Schmidt & Boshuizen explain these results by appeal to the idea that the way fourth and sixth year (or pre-clinical, intermediate) students comprehend the case depends on representing and accessing detailed pathophysiological concepts. Alternatively, they hypothesize, internists' (or the experts) comprehension of the case does not rely on this kind of intensive processing and recall. Instead experts approach the case by using certain shortcuts or cognitive heuristics that they have developed in the course of their extensive clinical experience. This training allows them to make patient diagnoses while also glossing over several pathophysiological processes picked up on by the students lacking the same kind of extensive clinical experience. This, in turn, explains why the internists were unable to recall as many propositions regarding the medical status of the patients in Schmidt & Boshuizen's experiments. They simply glossed over them when processing the case.

Are professional philosophers prone to the same sort of detrimental glossing effects as sports and medical experts? Here we do not have direct empirical evidence of this basic phenomenon among professional philosophers. The relevant experiments have not been conducted. However, many observations support this hypothesis and serve as good initial, indirect evidence that some areas of philosophy emit precisely this sort of glossing.

Consider, for instance, the literature in epistemology surrounding the issue of epistemic purism. Epistemic purism (also sometimes referred to as intellectualism) is the

thesis that, "for any two possible subjects S and S', if S and S' are alike with respect to the strength of their epistemic position regarding a true proposition p, then S and S' are alike with respect to being in a position to know that p." (Fantl & McGrath, 2007, p. 558). Purism states that whether or not one is in a position to know that p depends only on *truth conducive* factors. In other words, orthodoxy has long suggested that "the right sort" of factors for knowledge are ones like evidence or reliability, which increase the likelihood that a belief is true.

Recent work in contemporary epistemology and across experimental epistemology on knowledge ascription has suggested a very different picture. Researchers have found that knowledge ascriptions may often be governed by many of what epistemologists have historically considered the "wrong sort" of factors thought to make a difference for knowledge (for a review, see Buckwalter 2012a). By contrast, these results tend to support a general picture whereby many important non-truth-conducive factors influence knowledge ascription.

To cite one recent example, experimental work has found a powerful relationship between knowledge and actionability (Turri & Buckwalter under review). Specifically, these researchers find that whether a person should pursue a course of action has a powerful and direct relationship to knowledge ascription. In fact, mediation analysis conducted by these researchers suggests that judgments about "how one should act" influence judgments about "what one knows" as much as judgments about truth and evidence do. These results powerfully vindicate recent theoretical work in philosophy on the role of knowledge in licensing certain activities (Unger 1975; Williamson 2000;

Hawthorne & Stanley 2008), and lend support to the view that there is a deep and important connection between knowledge and actionability (Fantl & McGrath 2009).

Another example involves the apparent influence that normative factors—and in particular people's moral judgements—have on epistemic judgments (Beebe and Buckwalter 2010; Beebe and Jensen 2012; Buckwalter 2013). This basic effect on knowledge ascription, known as the "epistemic side-effect effect" takes its inspiration from Joshua Knobe's "side-effect effect" for intentionality (2003), arguably one of the most famous and well replicated effects in all of experimental philosophy. Specifically, what Beebe & Buckwalter have shown is participants are significantly less likely to agree that an agent knew that her actions would bring about a particular side-effect when the outcome was good, and more likely to attribute knowledge when the outcome was bad. This finding has been replicated and extended, and suggests a powerful relationship between normative or evaluative judgments on the one hand, and epistemic judgments on the other.

The effects of normative and pragmatic judgment on knowledge ascription have not been discovered until relatively recently. If one accepts the form of purism above, then these features should be dismissed when processing a thought experiment. On the traditional view, they are surface details of a thought experiment removed or altogether irrelevant to the deep structure of knowledge. To better approximate the deep structure of knowledge, philosophers for centuries have focused instead on factors such as truth or evidence. Some philosophers had hypothesized about these variables theoretically (Fantl & McGrath, 2009; Hawthorne & Stanley 2008; James 1948; Locke 1975), but it was only

recently, and aided by experimental methods, that philosophers did begin to discover how factors like actionability or normativity affect ordinary knowledge judgments.

Of course, the correct explanation or interpretation of these knowledge ascription practices might well remain a matter of debate (e.g., Alicke et al. 2011; Knobe 2010). For our present purposes, the key thing to realize about these effects, is that up until relatively recently, philosophical orthodoxy implicitly judged them to be irrelevant to the problem at issue. One explanation for this is that these factors—along with other factors recently discovered in contrast to purism—were potentially concealed by expert susceptibility to glossing when evaluating thought experiments. In other words, they were not processed, or processed and dismissed as factors of the case irrelevant to the philosophical question of inquiry. This situation is exactly what we would expect if expert intuitions in philosophy share the same limitations as experts in other fields.

Such findings constitute another step toward providing an adequate model of philosophical expertise. If philosophical expertise fits the model of expertise in other fields, we should expect glossing effects in professional philosophical inquiry. Thus future empirical work on philosophical expertise might profitably study susceptibilities to the kind of glossing effects found in the psychological literature. Future research might also compare the susceptibility of professional philosophers and non-philosophers when evaluating philosophical thought experiments. In the meantime, this research suggests that if philosophers are expert intuitors, their experience and extensive knowledge of philosophical phenomena will likely lead to glossing at the level of case processing, which has the potential to influence their intuitions as a result.

5.4 Expert Bias

David Lewis began "Elusive Knowledge" by considering the clash between infallibilism and the bulk of the propositions we ordinary take ourselves to know (1999). Classic infallibilism is the thesis that a subject must rule out just about every possibility of error in order to know that p—fallibilism denies this. If forced to choose between the two, Lewis reluctantly admits that he would give up the philosophically preferred position of infallibilism for fallibilism, rather than embrace almost total skepticism about everyday knowledge. But Lewis is not fully satisfied with the prospect of fallibilist knowledge. To convince philosophers who have accepted this thesis likewise, Lewis offers them the following plea:

> We can get used to it, and some of us have done. No joy there—we know that people can get used to the most crazy philosophical sayings imaginable. If you are a contented fallibilist, I implore you to be honest, be naive, hear it afresh. "He knows, yet he has not eliminated all possibilities of error." Even if you've numbed your ears, doesn't this overt, explicit fallibilism still sound wrong? (1999, p. 418-419).

What should we make of Lewis' request to be "honest" and "naive" when considering fallibilist conjunctions? One explanation, of course, is that introspection and intuition of this sort are often closely tied to biases that accompany one's prior philosophical commitments, training, and experience.

If philosophers are subject to this kind of professional bias, then they are again, in very good company. Research suggests that expertise in a number of specific domains can lead to specific errors within that domain due to the bias that accompanies greater domain knowledge. An indicative example of this tendency can be seen in Castel et al.'s (2007) "The Dark Side of Expertise." In their studies, Castel and colleagues present expert and novice participants (individuals with either high or low levels of knowledge about American football) with a random assortment of animal names in a recall task.

These animal names also happened to be the names of NFL teams (e.g. dolphins, broncos, falcons, etc).

Castel et al. predicted that although these objects were familiar to both groups of participants, high-knowledge and low-knowledge participants would process names differently. Presumably, they thought, experts will rely on their domain specific knowledge of football to organize the names of the animals, giving rise to superior performance in the memory task over the novices. But Castel et al. also hypothesized that this comes with a cost. Relying on this kind of schema for recalling the animal names may also lead experts to make more mistakes. Specifically they predicted that experts would misremember animal names that hadn't been presented on the list (or that they hadn't previously represented), just because they also happened to be NFL teams, too.

That's exactly what happened. What they found was that indeed, high-knowledge participants were able to correctly recall more animal names than low-knowledge subjects. However they also found that the high-knowledge participants, or the football experts, were also significantly more likely to incorrectly recall teams that were not represented. In other words, the experts were more susceptible to "intrusion" errors during recall. Castel et al. take theses results to show that "under some circumstances, the organizational processing that benefits experts also has a ''dark side''; specifically, it can lead to recall of domain-relevant information that was not presented" (p. 4).

Similar effects of expert bias have also been shown in remote association tasks (see Chi 2006, p. 27 for discussion). For instance, Wiley (1998) discovered that although expertise in a domain could increase performance in these kinds of tasks, it could also work to greatly interfere with creative problem solving by promoting the onset of

something known as "functional fixedness." Wiley's experiments featured participants who possessed a low degree or high degree of knowledge in baseball. Both groups were presented with several different sets of three words each. For example, included in one set of words was 'plate', 'broken', and 'shot'. They were then asked to supply the fourth word that can be successfully combined with each of the three individual words. (The correct answer for this set is 'glass'.)

But Wiley also varied something else about the sets: whether or not the words primed participant knowledge of baseball (the example above is a set of this type, since it begins with the word 'plate'). She found that when experts were primed in this way, they were significantly less likely to find the correct answer then low-knowledge participants when the correct answer had nothing to do with baseball. In this way, baseball knowledge was a strike against expert participants. It led them down a path where the correct answer was not likely to be found.

Findings from Wiley and others on bias, functional fixedness, and the "dark side" of expertise are perhaps best summarized by the Buddhist philosopher and Sōtō Zen monk Shunryu Suzuki, who writes, "In the beginner's mind there are many possibilities, in the expert's mind there are few" (1970). Experiments show that greater domain knowledge can constrain creative thinking. Specifically, expert knowledge often constrains the search space of answers to the domain of expertise. Sometimes this can be advantageous to problem solving. But it can also prove disadvantageous when more

general solutions to problems are called for. In short, novices outperform experts when experts favor biased solutions.[13]

It remains to be seen if professional philosophers directly suffer from the same kinds of biases when evaluating thought experiments. Though such a result would hardly be very surprising. After all, natural scientists quite generally have long identified the concern of research bias in designing and conducting experiments (see Rosenthal and Rubin 1978). We should expect no less from thought experiments. And as noted earlier, and unlike the natural sciences, we typically do not get direct or as robust feedback concerning intuitions in thought experiments. This would suggest potential for significant bias towards one's favored philosophical research program.

There is some direct empirical evidence that philosophers are susceptible to bias during the process of constructing philosophical thought experiments. For example, Strickland & Suben (2012) have shown the problem of experimenter bias extends to experimental philosophy itself.[14] To demonstrate this, Strickland & Suben conducted an experiment within an experiment. They divided up two groups of Yale undergraduates and asked them to design a study investigating a certain philosophical question. The

---

[13] Many studies have also shown bias in the medical domain, either in how outcome probabilities influence treatment decisions (Christensen et al. 1991) or how greater expertise in sub-specialty biases patient diagnoses (Hashem et al. 2003; Walther et al. 2003).

[14] Also see Hansen (2013), for work on bias in thought experiments involving shifting epistemic contexts.

philosophical question was about group mental states and the concept of phenomenal consciousness (e.g. "Can Acme Corp., the group, experience great joy?").[15] However they also supplied each group with the hypothesis they were meant to test. One group received the hypothesis that a certain factor, the biological embodiment of an entity, was a crucial psychological cue for phenomenal state ascription. The other group received the exact opposite assignment—to test the hypothesis that embodiment was *not* crucial for state ascription.

Strickland & Suben found that when they proceeded to conduct the experiments that each of these two groups of participants had designed, they found opposite results. Both experiments turned up some evidence in support of the (rival) hypotheses they were given! Strickland & Suben conclude that foreknowledge of the hypothesis, what the experimenter has in mind, and particularly the goals of the experiment, led to significant bias affecting the outcome of the experiment. These results point to an important source of bias, and suggest the need for important safeguards against experimental bias when conducting empirical research. Without such safeguards, Strickland & Suben write, researches have "ample opportunity to mold and craft one's stimuli in such a way that the expected or desired outcome is more likely to be obtained than it should be" (p. 3).

Professional philosophers not only evaluate thought experiments, but also construct them. They almost always construct them with foreknowledge of hypotheses, just as Strickland & Suben warn against. Moreover, they almost always construct and

---

[15] For more on the question of biological embodiment as it relates to phenomenal state ascription, see Knobe & Prinz (2008); Phelan & Buckwalter (2012).

evaluate thought experiments for the very purpose of achieving the intuition that they wish to achieve. This is not always a bad thing. Flexibility in constructing thought experiments can sometimes be a very good thing. But there are also clear risks associated with this practice. As results demonstrate, with greater flexibility in constructing and evaluating philosophical thought experiments comes greater chance for bias.

This kind of bias is right in line with the kind of bias accompanying expert professionals across a number of domains. At the same time, much more research needs to be done on professional bias in philosophical judgments. Such research might examine the influence on greater domain knowledge and commitments in philosophy on the construction and evaluation of thought experiments. One clear motivation for this research is that if philosophical expertise fits the model of expertise in other fields, we should expect, at least in some cases, that biased judgments follow as a result. In the meantime, if philosophers are expert intuiters, their expertise is most likely also accompanied by susceptibility to theoretical and experimental bias as a result.

6   The Wages of Expertise

We began with a shocking example. A surprising number of American patients seek out consumer reviews of trivial consumer goods at far greater rates than they do the qualifications of their surgeons. Unsurprisingly perhaps, the American College of Surgeons does not endorse this behavior. They recommend that patients get involved. You should, they advise, become familiar with your surgeon's qualifications, history, as well as the prescribed treatment and recovery procedures. After all, experts sometimes make mistakes.

Sometimes experts make mistakes because of their expertise. As paradoxical as this can sound, it is very well documented. I have suggested that this basic observation about the boundary conditions of expertise may serve as the basis for a novel approach to understanding philosophical expertise. Of course, experts across different domains might not share all the same features. It's beyond doubt that they differ in important ways too. But research has shown that they share many important properties, including some limitations. If philosophers tend to error when evaluating thought experiments in much the same way that other experts fail across domain, this fact may reinforce the analogy between philosophers and intuition, on the one hand, and expert performances in other fields, on the other. And although there is less evidence for the factors underlying expert failure rather than mediating success, we have seen a suggestive pattern.

The basic pattern that has emerged suggests that research on the failures of professional philosophical intuition can also be explained by the failures and limitations of genuine expertise across many fields. Philosophers, like genuine experts more generally, are not perfect. They may also suffer in situations when making intuitive judgments outside the specialized range of the domain of their expertise. Evidence to date suggests that they struggle to predict the judgments and philosophically relevant behaviors of non-experts. There seem to be many areas of philosophical inquiry where glossing can, or has, occurred. Evidence of and potential for bias lurks beneath vignette construction and evaluation.

These observations help to clarify the four challenges to studying philosophical expertise, discussed earlier. First and foremost, the approach from failure effectively responds to direct challenges to the expertise defense focusing on the performances of

professional philosophers. If the data on professional philosophical intuition challenges philosophical expertise, it must show how the performances of professional philosophers are different than what we might expect of expert performances in other fields. But this body of evidence points to the opposite conclusion. Much of the data uncovered by experimental philosophy previously taken to question the link between intuition and expertise can also be interpreted as supporting it. The data support the link between intuition and expertise because they can also be explained, and in some cases even predicted, by widely accepted accounts of expertise across the social sciences.

That said, only dedicated empirical studies can demonstrate that philosophical expert intuiters share the failures and shortcomings of genuine experts. And, research on the professional intuitions of philosophers in thought experiments is ongoing. Future studies might not support philosophical expertise or the link between philosophical intuition and the failures of experts in other fields. For instance, it's possible that the results collected so far are each explained by independent factors not associated with expertise. These are exciting questions for future research guided, in part, by this new approach to philosophical expertise.

The next challenge to studying philosophical expertise was that philosophers lack adequate understanding of the factors mediating expert philosophical judgments. On this score we have also made progress—though perhaps not in the way we might have expected. We made this progress by studying the factors mediating negative expert performance across a series of domains, and using those factors to cast light on the performances of professional philosophers when they construct or evaluate thought experiments and report their intuitions about them.

Understanding the factors that mediate negative expert performances can be used to begin to reverse engineer a new model of philosophical expertise. If one assumes with Williamson and others that philosophers are experts in the traditional sense, then in line with the analogy to experts in other fields, one should expect philosophers to emit of skills and abilities that generally resemble members of that class. And as we observed, for all their skills and abilities, genuine experts also generally tend to fail in certain areas, too. Again, many of these areas in which they fail appear to overlap with errors philosophers make as they evaluate thought experiments. This reinforces the analogy between expertise in philosophy and other disciplines. It also suggests that an accurate account of philosophical expertise most likely includes these failures and limitations. In other words, while more research must be conducted, an empirically successful model of philosophical expertise will most likely include, among other things, that expert philosophical intuition is domain-specific, highly susceptible to glossing and bias, and unreliable in predicting novice behavior.

The last two challenges to studying expertise noted by Weinberg et al. involved correctness and calibration. While it can sometimes be very difficult to discern correct philosophical intuitions, it is often very apparent when intuitions go wrong. The evidence from failure circumvents the issue of correctness entirely. Our focus is on when intuitions are most likely incorrect. The tools of experimental cognitive science have readily demonstrated several areas in which expert performances go awry. By demonstrating the limitations of expertise, these tools also provide an interesting new resource for objective feedback and calibration. Awareness and adaption by professional philosophers in

response to errors, such as poor prediction, glossing or bias, may provide just the sort of philosophical calibration that is currently available to experts in other fields.

Such observations reveal important features of philosophical activity and have the potential to help us do philosophy even better in the future. The use of experimental techniques helped defend and develop a new account of expertise in philosophy based on genuine expertise seen across many fields. The use of these methods need not threaten the basic idea that philosophers are experts. By contrast, they have helped us form more realistic expectations of expert performance, both inside and outside of philosophy. Within philosophy they help calibrate intuitions in precisely those areas in which expert intuiters fail, and, potentially lead to more reliable evaluation of thought experiments in the future.

Lastly, a model of philosophical expertise based on the limitations of genuine experts may place a series of constraints on the reliability of professional philosophical intuition. In the past, many experimental philosophers have argued that if professional philosophers are not expert intuiters, then their intuitions may not be reliable evidence in philosophical arguments (e.g. Weinberg et al. 2010). I have argued that many findings may be explained in virtue of expertise. But a constrained critique of the reliability of intuition may still be warranted. Consider that if philosophers are not expert intuiters, than we have reason to question their reliability. And alternatively, if philosophers are expert intuiters who share the failures and limitations of genuine experts, then we should decrease credence in the reliability of philosophical intuition, at least in certain circumstances, too. Namely, we should decrease credence in intuitions in those situations likely to involve glossing, professional prediction of non-professional behavior, bias, and

intuiting outside of the domain of expertise. In those circumstances it seems sensible to turn to the tools of empirical cognitive science to investigate intuitions more carefully.

These results may have interesting implications for the use of the expertise defense in philosophy. For instance, the expertise defense is often invoked for the purpose of highlighting the special evidentiary weight of professional philosophical intuition. But the same research in cognitive science suggesting that philosophers may be genuine experts also seems to suggest important limits on philosophical expertise, and in turn, important limits to viable deployments of the expertise defenses in philosophy. Sometimes, we should expect expert intuition to be less reliable in virtue of the factors mediating genuine expert performance than, as some experimental philosophers have argued, in the absence of it. In such contexts, the more creditable it is that philosophers are genuine experts, the less creditable it is that a viable "expertise defense" applies. Given this, future work might profitably question applications of the expertise defense in philosophy inside and outside of these contexts.

7   Conclusion

A full account of expertise in philosophy will require much more experimental data comparing philosophers to genuine experts of other fields. The framework provided by this paper is offered to help achieve this goal. Yet it is also perhaps no coincidence that the research conducted thus far points to well-known failures and limitations of expertise. Many results taken to challenge philosophical expertise may also be explained by genuine expertise. Such results also provide a promising new approach for understanding philosophical expertise and activity. In cataloguing when professional intuitions are likely to be wrong, we may be able to increase our confidence in their expertise, in the

traditional sense Williamson and others have advocated. At the same time, the failures and limitations of genuine expertise suggest important limits on expert intuition and intuitional reliability, which experimental methods in psychology and cognitive science can help to reveal and address. Attending to these limitations and turning to those empirical methods not only helps us begin to approximate philosophical expertise and activity, but can help us do philosophy even better in the future.[16]

References

Ackerman, P. L. (1990). A correlational analysis of skill specificity: Learning, abilities, and individual differences. *JEP: Learning, Memory, and Cognition* 16, 883–901.

Adelson, B. (1984). When novices surpass experts: The difficulty of a task may increase with expertise. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10, 483–495.

Alicke, M., Rose, D. and Bloom, D. (2011). Causation, norm violation, and culpable control. *Journal of Philosophy* 108 (12):670-696.

Arico, A.J., & Fallis, D. (2013). Lies, damned lies, and statistics: An empirical investigation of the concept of lying. *Philosophical Psychology* 26 (6):790 - 816.

Austin, J.L. (1956). A plea for excuses. *Proceedings of the Aristotelian Society*, new series, 57, 1–30.

Batson C. D. (2008). Moral masquerades: Experimental exploration of the nature of moral motivation. *Phenomenology and the Cognitive Sciences*, 7, 1.

Boshuizen, H. P. A., Schmidt, H. G., & Coughlin, L. D. (1987). On-Line representation of a clinical case and the development of expertise. Paper presented at the Annual Meeting of the American Educational Research Association, Washington, D.C.

Braddock, M. (2010). Constructivist experimental philosophy on well-being and virtue. *Southern Journal of Philosophy*, 48 (3) 295-323.

Beebe, J., & Buckwalter, W. (2010). The epistemic side-effect effect. *Mind & Language*, 25(4), 474-498.

Beebe, J., & Jensen, M. (2012). Surprising connections between knowledge and action: The robustness of the epistemic side-effect effect. *Philosophical Psychology*, 25(5), 689-15.

Bengson, J. (forthcoming). How philosophers use intuition and 'intuition.' *Philosophical Studies*.

Bengson, J., Moffett, M.A., Wright, J. C. (2009). The folk on knowing how. *Philosophical Studies*, 142 (3):387–401.

Buckwalter, W. (2010). Knowledge isn't closed on saturdays. *Review of Philosophy and Psychology*, 1, 395-406.

Buckwalter, W. (2012a). Non-traditional factors in judgments about knowledge. *Philosophy Compass*, 7/4, 278–289.

Buckwalter, W. (2012b). Surveying philosophers: A response to Kuntz & Kuntz. *Review of Philosophy and Psychology*, 3, 515-524.

Buckwalter, W. (2013). Gettier made ESEE. *Philosophical Psychology*, doi:10.1080/09515089.2012.730965.

Buckwalter, W. (2014). The mystery of stakes and error in ascriber intuitions. In (ed. J. Beebe) *Advances in Experimental Epistemology*. Bloomsbury Press.

Buckwalter, W., & Phelan, M. (2013). Function and feeling machines: A defense of the philosophical conception of subjective experience. *Philosophical Studies* 166(2): 349-361.

Buckwalter, W., & Schaffer, J. (2013). Knowledge, stakes, and mistakes. *Noûs* doi: 10.1111/nous.12017.

Buckwalter, W., Rose, D., & Turri, J (2013). Belief through thick and thin. *Noûs* doi: 10.1111/nous.12048.

Chalmers, D. (forthcoming). Intuitions in philosophy: A minimal defense. *Philosophical Studies*.

Cappelen, H. (2012). *Philosophy without intuitions*. Oxford University Press.

Castel, A.D., McCabe, D.P., Roediger, H.L., Heitman, J.L. (2007). The dark side of expertise: Domain-specific memory errors. *Psychological Science*, 18 (1), 3-5 doi: 10.1111/j.1467-9280.2007.01838.x.

Chi, M. T. H. (2006). Two approaches to the study of experts' characteristics. In: Ericsson KA, Charness N, Feltovitch PJ, Hoffman RR, Ed. *The Cambridge Handbook of Expertise and Expert Performance*. Cambridge: Cambridge University Press. pp. 21–30.

Chi, M. T. H., Feltovich, P., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 5, 121–5 2.

Christensen, C., Heckerling, P. S., Mackesy, M. E., Berstein, L. M., & Elstein, A. S. (1991). Framing bias among expert and novice physicians. *Academic Medicine*, 66, S76–S78.

Church, B., Gaa, J.C., Khalid Nainar, S. M., & Shehata, M. M. (2005). Experimental evidence relating to the person-situation interactionist model of ethical decision making. *Business Ethics Quarterly*, 15 (3), 363-383.

Collins, J., Hall, N., & Paul, L. A. (2004). *Causation and counterfactuals*. Cambridge, Mass.: MIT Press.

Davis, H. L., Hoch, S. J., & Ragsdale, E. K. (1986). An anchoring and adjustment model of spousal predictions. *Journal of Consumer Research*, 13, 25-37.

Dawes, R. M. (1971). A case study of graduate admissions: Application of three principles of human decision making. *American Psychologist*, 26, 180–188.

DeRose, K. (2005). The ordinary language basis for contextualism and the new invariantism. *Philosophical Quarterly*, 55, 172-98.

Deutsch, M. (2009). Experimental philosophy and the theory of reference. *Mind & Language*, 24(4), 445-466.

Devitt, M. (2006). Intuitions in linguistics. *British Journal for the Philosophy of Science*, 57 (3), 481-513.

Ericsson, K. A., & Lehmann, A. C. (1996). Expert and exceptional performance: evidence on maximal adaptations on task constraints. *Annual Review of Psychology*, 47, 273–305.

Fantl, J., & McGrath, M. (2007). On pragmatic encroachment in epistemology. *Philosophy and Phenomenological Research*, 75, 558–89.

Fantl, J., & McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford University Press.

Faraci, D. & Shoemaker, D. (2010). Insanity, deep selves, and moral responsibility: The case of JoJo. *Review of Philosophy and Psychology* 1 (3), 319-332.

Feltz, A., Cokely, E.T. & Nadelhoffer, T. (2009). Natural compatibilism versus natural incompatibilism: Back to the drawing board. *Mind & Language* 24 (1), 1-23.

Feltz, A., & Zarpentine, C. (2010). Do you know more when it matters less? *Philosophical Psychology* 23, 5, 683–706.

Grundmann, T. (2010). Some hope for intuitions: A reply to Weinberg. *Philosophical Psychology* 23, 481-509.

Gobet, F., & Simon, H. A. (1996a). Recall of random and distorted chess positions: Implications for the theory of expertise. Memory & Cognition, 24 (4): 493-503.

Gobet, F., & Simon, H. A. (1996b). Recall of rapidly presented random chess positions is a function of skill. *Psychonomic Bulletin and Review* 3, 159–163.

Gobet, F. & Simon, H. A. (1996c). Templates in chess memory: A mechanism for recalling several boards. *Cognitive Psychology* 31, 1-40.

Hales, S. D. (2006). *Relativism and the foundations of philosophy*, Cambridge, MA: MIT Press.

Hansen, N. (2013). A slugfest of intuitions: Contextualism and experimental design. *Synthese* 190 (10):1771-1792.

Hashem, A., Chi, M. T. H., & Friedman, C. P. (2003). Medical errors as a result of specialization. *Journal of Biomedical Informatics* 36, 61–69.

Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.

Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *Journal of Philosophy* 105(10), 571.

Hinds, P. J. (1999). The curse of expertise: The effects of expertise and debiasing methods on prediction of novice performance. *Journal of Experimental Psychology: Applied*, 5 (2), 205– 221.

Hoffman, R. R. (1998). How can expertise be defined?: Implications of research from cognitive psychology. In R. Williams, W. Faulkner, & J. Fleck (Eds.), *Exploring expertise* (p. 81–100). New York: Macmillan.

Horvath, J. (2010). How (not) to react to experimental philosophy. *Philosophical Psychology*, 23, 447-480.

Howe, M.J., Davidson, J. W., & Sloboda, J, A. (1998). Innate talents: reality or myth? *Behavioral and Brain Sciences* 21(3): 399-407.

James, W. (1948). The sentiment of rationality. In A. Castell (Ed.), *Essays in pragmatism*. New York: Hafner Press.

Johnson, E. J. (1988). Expertise and decision under uncertainty: Performance and process. In M. T. H. Chi, R. Glaser, & M. J. Farr (Eds.), *The nature of expertise* (p. 209–228). Hillsdale, NJ: Erlbaum.

Knobe, J. (2003). Intentional action and side effects in ordinary language. *Analysis* 63, 190–204.

Knobe, J. (2010). Person as scientist, person as moralist. *Behavioral and Brain Sciences* 33, 315–29.

Knobe, J., Buckwalter, W., Nichols, S., Robbins, P., Sarkissian, H., & Sommers, T. (2012). Experimental philosophy. *Annual Review of Psychology* 63, 81-99.

Knobe, J. and Prinz, J. (2008). Intuitions about consciousness: Experimental studies. *Phenomenology and the Cognitive Sciences* 7, 67–83.

Kuntz, J., R., & Kuntz, J., R., C. (2011). Surveying philosophers about philosophical intuition. *Review of Philosophy and Psychology* 4, 643-665.

Lewis, D. (1999). Elusive knowledge. Reprinted in *Papers in Metaphysics and Epistemology,* Cambridge: Cambridge University Press, pp. 418-445.

Liao, M., Wiegmann, A., Alexander, J., and Vong, G. (2012). Putting the trolley in order: Experimental philosophy and the loop case. *Philosophical Psychology* 25 (5):661-671.

Livengood, J. & Machery, E. (2007). The folk probably don't think what you think they think: Experiments on causation by absence. *Midwest Studies in Philosophy* 31 (1):107–127.

Livengood, J., Sytsma, J., Feltz, A., Scheines, R., & Machery, E. (2010). Philosophical temperament. *Philosophical Psychology* 23(3), 313-330.

Locke, J. (1975). An essay concerning human understanding. (P. H. Nidditch, Ed.). Oxford: Clarendon Press.

Longfellow, H. W. (1848). *Hyperion, a romance*. W. D. Ticknor & Company: Cambridge.

Ludwig, K. (2007). The epistemology of thought experiments: First person versus third person approaches. *Midwest Studies in Philosophy* 31, 128-159.

Machery, E. (2011). Thought experiments and philosophical knowledge. *Metaphilosophy* 42, 191-214.

Machery, E. (2012). Expertise and intuitions about reference. *Theoria* 27 (1): 37-54.

May, J., Sinnott-Armstrong, W., Hull, J. G., & Zimmerman, A. (2010). Practical interests, relevant alternatives, and knowledge attributions: An empirical study. *Review of Philosophy and Psychology* 1, 265-273.

Moore, G.E. (1942). A Reply to My Critics. In *The Philosophy of G.E. Moore* (La Salle, Ill.: Open Court), ed. P. Schilpp.

Murray, D., Sytsma, J., Livengood, J. (2013). God knows (but does God believe?). *Philosophical Studies* 166 (1):83-107.

Myers-Schulz, B. & Schwitzgebel, E. (2013). Knowing that P without believing that P. *Noûs* 47 (2):371-384.

Nado, J. (2012). Why intuition? *Philosophy and Phenomenological Research*. doi: 10.1111/j.1933-1592.2012.00644.x.

Nickerson, R. S., Baddeley, A., & Freeman, B. (1987). Are people's estimates of what other people know influenced by what they themselves know? *Acta Psychologica* 64, 245-259.

Patel, V. L., Evans, D. A., & Groen, G. J. (1989). Biomedical knowledge and clinical reasoning. In Patel, V. L., and Evans, D. A. (eds.), *Cognitive Science in Medicine: Biomedical Modeling*, MIT Press, Cambridge, MA., pp. 53-112.

Phelan, M. & Buckwalter, W. (2012). Analytic functionalism and mental state attribution. *Philosophical Topics*, 40(2), 129–154.

Plato. (1967). Plato in twelve volumes, Vol. 3 translated by W.R.M. Lamb. Cambridge, MA, Harvard University Press; London, William Heinemann Ltd.

Pinillos, N, Á. (2011). Some recent work in experimental epistemology. *Philosophy Compass*, 10, 675–688.

Pinillos, N, Á. (2012). Knowledge, experiments and practical interests. in *New Essays On Knowledge Ascriptions* (Eds. Jessica Brown and Mikkel Gerken) Oxford University Press.

Reid, T. (1785/2002). *Essays on the intellectual powers of man*. (D. Brookes, Ed.) University Park: Pennsylvania State University Press.

Reuter, K. (2011). Distinguishing the appearance from the reality of pain. *Journal of Consciousness Studies* 18 (9-10), 94-109.

Richman, H.B., Staszewski, J., & Simon, H.A. (1995). Simulation of expert memory with EPAM IV. *Psychological Review* 102, 305-330.

Roediger, H.L., & McDermott, K.B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 21, 803– 814.

Rosenthal, R., & Rubin, D.B. (1978). Interpersonal expectancy effects: the first 345 studies. *The Behavioral and Brain Sciences* 3, 377–386.

Ross, W.D., (1984). *The complete works of Aristotle*. The Revised Oxford Translation, vol. 2, Jonathan Barnes, ed., Princeton: Princeton University Press.

Schaffer, J., & Knobe, K. (2012). Contrastive knowledge surveyed. *Noûs* 46 (4):675-708.

Schmidt, H. G., & Boshuizen, P. A. (1993a). On the origin of intermediate effects in clinical case recall. *Memory & Cognition* 21 (3), 338-351.

Schmidt, H. G., & Boshuizen, P. A. (1993b). On acquiring expertise in medicine. *Educational Psychology Review* 5, 205–220.

Schultz, E., Cokely, E. T., & Feltz, A. (2011). Persistent bias in expert judgments about free will and moral responsibility: A test of the Expertise Defense. *Consciousness and Cognition* 20 (4):1722-1731.

Schwitzgebel, E. & Cushman, F. Expertise in Moral Reasoning? (2012). Order effects on moral judgment in professional philosophers and non-philosophers. *Mind & Language.* 27 (2):135-153.

Smart, J.J.C. & Williams, B. (1973). *Utilitarianism: For and against.* Cambridge University Press.

Smith, S.M., Ward, T.B., Tindell, D.R., Sifonis, C.M., & Wilkenfeld, M.J. (2000). Category structure and created memories. *Memory & Cognition* 28, 386–395.

Sripada, C., & Stanley, J. (2012). Empirical tests of interest-relative invariantism. *Episteme* 9 (1):3-26.

Stanley, J. (2005). *Knowledge and practical interests*, Oxford University Press.

Starmans, C., & Friedman, O. (2012). The folk conception of knowledge. *Cognition* Online First http://dx.doi.org/10.1016/j.cognition.2012.05.017.

Strickland, B. & Suben, A. (2012) Experimenter philosophy: The problem of experimenter bias in experimental philosophy. *Review of Philosophy and Psychology.* doi:10.1007/s13164-012-0100-9.

Suzuki, S. (1970). *Zen mind, beginner's mind: Informal talks on zen meditation and practice*. Weatherhill Press.

Swain, S., Alexander, J., & Weinberg, J. (2008). The instability of philosophical intuitions: Running hot and cold on Truetemp. P*hilosophy and Phenomenological Research* 76, 138–55.

Sytsma, J. & Machery, E. (2010). Two conceptions of subjective experience. *Philosophical Studies* 151 (2):299-327.

Tobia, K., Buckwalter, W. & Stich, S. (2013). Moral intuitions: Are philosophers experts? *Philosophical Psychology* 26 (5): 629-638.

Tobia, K, Chapman, G., and Stich, S. (2013). Cleanliness is next to morality, even for philosophers. *Journal of Consciousness Studies* 20 (11-12): 195-204.

Turri, J. (2013). The test of truth: an experimental investigation of the norm of assertion. *Cognition* 129: 279–91.

Turri, J. (2013b). A conspicuous art: putting Gettier to the test. *Philosophers' Imprint*.

Turri, J., & Buckwalter, W. (2014). Descartes's schism, Locke's reunion: completing the pragmatic turn in epistemology. Manuscript under review.

Turri, J., Buckwalter, W., Blouw, P. (in press). Knowledge and Luck. *Psychon Bull Rev* DOI 10.3758/s13423-014-0683-5.

Unger, P. (1975). *Ignorance: A case for skepticism*. Oxford: Oxford University Press.

Vicente, K. J. (1992). Memory recall in a process control system: A measure of expertise and display effectiveness. *Memory and Cognition* 20, 356–373.

Vicente, K. J., & Wang, J. H. (1998). An ecological theory of expertise effects in memory recall. *Psychological Review* 105, 33–57.

Voss, J. F., Vesonder, G., & Spilich, H. (1980). Text generation and recall by high-knowledge and low-knowledge individuals. *Journal of Verbal Learning and Verbal Behavior* 19, 651– 667.

Voss J. F., Greene, T.R., Post, T.A., & Penner, B.C. (1983). Problem-solving skill in the social sciences. In *The Psychology of Learning and Motivation: Advances in Research and Theory*, ed. GH Bower, 17, 165–213. New York: Academic.

Walther, E., Fiedler, K., & Nickel, S. (2003). The influence of prior knowledge on constructive biases. *Swiss Journal of Psychology* 62(4), 219–231.

Weinberg, J., Gonnerman, C., Buckner, C., & Alexander, J. (2010). Are philosophers expert intuiters? *Philosophical Psychology* 23. 331-355.

Weinberg, J. M., Alexander, J., Gonnerman, C., & Reuter, S. (2013). Restriction & reflection: Challenge deflected, or simply redirected? *The Monist* 95 (2):200-222.

Wiley, J. (1998). Expertise as mental set: The effects of domain knowledge in creative problem solving. *Memory and Cognition* 26, 716–730.

Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.

Williamson, T. (2005). Armchair philosophy, metaphysical modality and counterfactual thinking. *Proceedings of the Aristotelian Society* 105, 1-23.

Williamson, T. (2007). *The philosophy of philosophy*. Oxford: Blackwell

Williamson, T. (2011). Philosophical expertise and the burden of proof. *Metaphilosophy* 42, 3, 215-229.