

Experimental Epistemology

James R. Beebe (University at Buffalo)

Word limit: 6,000 Current word count: 8154

Experimental epistemology is the use of the experimental methods of the cognitive sciences to shed light on debates within epistemology, the philosophical study of knowledge and rationally justified belief. Some philosophers contend that ‘experimental epistemology’ (or ‘experimental philosophy’ more generally) is an oxymoron. If you are doing experiments, they say, you are not doing philosophy. You are doing psychology or some other scientific activity. It is true that the part of experimental philosophy that is devoted to carrying out experiments and performing statistical analyses on the data collected is science rather than philosophy. But experimental philosophy is a two-part activity. Because experiments are designed to shed light on debates within philosophy, when experimental philosophers have completed their experiments, their job as experimental philosophers is not yet finished. They must take their results and inject them into an area of philosophical debate, reflecting on what their scientific investigations may have revealed. This part of experimental philosophy is indeed philosophy—not philosophy as usual perhaps, but philosophy nonetheless.

A variety of other misconceptions about experimental philosophy contribute to making it far more controversial than it should be. For example, no experimental philosopher has ever claimed that experimentation should completely replace philosophical theorizing. Yet experimental philosophers are continually faced with the following challenge from would-be critics: “If we surveyed everyone and discovered that they believe that skepticism is false (or that it’s rational to believe in God, that we have free will, etc.), how is this fact supposed to put an end to the centuries-old philosophical debate?” The simple answer is “It is not.” The empirical data gathered by experimental philosophers is supposed to *inform* rather than *replace* philosophical debate. Experimental philosophers also do not claim that their methods and results will necessarily be relevant to every area of philosophy. Again, however, it is common for critics to try to think of areas of philosophical debate where experimentation would not seem to be relevant and present them as evidence for the lack of worth of experimental philosophy. Yet consider the fact that no philosopher would dream of offering the following argument: “Insights from modal logic are not relevant to every area of philosophy; therefore, modal logic has no philosophical value and should not be practiced.” It turns out that the experiments being performed by experimental philosophers can shed light on surprisingly wide swaths of philosophical debate, but there is no claim that they must somehow be relevant to every dispute. What follows is an overview of the main areas of epistemological debate to which experimental philosophers have been contributing and the larger, philosophical challenges these contributions have raised.

I. Gettier and Truetemp

Most of the major movements and innovations of the last fifty years of contemporary epistemological debate have relied heavily upon intuitions elicited by key thought experiments. Edmund Gettier (1963), for example, appeared to successfully undermine the analysis of

knowledge as justified true belief with two simple thought experiments in which the protagonists seemed to have justified true beliefs but lacked knowledge. The externalist theories of epistemic justification that appeared in the 1970s and 1980s were attacked primarily on the grounds that they seemed to conflict with widely shared intuitions about cases such as Norman the clairvoyant (BonJour 1980), Truetemp the amazingly accurate temperature perceiver (Lehrer 1990), and victims of evil demon deception (Cohen 1984). More recently, epistemic contextualism has been both defended and attacked on the grounds that it comports well or poorly with common intuitions about key cases (DeRose 1992; 1995; 2005; Hawthorne 2004; Stanley 2005).

All of these uses of philosophical thought experiments are based on the assumption that the intuitions they elicit will be widely shared—indeed, that they *ought* to be shared by anyone who possesses the concepts of knowledge and justified belief and who has at least minimal capacities for reflection upon the correct application of those concepts. Recently, some work in experimental epistemology has put this simple assumption to the test. More precisely, experimental epistemologists have gathered data about people’s intuitive responses to philosophically influential thought experiments in a controlled fashion, and the results have often been surprising.

Gettier Cases

Having a justified true belief usually means having knowledge. However, Gettier (1963) famously introduced a class of cases in which cognitive agents have justified true beliefs that do not appear to count as knowledge. In what is usually considered to be the founding document of experimental epistemology, Jonathan Weinberg, Shaun Nichols and Stephen Stich (2001) discovered that, while most American college students of European ancestry (i.e., ‘Westerners’) gave the “correct” or typical response to Gettier cases, many American college students of East Asian (i.e., Korean, Japanese and Chinese) and South Asian (i.e., Indian, Pakistani, Bangladeshi) descent did not. Weinberg, Nichols and Stich presented participants with the following version of one of Gettier’s original cases:

Bob has a friend, Jill, who has driven a Buick for many years. Bob therefore thinks that Jill drives an American car. He is not aware, however, that her Buick has recently been stolen, and he is also not aware that Jill has replaced it with a Pontiac, which is a different kind of American car. Does Bob really know that Jill drives an American car, or does he only believe it?

Bob’s belief is justified because of his past familiarity with Jill’s driving habits, and his belief is true because Jill really does drive an American car. However, according to the overwhelming majority of epistemologists, the fact that makes Bob’s belief justified and the fact that makes it true are not related in an epistemically proper fashion. 74% of Western participants surveyed seemed to agree, as they indicated that Bob only believed but did not really know that Jill drives an American car. 53% of East Asians and 61% of South Asians, however, indicated that Bob really knows this fact (cf. Figure 1).

Gettier Case	Really Knows	Only Believes
--------------	--------------	---------------

Westerners	26%	74%
East Asians	53%	47%
South Asians	61%	39%

Figure 1

Christina Starman and Ori Friedman (unpublished) presented subjects with similar Gettier cases and found a significant gender difference among the responses. Males were more likely than females to deny that the protagonists in Gettier cases possessed knowledge. This difference did not result because men were generally reluctant to attribute knowledge or because women were generally inclined to attribute it. Starman and Friedman hypothesize that this sex difference might be due in part to the fact that women are generally more empathetic and more prone to adopt others' perspectives than men.

When intuitive responses are found to diverge in cases where it had been previously assumed that they would be unanimous, a challenge is posed to the evidential and argumentative force of these cases. If everyone who possessed the concept of knowledge agreed that the protagonists in Gettier cases lacked knowledge, the cases could be persuasively used to impugn the 'justified true belief' account of knowledge. But if there is significant disagreement, matters become more complicated. It could be that some respondents are simply confused or made some kind of performance error that prevents their responses from adequately reflecting their conceptual competence. Or it may be that some participants (e.g., from one culture) are operating with one concept of knowledge, whereas other participants (e.g., from another culture) are operating with a different one. Some have suggested that in cases of disagreement greater weight should be given to the intuitions of experts than to those of the philosophically untrained. Several experimental epistemologists have suggested that the diversity and instability of epistemic intuitions point to a more radical conclusion, viz., that intuitions should not be used as evidence in philosophical theorizing at all.

Not everyone is convinced that the results cited above pose a significant philosophical challenge. Simon Cullen (2010), for example, replicated the Gettier studies of Weinberg, Nichols and Stich but instructed participants to choose between saying that Bob *knows* or *does not know* that Jill drives an American car instead of offering them the choices *really knows* and *only believes*. Cullen correctly notes that 'really knows' seems to be distinct from 'knows' and is perhaps more akin to 'knows with certainty.' When Western participants were offered the dichotomous choice between knowing and not knowing, 42% chose "knows"—significantly more than chose "really knows" in the Weinberg, Nichols and Stich study. Cullen's replication shows that conclusions drawn about participants' concept of 'knowledge' should not be drawn from participant responses to questions about 'really knowing,' but it should be kept in mind that this does not undermine conclusions one might want to draw about participants' concept of 'really knowing.'

Jennifer Nagel (2007) cites empirical evidence from a variety of areas of cognitive scientific research that seem to point toward the hypothesis that the cognitive processes underlying epistemic evaluation should be universal and relatively independent of culture. She takes this to indicate that further empirical investigation will overturn the claim that there is philosophically significant, cross-cultural variation in epistemic intuitions concerning Gettier

cases. Nagel and her collaborators are currently carrying out experiments with the aim of showing this.

While the verdict is still out concerning the details and possible variation of Gettier case intuitions across demographic categories and the psychological processes underlying them, epistemologists should at the very least heed the following cautionary message offered by Nagel (2007, p. 802): “reactions to sketchy cases can involve a complex array of factors, and one should not be hasty to assume that one’s own initial reactions are always definitive.”

Truetemp Cases

One of the more prominent areas of debate within contemporary epistemology has been the dispute between epistemic internalism and externalism. Describing the distinction between internalism and externalism about epistemic justification, Laurence Bonjour (1992, p. 132) writes:

The most generally accepted account of this distinction is that a theory of justification is *internalist* if and only if it requires that all of the factors needed for a belief to be epistemically justified for a given person be *cognitively accessible* to that person, internal to his cognitive perspective; and *externalist*, if it allows that at least some of the justifying factors need not be thus accessible, so that they can be *external* to the believer’s cognitive perspective, beyond his ken.

The most common form of epistemic externalism is reliabilism, which claims that beliefs are justified just when they are produced by cognitive processes that are highly reliable or truth-conducive (cf. Goldman 1986). Reliabilism does not require that subjects know or be able to recognize that their cognitive processes are reliable. They must simply be reliable. This feature of reliabilism makes it a form of epistemic externalism, and it has been the target of most of the objections lodged against the theory.

Critics of reliabilism (and externalism more generally) have used thought experiments in which a hypothetical cognitive agent satisfies the reliabilist (or otherwise externalist) conditions for knowledge or justified belief, yet intuitively seems to lack knowledge or justification. One of the most widely discussed such thought experiments is Keith Lehrer’s (1990) story of Mr. Truetemp. Weinberg, Nichols and Stich (2001) employed the following version of the story in one of their experiments:

One day Charles is suddenly knocked out by a falling rock, and his brain becomes re-wired so that he is always absolutely right whenever he estimates the temperature where he is. Charles is completely unaware that his brain has been altered in this way. A few weeks later, this brain re-wiring leads him to believe that it is 71 degrees in his room. Apart from his estimation, he has no other reasons to think that it is 71 degrees. In fact, it is at that time 71 degrees in his room. Does Charles really know that it was 71 degrees in the room, or does he only believe it?

Almost all epistemologists have maintained that Truetemp-style subjects like Charles lack justification for their beliefs and, since justification is necessary for knowledge, they take him to lack knowledge as well. Among Western participants surveyed by Weinberg, Nichols and Stich, 68% of them agree. But an even greater proportion of East Asians agree (cf. Figure 2). The difference between Eastern and Western responses is statistically significant.

Individualistic Truetemp Case	Really Knows	Only Believes
Westerners	32%	68%
East Asians	12%	88%

Figure 2

A key feature of Charles' epistemic situation—and the original Truetemp story it was patterned after—is that Charles has a belief-forming process that is not shared by anyone else in his community. Drawing upon work in social psychology that suggests that people from East Asian cultures tend to be more collectivist in their thinking and less inclined toward understanding objects and individuals in detachment from their contexts than Westerners, Weinberg, Nichols and Stich constructed some other Truetemp-style cases that were less individualistic. In one version, the rock that gave Charles his new perceptual ability was replaced by a team of well-meaning scientists that are sent by the elders in his community. In another version, the entire community shares the new perceptual process in question. In both cases where some kind of community-based sanction is introduced, the statistically significant difference between Westerners and East Asians disappears. 75% of East Asians responded that the protagonist whose brain has been rewired with elder approval only believed the proposition in question, and 68% of East Asians said that the protagonist who shared his new perceptual process with others in his community only believed and did not really know (cf. Figures 3 and 4).

Elders Truetemp Case	Really Knows	Only Believes
Westerners	35%	65%
East Asians	25%	75%

Figure 3

Community Wide Truetemp Case	Really Knows	Only Believes
Westerners	20%	80%
East Asians	32%	68%

Figure 4

Swain, Alexander and Weinberg (2008) found that intuitions given in response to the basic Truetemp case are also subject to an ordering effect. If participants are first presented with a clear case of knowledge before considering the Truetemp case, they are less willing to attribute knowledge to the protagonist in the Truetemp case. But if they are first presented with a clear case of non-knowledge, they are more willing to attribute knowledge in the Truetemp Case.

Because Truetemp intuitions are thus unstable, Swain, Alexander and Weinberg suggest that they are unsuitable for use in philosophical argumentation.

Jen Wright (forthcoming), however, highlights the fact that not all of the intuitive responses gathered by Swain, Alexander and Weinberg are unstable. The clear cases of knowledge and non-knowledge were consistently recognized as knowledge and non-knowledge, respectively, regardless of the order in which they were presented. Building upon this observation, Wright replicated the findings of Swain, Alexander and Weinberg but asked participants to rate how confident they were in their judgments about the various cases. She found that participants were significantly more confident in their judgments about the clear cases than they were about Truetemp. Wright claims that her findings provide an answer to the question raised by many experimental philosophers about whether there is any way to calibrate intuitions and distinguish the reliable ones from the unreliable. Wright argues that because high degrees of confidence track stability and low degrees of confidence track instability, participants have introspective access to features that be used to distinguish those intuitions that are less subject to biasing influences from those that are not.

Cullen (2010) also replicated Swain, Alexander and Weinberg's study but instructed participants to evaluate each vignette separately. When this was done, the order effect found by Swain, Alexander and Weinberg disappeared. Cullen argues that when participants are asked a question and the point of the question is not terribly clear to them, they will look around in the context surrounding the question (where this includes other vignettes previously encountered in a questionnaire) in search of clues that can guide them in how to answer. Cullen hypothesizes that participants who are not specifically instructed to consider each vignette separately will tend to think that being shown the Truetemp case immediately after an apparently obvious case of non-knowledge is a clue that they should compare the two cases. Cullen contends that participants who compare the Truetemp case to a clear case of knowledge are in effect answering a different question from participants who are comparing the Truetemp case to a clear case of non-knowledge. When participants are given the explicit instruction to evaluate vignettes separately, however, Cullen maintains they come closer to answering the question researchers want them to answer.

In another replication of the Charles experiment above, Cullen had Western participants choose between "Charles knows" and "Charles does not know" rather than "Charles really knows" and "Charles merely believes." Cullen found that 57% of participants answered that Charles knows—almost twice as many as those who claimed that Charles *really* knows in Weinberg, Nichols and Stich's original experiment.

An often unremarked feature of the empirical findings on Truetemp cases is that before experimental epistemology came onto the scene, epistemologists unanimously agreed that the intuition that Truetemp does not know is obviously correct and is one that would be universally shared. Even Alvin Goldman (1994) and William Alston (1989)—two of the foremost defenders of reliabilism—shared this opinion and agreed that the Truetemp case presented a deep and significant challenge to their theory. However, it seems that a substantial portion of ordinary subjects view the case differently. One of the basic contributions experimental epistemology has

made has been to test a variety of empirical assumptions made by contemporary epistemologists and show how the empirical data can often surprise us.

II. Error Possibilities and Stakes

The versions of epistemic contextualism developed by Keith DeRose (1992, 1995, 2005), Stewart Cohen (1988, 1999) and David Lewis (1996) have been at the forefront of epistemological debate for the last two decades. Contextualists and their contrastivist (e.g., Schaffer 2004) or subject-sensitive invariantist critics (e.g., Hawthorne 2004, Stanley 2005) maintain that it can be true to assert ‘Bob knows that Jill drives an American car’ or ‘Mike knows that the animal is a zebra’ in some conversational contexts but false to assert them in other contexts. Although these various theorists disagree about the semantics of knowledge attributions, they generally agree that when error possibilities are made sufficiently salient in conversational contexts of certain kinds, it may no longer be true to say that someone knows, even if it would have been true before those possibilities were raised. They also claim that when the stakes are raised—i.e., when the cost of someone’s belief being wrong is high—it may be false to say that someone knows certain propositions, even though it will be true to say that someone knows those propositions in contexts where the stakes are low.

Perhaps more than any other recent position in epistemology, contextualism has made clear its aim to be grounded firmly on the epistemic intuitions of the average person. DeRose (2005, p. 172), for example, claims:

The best grounds for accepting contextualism concerning knowledge attributions come from how knowledge-attributing (and knowledge-denying) sentences are used in ordinary, non-philosophical talk: what ordinary speakers will count as ‘knowledge’ in some non-philosophical contexts they will deny is such in others.

Jason Stanley (2004, p. 11) contends that his invariantist alternative to contextualism can be shown to be superior to contextualism in part because intuitive responses to his favored thought experiments will “provide powerful intuitive evidence for [the] antecedently plausible principle concerning the relation between knowledge and action” that he wishes to defend. These claims are obviously ripe for empirical testing.

Because several leading epistemological theories predict that getting subjects to think about possibilities in which their beliefs are mistaken should make them less willing to attribute knowledge to themselves or others, Weinberg, Nichols and Stich (2001) presented participants with the following cases in which error possibilities have been raised:

It’s clear that smoking cigarettes increases the likelihood of getting cancer. However, there is now a great deal of evidence that just using nicotine by itself without smoking (for instance, by taking a nicotine pill) does not increase the likelihood of getting cancer. Jim knows about this evidence and as a result, he believes that using nicotine does not increase the likelihood of getting cancer. It is possible that the tobacco companies dishonestly made up and publicized this evidence that using nicotine does not increase

the likelihood of cancer, and that the evidence is really false and misleading. Now, the tobacco companies did not actually make up this evidence, but Jim is not aware of this fact. Does Jim really know that using nicotine doesn't increase the likelihood of getting cancer, or does he only believe it?

Mike is a young man visiting the zoo with his son, and when they come to the zebra cage, Mike points to the animal and says, "that's a zebra." Mike is right—it is a zebra. However, as the older people in his community know, there are lots of ways that people can be tricked into believing things that aren't true. Indeed, the older people in the community know that it's possible that zoo authorities could cleverly disguise mules to look just like zebras, and people viewing the animals would not be able to tell the difference. If the animal that Mike called a zebra had really been such a cleverly painted mule, Mike still would have thought that it was a zebra. Does Mike really know that the animal is a zebra, or does he only believe that it is?

In both cases, a purely hypothetical scenario involving deception is made salient to the reader, even though the reader is told that the scenario is not actual. Contrary to the widely shared assumption in epistemology that all participants should be equally disinclined to attribute knowledge in such cases, Weinberg et al. found that South Asians appear to be much less likely than their Western counterparts to deny that the protagonists have knowledge (cf. Figures 5 and 6).

Cancer Conspiracy Case	Really Knows	Only Believes
Westerners	11%	89%
South Asians	30%	70%

Figure 5

Zebra-in-Zoo Case I	Really Knows	Only Believes
Westerners	31%	69%
South Asians	50%	50%

Figure 6

Weinberg, Nichols and Stich also found differences between the responses of high and low socioeconomic status participants concerning the Cancer Conspiracy Case above and a variation of the zebra-in-the-zoo case. High socioeconomic status participants were significantly more likely than low socioeconomic status participants to deny that the cognitive agents in these cases really know the propositions in question (cf. Figures 7 and 8).

Cancer Conspiracy Case	Really Knows	Only Believes
High SES	17%	83%
Low SES	50%	50%

Figure 7

Zebra-in-Zoo Case II	Really Knows	Only Believes
-----------------------------	---------------------	----------------------

High SES	12%	88%
Low SES	33%	67%

Figure 8

Nichols, Stich & Weinberg (2003) also found that a significant majority of American college students who had taken three or more philosophy courses thought that the protagonist in a typical brain-in-a-vat case only believed and did not know that he was not a virtual-reality brain, while a narrow majority of students who had taken two or less thought that he really knew this fact.

After reviewing the above findings concerning the variability of epistemic intuitions about skeptical scenarios, Nichols, Stich and Weinberg (2003, p. 243) conclude:

Our predicament is in some ways analogous to the predicament of a person who is raised in a homogeneous and deeply religious culture and finds the truth of certain religious claims to be obvious or compelling. When such a person discovers that other people do not share his intuitions, he may well come to wonder why his intuitions are any more likely to be true than theirs.

In addition to casting doubt upon the reliability of our intuitions, Nichols, Stich and Weinberg (2003, p. 246) also think that the foregoing data should make us question the central place that debates about skepticism have occupied in western philosophy:

For if people in different cultural and SES groups and people who have had little or no philosophical training do not share ‘*our*’ intuitions (that is, the intuitions of the typical analytic philosopher who is white, western, high SES and has had *lots* of philosophical training) then they are unlikely to be as convinced or distressed as ‘*we*’ are by arguments [in support of skepticism] whose premises seem plausible only if one has the intuitions common in our very small cultural and intellectual tribe. *Pace* McGinn’s ‘anthropological conjecture,’ skepticism is neither primitive nor inevitable. And *pace* Stroud there is no reason to think that skepticism “appeals to something deep in our nature.” Rather, it seems, its appeal is very much a product of our culture, our social status and our education!

Wesley Buckwalter (2010) tested for the effects of error possibilities and high stakes by presenting participants with three versions of DeRose’s “bank” cases, one of which is the following:

Bank. Sylvie and Bruno are driving home from work on a Friday afternoon. They plan to stop at the bank to deposit their paychecks, but as they drive past the bank they notice that the lines inside are very long. Although they generally like to deposit their paychecks as soon as possible, it is not especially important in this case that they be deposited right away. Bruno tells Sylvie, “I was just here last week and I know that the bank will be open on Saturday.” Instead, Bruno suggests that they drive straight home and return to deposit their paychecks on Saturday. When they return to the bank on Saturday, it is open for business.

In the ‘High Stakes’ variant of this case, instead of being told that “it is not especially important in this case that [their paychecks] be deposited right away,” participants were told “Bruno has written a very large check, and if the money from his pay is not deposited by Monday, it will bounce, leaving Bruno in a very bad situation with his creditors.” In the ‘High Standards’ variant, participants were given the following, additional piece of information: “Sylvie says, ‘Banks are typically closed on Saturday. Maybe this bank won’t be open tomorrow either. Banks can always change their hours, I remember that this bank used to have different hours.’” Thus, the costs of being wrong are high for Bruno only in High Stakes, and an error possibility is raised only in High Standards.

Buckwalter found that while 74% of participants agreed that Bruno’s assertion “I know that the bank will be open on Saturday” was true in Bank, 69% of participants in High Stakes and 66% in High Standards also thought that Bruno’s assertion was true. Statistical analysis reveals that the mean responses in each case are significantly above the midpoint—i.e., most people agree that Bruno’s knowledge attribution is true in all three cases—but that there is no significant difference between the means of the three sets of responses. At least in the cases studied by Buckwalter, raising error possibilities and stakes had no appreciable effect on participants’ willingness to attribute knowledge.

Joshua May, Walter Sinnott-Armstrong, Jay G. Hull and Aaron Zimmerman (2010) ran a similar experiment, which included an additional case that combined both error possibilities and stakes, and came up with comparable results. May et al. found that “neither raising the possibility of error nor raising stakes moves most people from attributing knowledge to denying it.” However, even though participants generally attributed knowledge in both high and low stakes cases, they were more strongly inclined to attribute knowledge in low stakes cases. May et al. found no such effect for error possibilities.

One important difference between Buckwalter’s study and that of May et al. is that Buckwalter had participants evaluate the correctness of a knowledge attribution made by a character within the vignette, whereas May et al. simply asked participants whether they thought the character had knowledge. Because one of the most important lines of difference between contextualists and many of their critics concerns the question of whether the epistemic standards of putative knowers or the standards of attributors and deniers of knowledge should determine the truth values of knowledge attributions, the bearing that empirical results from experimental epistemology are taken to have upon extant theories of knowledge or justification can be greatly affected by which kind of probe question was asked. Subject-sensitive invariantists, for example, maintain that there should be no significant difference between the “correct” participant response to the question of whether a vignette’s protagonist knows and the “correct” response to the question of whether an attribution of knowledge to the protagonist made by another character in the vignette is correct. This is because the truth conditions for knowledge attributions, according to subject-sensitive invariantism, are not affected by how much is at stake or whether error possibilities are salient for those who are merely reading and commenting on the vignettes.

However, according to contextualism, the truth conditions for knowledge attributions are tied to the circumstances of those who are attributing or denying knowledge. If a character in a

vignette is attributing knowledge, the epistemic standards in place at that character's conversational context determine the truth value of the knowledge attribution. But if the reader of a vignette is asked to attribute or deny knowledge, it will be the standards of the reader's context that determine the attribution's truth value. Furthermore, DeRose (2010) contends that it is not as clear as many epistemologists think about what predictions extant versions of contextualism will make regarding participant responses when they are simply asked whether a vignette character knows or does not know. In fact, DeRose thinks that no version of contextualism is sufficiently developed and detailed to make any clear prediction.

In other research, Adam Feltz and Chris Zarpentine (forthcoming) ran a set of experiments that tested for the effect of high stakes upon knowledge attributions but found no effect in their studies. Ram Neta and Mark Phelan (forthcoming) ran a related set of experiments that looked at how strong participants thought the evidence of protagonists was and again found that raising stakes did not affect folk attributions, as long as the cases were presented individually. However, they did find an effect when high and low stakes cases were presented in a juxtaposed fashion. Neta and Phelan take the fact that no effect was found in individual cases to indicate that stakes do not in general factor into people's assessments of strength of evidence.

Researchers in any area of scientific inquiry need to be careful about the conclusions they draw from null results—i.e., results that fail to pass tests of statistical significance. If, for example, I look very carefully at my hands but fail to see any microbes, obviously I am not entitled to conclude that there are no microbes on my hands. Similarly, experimental epistemologists need to be careful about the conclusions they draw from studies in which no statistical differences were found between participants' responses to sets of cases. One reason why no such differences were found may be that the tools researchers used to probe for such differences were not sufficiently fine-grained or otherwise attuned to the phenomena under investigation.

In contrast to the foregoing studies that failed to find that making error possibilities salient had any effect on knowledge attributions, a study by Jonathan Schaffer and Joshua Knobe (forthcoming) did reveal such an effect when possibilities of error were presented in what they claim was “a concrete and vivid fashion.” Instead of having one character in a Bank case simply mention the abstract possibility that banks might change their hours and thus be closed on one Saturday after having been open on another, Schaffer and Knobe had one of the characters in their vignettes make the following statements:

Well, banks do change their hours sometimes. My brother Leon once got into trouble when the bank changed hours on him and closed on Saturday. How frustrating! Just imagine driving here tomorrow and finding the door locked.

Even though all participants were told that the cognitive agent whose belief was in question remained “just as confident” as he or she was that the bank will be open on Saturday, participants were less inclined to think that the character knew the bank would be open when the possibility of error was presented in this concrete fashion (mean rating: 3.05 out of 7) than when the possibility of error was presented more abstractly (mean rating: 5.54 out of 7). There is reason to suspect that the same thing can be shown to be true for the effects of raising stakes.

DeRose (2010), in fact, claims that the extra information provided about Leon above may well have the (unintended) effect of raising the stakes in the vignette. He writes:

It seems it could just as easily be thought that what [Schaffer and Knobe] are pounding home to the survey takers is thoughts about the practical consequences or stakes of being wrong. Sarah tells Hannah about poor Leon and the frustration he had to endure when he was wrong about the bank's hours, and then, turning to Hannah's and her own situation, adds, "Just imagine driving here tomorrow and finding the doors locked." Is this not at least hinting that the stakes may be quite high and/or encouraging respondents to focus on the matter of the practical consequences or stakes of being wrong?... [T]hat Sarah would carry on as she does in [Schaffer and Knobe's] beefed up HIGH case might well indicate to the survey takers that the stakes are high somehow.

If DeRose is right, then experimental epistemologists already possess data that shows that raising stakes can decrease participants' willingness to attribute knowledge. However, further empirical investigation of this issue is obviously required.

III. Knowledge and Action

If stakes were to affect knowledge attributions, it would represent one kind of connection between knowledge and action, since the practical costs of failing to know are costs associated with the actions one is undertaking. Although the experiments described above did not find such a connection, a different sort of connection between knowledge and action was found by Beebe and Buckwalter (forthcoming) and Beebe and Jensen (forthcoming). Beebe and Buckwalter initially presented participants with either the help or the harm versions of the following vignette (based upon Knobe 2003a's original study):

The vice-president of a company went to the chairman of the board and said, 'We are thinking of starting a new program. We are sure that it will help us increase profits, and it will also *help/harm* the environment.' The chairman of the board answered, 'I don't care at all about *helping/harming* the environment. I just want to make as much profit as I can. Let's start the new program.' They started the new program. Sure enough, the environment was *helped/harmed*. Did the chairman know that the new program would *help/harm* the environment?

Participants were significantly more likely to attribute knowledge to the chairman in the harm case than in the help case. In other words, the goodness or badness of the action the chairman was undertaking in light of his belief significantly affected whether participants thought that his belief counted as knowledge.

Buckwalter (forthcoming) found a significant gender difference in how participants responded to the chairman case. Women were significantly less likely than men to attribute knowledge to the chairman in the help condition, which means that the difference between helping and harming had more of an overall effect on how women responded than men. Beebe

and Jensen also found that subjects were more likely to attribute knowledge when the side-effect in question involved aesthetic or prudential (as opposed to moral) harm.

These findings suggest that the practice of making epistemic evaluations may be more closely related—at least psychologically—to other practices of normative assessment than epistemologists have thought. According to the traditional epistemological view of the relationship between knowledge and action, whether a subject knows a proposition is completely independent of whatever actions that subject may undertake in light of believing that proposition. Recent proponents of ‘pragmatic encroachment’ in epistemology (e.g. Fantl & McGrath, 2002, 2007; Hawthorne, 2004; Stanley, 2005), however, have challenged this view and argued that whether a true belief counts as knowledge can depend in part upon non-epistemic facts about how much is at stake for a subject concerning the truth of the belief. Yet it is unlikely that any of these scholars would embrace the view that the goodness or badness of actions performed in light of a belief can affect that belief’s status of knowledge. However, many defenders of pragmatic encroachment (e.g., Hawthorne and Stanley) place a premium on making their own epistemological theories square with the epistemic intuitions of ordinary people. By demonstrating another respect in which folk epistemic intuitions diverge from *a priori* expectations concerning them, the foregoing results place pressure on anyone wishing to maintain this combination of views.

IV. Larger Methodological Issues

Although the findings of experimental epistemologists obviously raise challenges to the use of this or that thought experiment for this or that purpose in epistemology, experimental philosophy is most often associated with more global methodological challenges. Consider the following, widely endorsed theses:

- (i) Whether a true belief counts as knowledge depends only upon epistemic factors such as evidence or reliability.
- (ii) Because the target of philosophical analyses of knowledge is the ordinary person’s concept of knowledge, such analyses should be answerable to data about ‘what the ordinary person would say’ in response to various epistemological thought experiments.

Many experimental epistemologists believe that the variability, instability and seeming irrationality of folk responses to thought experiments makes the conjunction of (i) and (ii) increasingly difficult to maintain. While it may be possible to dismiss a small class of the surprising patterns of variation as due to performance errors or noise, if more and more experimental data is gathered that shows that ordinary peoples’ knowledge attributions are influenced by a variety of non-epistemic factors such as culture, education, socioeconomic status, and the moral properties of actions, this line will become ever more difficult to hold.

The strongest form of the ‘experimentalist’s challenge’ to standard philosophical practice has been dubbed the ‘restrictionist view,’ according to which “the results of experimental

philosophy should figure into a radical restriction of the deployment of intuitions as evidence” (Alexander & Weinberg 2007, p. 61). Restrictionists maintain that “the problem with standard philosophical practice is that experimental evidence seems to point to the unsuitability of intuitions to serve as evidence at all” (Alexander & Weinberg 2007, p. 63). Weinberg, Nichols and Stich (2001) claim, “a sizeable group of epistemological projects—a group which includes much of what has been done in epistemology in the analytic tradition—would be seriously undermined if one or more of a cluster of empirical hypotheses about epistemic intuitions turns out to be true.” Critics of experimental philosophy have responded to these challenges in a variety of ways.

Surface Intuitions vs. Robust Intuitions

Antti Kauppinen (2007, p. 105) questions whether the intuitive responses gathered by experimental philosophical research are sufficiently robust to underwrite the experimentalist’s challenge:

There is no support to be had from responses of those non-philosophers who only appear to understand the question, who may have an imperfect grasp of the concept in question, who may or may not think hard about the application of the concept in circumstances that may or may not be conducive to avoiding conceptual mistakes, who may or may not rush in their judgements, and who may or may not be influenced by various pragmatic factors.... [T]he *actual studies* conducted so far have failed to rule out competence failures, performance failures, and the potential influence of pragmatic factors, and as such do not yield the sort of results that could support or raise doubts about philosophical appeals to conceptual intuitions.

Kauppinen is correct that experimental philosophers have sometimes been too quick to draw conclusions about the conceptual competence and narrative comprehension of their participants. For example, in the Cancer Conspiracy Case discussed above, one explanation for why the responses of low socioeconomic status did not differ significantly from chance is simply that they did not understand the task or vignette sufficiently well. Weinberg (personal communication) now cautions against relying upon this particular result precisely because the possibility of poor comprehension was not ruled out.

However, when Kauppinen goes on to express grave doubts about whether it is possible for experimental epistemologists to deal with the worries he raises, his case becomes much less convincing. Kauppinen does not think that the methods used by experimentalists allow for the possibility of (i) testing for how well participants grasp key concepts, (ii) providing sufficient motivation for participants to display a high level of performance, or (iii) testing for the effects of pragmatic factors in research materials. However, many experimental philosophers have already been testing these factors. For example, Weinberg, Nichols and Stich (2001) and Swain, Alexander and Weinberg (2008) include the following vignette to test the conceptual competence of their participants:

Dave likes to play a game with flipping a coin. He sometimes gets a “special feeling” that the next flip will come out heads. When he gets this “special feeling,” he is right about half the time, and wrong about half the time. Just before the next flip, Dave gets that “special feeling,” and the feeling leads him to believe that the coin will land heads. He flips the coin, and it does land heads.

When participants were asked whether Dave knew that the coin was going to land heads, Swain, Alexander and Weinberg excluded from further analysis any participant who answered in the affirmative, on the presumption that the participants either did not understand the question or were operating with alternative conceptions of knowledge. Comprehension checks are standard fare in social psychology and can be deployed any time concerns about participant comprehension are germane.

Regarding the effects that pragmatic factors may have on participant responses, Knobe (2003b) and Nadelhoffer (2004) have already run experiments that have tested for and apparently ruled out pragmatic effects. If researchers suspect that one kind of vignette used in experimental philosophical research has a certain pragmatic implicature and that participants may be responding to that implicature rather than to the semantic content alone (as perhaps intended by the original researchers), modified versions of the vignette that uncontroversially lack such an implicature can be used to test this suspicion. Thus, the difficulties raised by Kauppinen seem to be practical rather than principled.

The Different Concepts Response

Ernest Sosa (2007, pp. 102-103) offers the following reply to the experimentalist’s challenge:

The bearing of these surveys on traditional philosophical issues is questionable, however, because the experimental results really concern in the first instance only people’s responses to certain words. But verbal disagreement *need* not reveal any substantive, real disagreement, if ambiguity and context might account for the verbal divergence.... The experimentalists have not yet done enough to show that they have crossed the gaps created by such potential differences in meaning and context, so as to show that supposedly commonsense intuitive belief is really not as widely shared as philosophers have assumed it to be.

Sosa is certainly correct that too often experimental philosophers have tried to support far-reaching conclusions on the basis of too few studies and should do more to rule out alternative, less radical explanations of their experimental data. However, it is also important to note that pointing to the bare possibility that participants who offer different responses to survey questions may be parties to a merely verbal dispute does nothing to show that this is indeed the correct explanation of the surprising data.

Sosa (2005) also raises the following, related objection:

When we read fiction we import a great deal that is not explicit in the text. We import a lot that is normally presupposed about the physical and social structure of the situation as we follow the author's lead in our own imaginative construction.... Given that these subjects are sufficiently different culturally and socio-economically, they may because of this import different assumptions as they follow in their own imaginative construction the lead of the author of the examples, and this may result in their filling the crucial [description of a protagonist's epistemic condition] differently. But if [this description] varies across the divide, then the subjects may not after all disagree about the very same content.

Sosa's objections here are versions of the 'different concepts' response to the experimentalist's challenge. According to this response, if it can be shown that people from different demographic groups (e.g., East Asians vs. Westerners, high vs. low socioeconomic status, male vs. female) repeatedly respond to philosophical thought experiments in systematically different ways, then the two groups may be deploying nonequivalent concepts. If they are using nonequivalent concepts, the attempt by some experimental philosophers to use cross-demographic variation to challenge current ways of thinking in epistemology will come to naught, since the different groups will not even be talking about the same thing. Again, however, the 'two concepts' response merely points to a hypothetical possibility without providing any reason for believing the possibility is real. If this possibility were realized, the experimentalist's challenge would indeed be neutralized. But in the absence of reasons to think that it is, the experimentalist's challenge stands.

The Expert Response

A more common response to the experimentalist's challenge is to try to find some reason to privilege the intuitions of those who are experts concerning the application of the concepts in the relevant domain. Michael Devitt (2006, p. 103) takes up this line of response and argues that intuitions are "are empirical theory-laden central-processor responses to phenomena, differing from many other such responses only in being fairly immediate and unreflective, based on little if any conscious reasoning." He argues that we should trust a person's intuitions to the degree that we should trust the theory and experience underwriting those intuitions:

Sometimes the folk may be as expert as anyone: intuitions laden with "folk theory" are the best we have to go on. Perhaps this is the case for a range of psychological kinds. For most kinds, it clearly is not: we should trust intuitions laden with established scientific theories. Consider, for example, a paleontologist in the field searching for fossils. She sees a bit of white stone sticking through grey rock, and thinks "a pig's jawbone." This intuitive judgment is quick and unreflective. She may be quite sure but unable to explain just how she knows. We trust her judgment in a way that we would not trust folk judgments because we know that it is the result of years of study and experience of old bones; she has become a reliable indicator of the properties of fossils. Similarly we trust the intuitions of the physicist over those of the folk about many aspects of the physical world where the folk have proved notoriously unreliable. (Devitt 2006, pp. 104-105)

One can grant that Devitt's proposal sounds plausible for disciplines like paleontology and physics and yet wonder whether there is anyone who has comparable expertise in matters philosophical. The mere fact that philosophers spend more time thinking about philosophical concepts does not guarantee that time spent translates into expertise. Alexander and Weinberg (2007) note that extended reflection might simply reinforce intuitive judgments philosophers already made before engaging in reflection—i.e., that philosophical reflection might not be what produces the intuitions of philosophers at all.

Another possibility is that that long hours of participating in philosophical debate may have an effect more akin to enculturation or socialization than enlightenment. Extended practice in philosophy may simply enable one to participate successively in the culture of philosophy, where giving certain kinds of recognized responses to philosophical thought experiments is part of what makes one a full member of the culture. Experimental philosophers are not committed to the view that this is all there is to being a professional philosopher, but the experimentalist's challenge calls upon proponents of the expert response to provide non-question-begging reasons to rule out these alternatives. It is widely agreed that reasons of this sort have not been forthcoming. Moreover, Thomas Nadelhoffer and Eddy Nahmias (2007, p. 129) note that “to establish that pre-philosophical folk intuitions should *not* be trusted and that philosophically informed intuitions *should* be trusted would require more, not less, experimental research” because it is only through rigorous empirical investigation that we could establish the putative unreliability of folk intuitions.

V. Conclusion

Experimental epistemologists have begun to gather a variety of interesting data about epistemic cognition—i.e., about how ordinary people think about knowledge, evidence and related epistemic notions. Sometimes their experiments confirm established philosophical opinion, while at other times they surprise us. Contrary to what some philosophers may think, it is not the case that experimental epistemology can be deemed profitable only if it succeeds in radically overturning traditional philosophical methodology. To the degree that epistemologists engage with the ordinary person's notions of knowledge, evidence and justified belief, we need to understand what those notions are—in all of their (perhaps messy) details. Experimental epistemology is one attempt to provide us with this important kind of understanding.

References

- Alston, W. P.: 1989, 'An Internalist Externalism', in *Epistemic Justification: Essays in the Theory of Knowledge*, Cornell University Press, Ithaca, NY, pp. 227–245.
- Alexander, Joshua & Jonathan M. Weinberg, "Analytic Epistemology and Experimental Philosophy" *Philosophy Compass*, 2(2007): 56-
- Beebe, James & Buckwalter, Wesley. forthcoming. The epistemic side-effect effect," *Mind & Language*.
- Beebe, James., Jensen, Mark., manuscript. Surprising connections between knowledge and intentional action: The robustness of the epistemic side-effect effect."

- Bishop, M.A., and Trout, J.D. (2005) *Epistemology and the Psychology of Human Judgment* (Oxford and New York: Oxford University Press)
- BonJour, Laurence. 1980. "Externalist Theories of Empirical Knowledge." In P. French, T. Uehling & H. Wettstein (eds.), *Midwest Studies in Philosophy*, Vol. 5: Epistemology, University of Minnesota Press, Minneapolis, pp. 53-73.
- BonJour, Laurence. 1992. "Externalism/Internalism." In Jonathan Dancy & Ernest Sosa (eds.), *A Companion to Epistemology*, Blackwell, Oxford, pp. 132-136.
- Buckwalter, Wesley. 2010. "Knowledge Isn't Closed on Saturday: A Study in Ordinary Language." *Review of Philosophy and Psychology* 1.
- Buckwalter. forthcoming. Gender and epistemic intuition.
- Cohen, S. 1984. Justification and Truth. *Philosophical Studies* 46: 279-96.
- Cohen, Stewart 1988. "How to be a Fallibilist," *Philosophical Perspectives* 2: 91-123.
- Cohen, Stewart 1999. "Contextualism, Skepticism, and the Structure of Reasons," *Philosophical Perspectives* 13: 57-89.
- Cullen, Simon. 2010. "Survey-Driven Romanticism." *Review of Philosophy and Psychology* 1.
- DeRose, Keith. 1992. Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52, 4 pp. 913-929.
- DeRose, Keith. 1995 "Solving the Skeptical Problem." *Philosophical Review*, pp. 1-52
- DeRose, Keith. 2005. The ordinary language basis for contextualism and the new invariantism. *The Philosophical Quarterly*, 55: 172-198.
- DeRose, Keith. 2010. "Contextualism, Contrastivism, and X-Phi Surveys." Presented at the 39th Oberlin College Colloquium in Philosophy. May 9, 2010.
- Devitt, Michael. (2006). *Ignorance of Language*. Oxford: Clarendon Press.
- Fantl, J., and McGrath, M. 2002: Evidence, pragmatics, and justification. *Philosophical Review*, 111, 67-94.
- Fantl, J., and McGrath, M. 2007: On pragmatic encroachment in epistemology. *Philosophy and Phenomenological Research*, 75, 558-589.
- Feltz, Adam & Chris Zarpentine. forthcoming. "Do You Know More When It Matters Less?" *Philosophical Psychology*.
- Gettier, E. L. 1963. Is justified true belief knowledge? *Analysis*, 23, 121-123.
- Goldman, Alvin. 1986, *Epistemology and Cognition*, Harvard University Press, Cambridge, MA.
- Goldman, A. I.: 1994, 'Naturalistic Epistemology and Reliabilism', in P. French, T. Uehling, and H. Wettstein (eds.), *Midwest Studies in Philosophy*, Vol. 19. University of Minnesota Press, Minneapolis, pp. 301-320.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Kauppinen, Antti. 2007. "The Rise and Fall of Experimental Philosophy." *Philosophical Explorations* 10:95-118.
- Knobe, J. 2003a: Intentional action and side effects in ordinary language. *Analysis*, 63, 190-193.
- Knobe, J. 2003b. Intentional action in folk psychology: An experimental investigation. *Philosophical Psychology* 16 (2): 309-23
- Lehrer, K. (1990). *Theory of Knowledge*. Boulder and London: Westview Press
- Lewis, D. 1996: Elusive knowledge. *Australasian Journal of Philosophy*, 74, 549-567.
- Liao, Matthew. 2008. "A Defense of Intuitions," *Philosophical Studies* 140:247-262.
- May, Joshua, Sinnott-Armstrong, Walter, Hull, Jay G. & Zimmerman, Aaron. (2010). "Practical interests, relevant alternatives, and knowledge attributions: An empirical study." *Review of Philosophy and Psychology* 1.

- Nadelhoffer, Thomas. 2004. The Butler problem revisited. *Analysis* 64 (3): 277-84
- Nadelhoffer, Thomas & Eddy Nahmias. 2007. "The Past and Future of Experimental Philosophy." *Philosophical Explorations* 10:123-149.
- Nagel, Jennifer. 2007. "Epistemic Intuitions." *Philosophy Compass* 2:792-819.
- Neta, Ram and Mark Phelan manuscript. "Evidence that Stakes Don't Matter for Evidence."
- Nichols, Shaun, Stephen Stich & Jonathan M. Weinberg. 2003. "Metaskepticism: Meditations in Ethno-Epistemology." In Stephen Luper (ed.), *The Sceptics*. Ashgate Press, pp. 227-247.
- Schaffer, Jonathan. 2004. "From Contextualism to Contrastivism." *Philosophical Studies* 119:73-103.
- Schaffer, Jonathan., and Knobe, Joshua., manuscript. Contrastive Knowledge Surveyed. *Nous*
- Sosa, E. (2005). A defense of the use of intuitions in philosophy, in D. Murphy and M. Bishop (eds), *Stich and his critics*, Blackwell, Oxford, pp. 101-112.
- Sosa, Ernest "Experimental Philosophy and Philosophical Intuition" *Philosophical Studies* 132 (2007): 99-107.
- Stanley, Jason. 2005. *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Starmans, Christina & Ori Friedman. (unpublished manuscript). "A Sex Difference in Adults' Attributions of Knowledge."
- Swain, Stacey, Joshua Alexander & Jonathan M. Weinberg. 2008. "The Instability of Philosophical Intuitions: Running Hot and Cold on Truetemp." *Philosophy and Phenomenological Research* 76:138-155.
- Weinberg, Jonathan M., Shaun Nichols & Stephen Stich. 2001. "Normativity and Epistemic Intuitions." *Philosophical Topics* 29: 429-460.
- Wright, Jen. (forthcoming). "On intuitional stability: The clear, the strong, and the paradigmatic." *Cognition*.