Philosophical Psychology, 2014
Vol. 27, No. 3, 368–383, http://dx.doi.org/10.1080/09515089.2012.730965

Routledge
Taylor & Francis Group

Gettier made ESEE

Wesley Buckwalter

Previous research in experimental philosophy has suggested that moral judgments can influence the ordinary application of a number of different concepts, including attributions of knowledge. But should epistemologists care? The present set of studies demonstrate that this basic effect can be extended to overturn intuitions in some of the most theoretically central thought experiments in contemporary epistemology: Gettier cases. Furthermore, experiment 3 shows that this effect is unlikely to be mediated by a simple desire to blame, suggesting that a correct psychological account of ordinary knowledge attribution may include moral judgment.

Keywords: Attribution; Experimental Philosophy; Gettier; Intuitions; Knowledge; Moral Judgment; Non-Epistemic Factors; Ordinary Language; Theory of Mind

1. Introduction

Philosophers have considered a wide range of factors when theorizing about knowledge. Traditionally, epistemologists have focused their analyses on things like justification, evidence, or reliability—those factors thought to increase the likelihood that an agent's belief is true. However, a popular new trend in epistemology has encouraged a different kind of approach. Epistemologists working under this new approach have begun placing more focus on ordinary language. They argue for epistemic theories based, in part, on their ability to account for everyday uses of the word 'know'—or the purported factors that influence ordinary judgments concerning knowledge ascription.

What these theorists have suggested is that in addition to traditional things like justification or evidence, knowledge attribution may also be sensitive to some very surprising factors. For example, leading pragmatist views in epistemology claim that the truth of a knowledge ascription is sensitive to stakes, or the practical interests of an ascribee (Fantl & McGrath, 2010; Hawthorne, 2004; Stanley, 2005). On the other hand, many contextualists claim that the truth of a knowledge ascription depends on

Wesley Buckwalter is a PhD candidate at the City University of New York.
Correspondence to: Wesley Buckwalter, Philosophy Program, The Graduate Center, CUNY, 365 Fifth Ave, Rm 7113, New York, NY 10016, USA. Email: jbuckwalter@gc.cuny.edu

the degree to which error possibilities have or have not been made salient to an attributor (Cohen, 2004; DeRose, 2009).

Since these claims are largely descriptive in nature, this new trend emphasizing ordinary ascription has also inspired new collaborations between epistemologists and experimental philosophers (DeRose, 2011; Nagel, forthcoming; Schaffer & Knobe, 2010). Their efforts have mostly focused on testing the predictions of pragmatist and contextualist theories in epistemology. What researchers have discovered is that when subjected to empirical scrutiny in properly controlled experiments, there is some evidence to support the claims epistemologists have made regarding practical interests (Pinillos, 2012; Sripada & Stanley, 2012) and error salience (Buckwalter, forthcoming; Schaffer & Knobe, 2010).[1]

But importantly, experimental philosophers have also discovered that these are not the only surprising factors at work in ordinary attributions.[2] Following the work of Knobe (2003) on intentionality judgments, it has also been suggested that there is a *moral component* to knowledge attribution (Beebe & Buckwalter, 2010). This effect, known as the *Epistemic Side-Effect Effect* (hereafter ESEE), demonstrates that people's prior evaluative judgments—and in particular, their moral judgments—can influence judgments they make about knowledge attribution. Though there has been significant debate among cognitive scientists regarding the best explanation of this effect, what researchers have found time and again is that morality seems to have a robust impact on what we take others to know (Beebe & Jensen, 2012).

In comparison to the attention given to empirical data in experimental epistemology regarding stakes or error salience, though, the impact of morality on epistemic judgments has gone largely overlooked. Should epistemologists also care about the influence that morality seems to have on knowledge attribution practices? This paper presents new experimental evidence suggesting that the moral component of knowledge attribution may have a series of important implications for epistemologists sympathetic to the ordinary language approach. These data show that the general effect shown in ESEE can be extended, to overturn standard intuitions in some of the most theoretically central thought experiments used in contemporary epistemology: Gettier cases.

In what follows, section 2 reviews the extant literature in experimental philosophy regarding the moral component of knowledge attribution, as well as the disagreement surrounding rival psychological explanations for this effect. Section 3 is a short discussion of the role of Gettier intuitions in epistemology. Section 4 presents three studies demonstrating that when typical Gettier cases are *made ESEE* (in other words, when Gettier cases include evaluative features such as those found in epistemic side-effect cases), moral valence significantly influences participants to attribute knowledge to Gettier subjects. In section 5, it is argued that while this research is currently ongoing, the correctness of each rival psychological explanation reviewed in section 2 will bear heavily on the typical use of Gettier intuitions in epistemology, and the recent practice of using data concerning ordinary knowledge attribution as evidence for theories of knowledge.

2. Prior Work on ESEE Judgments

Perhaps one of the most famous discoveries in experimental philosophy to date is that moral judgments can have a pervasive impact on the application of non-moral concepts (Knobe et al., 2012). For instance, in Knobe's (2003) well-known study on intentionality judgments, he constructed two vignettes that described the circumstances of a company chairman learning of a potential new program to increase profits. However in each case, adopting this new program has a side effect—it will either help or harm the environment. Since the chairman is indifferent toward these side effects, he decides to start the new program in both cases. The result is that the environment ends up being helped or harmed, respectively. What Knobe found is that people's moral judgments in these cases had a large impact on their intentionality judgments. Specifically, he found that 82% of participants in the harm condition agreed that the chairman intentionally harmed the environment, whereas 77% of subjects in the help condition said that the chairman did not intentionally help the environment. This effect is known as the side-effect effect.

What Beebe and Buckwalter (2010) have demonstrated in ESEE is that moral considerations like those prominent in the side-effect effect can also affect people's intuitions about the chairman's knowledge. In their study, they gave undergraduate students ($N = 749$) the same cases Knobe used above, but instead asked participants to ascribe knowledge to the chairman regarding the occurrence of either the good or bad side effect. And surprisingly, they found the same basic result. Participants were more likely to say that the chairman in the vignettes knew that his actions would bring about the side effect when the outcome was bad, and less likely to attribute knowledge when the side effect was good. This effect has been widely replicated, suggesting that just like the intentional side effect, people's prior moral judgments can also play an important role in their knowledge attributions (Beebe & Jensen, 2012).

The general phenomenon regarding the influence of morality on intentionality or knowledge judgments is now reasonably well understood. However, the psychological explanation for why these effects occur has become the subject of continuing and ever-broadening controversy (Knobe, 2010). Rival explanations are often divided into two broad families of views. One family of views holds that moral considerations *distort* the ordinary application of non-moral concepts like knowledge. Though these views have been developed in a number of different ways, they all posit some sort of additional cognitive process triggered by the presence of moral considerations that bias people's judgments. For instance, Nadelhoffer (2006) argues that moral considerations give rise to certain affective or emotional responses, and these emotional responses in turn bias the judgments participants make. Alternatively, Adams and Steadman (2004) suggest that this effect is due entirely to conversational pragmatics. Still others argue that people only apply the concepts in question because they are looking for a way to justify their desire to blame, or hold agents morally responsible for their actions (Alicke, 2008; Malle & Nelson, 2003).[3]

The other family of views holds that morality is an important part of people's core conceptual competences. As opposed to distortion accounts, *competence explanations* claim that there is no additional cognitive process distorting ordinary judgments. Under this type of explanation, the influence of morality is said to reflect the genuine use of the concepts in question (Knobe, 2010). For example, one prevalent theory along these lines by Halpern and Hitchcock (2010) suggests that moral judgments play an important role in counterfactual reasoning, and that these counterfactuals, in turn, underlie the correct application of the concepts in question. Similarly, Schaffer and Knobe (2010) argue that moral considerations fix the class of alternative outcomes people consider relevant, and these contrast classes affect the truth of knowledge ascriptions. Generally, though, the basic idea is that moral considerations are an important part of the genuine use of concepts like intentionality or knowledge.

Yet, despite the proliferation of data for and against theories of both sorts, no single theoretical proposal in either family has emerged the clear victor (Knobe et al., 2012). It is fair to say that this debate has reached something of an impasse. The impasse has led some researchers to reconsider whether experimental philosophers actually possess the necessary tools for classifying psychological effects in terms of competence or performance at all (Alexander, Mallon, & Weinberg, 2010a, 2010b). Some have posited that there are stable individual differences responsible for the effects in question (Nichols & Ulatowski, 2007; Pinillos, Smith, Nair, Marchetto, & Mun, 2011). Still others argue that there is unlikely to exist any one unified explanation for these phenomena, and that many overlapping factors are involved in the impact moral judgment has on the ordinary application of concepts like intentionality or knowledge (Phelan, 2011).

Research into the best explanation for the role of morality in knowledge attribution is ongoing, and debate is sure to persist among philosophers and cognitive scientists. Presently, our interest is in the lessons that epistemologists can learn from the moral component of knowledge attribution in the meantime. It is argued that at the current state of research, the general correctness of *either* competence- or performance-style explanations may have important implications for epistemologists interested in the conditions under which knowledge is ordinarily ascribed. To help frame these implications, and to further explore the role of morality in people's epistemic judgments, we revisit the basic effect discovered in ESEE. What we find is that ESEE case results can be expanded to significantly influence central thought experiments of the philosophical tradition: Gettier cases.

3. Gettier Cases in Epistemology

In contemporary epistemology, many refer to Gettier cases as standard counterexamples to the previously received justified-true-belief (JTB) account of knowledge (Gettier, 1963). More specifically, Gettier cases are thought experiments aimed at eliciting the following intuition. If an epistemic agent in a Gettier case does not know a true proposition (that he or she believes and is justified to believe), then this is

evidence against the view that knowledge is justified true belief. Consider the following example:

> Sam hears on the news that a beneficial new chemical has found its way into the town reservoir. As the water pump operator, Sam continues his duties believing that pumping the water will help the local crops. Sure enough, the crops thrive. So Sam's belief is true. However, now imagine everyone was wrong about the beneficial chemical, and instead a mysterious fungus was responsible for the positive results. *Did Sam know the crops would thrive?*

In the thought experiment above, Sam has a certain belief that his actions will affect the crops, and that belief is supported by relatively good evidence. Sam's belief is justified, but his belief is only true by luck. The target intuition of the Gettier case is that no matter the evidence or justification of the putative knower, Sam did not know that by pumping the water the crops would thrive. Typically, this intuition is taken to be unanimous by contemporary epistemologists.[4] The content of the intuition is then taken as obviously true, and imported as strong evidence against the position that knowledge consists of justification, truth, and belief alone.

But while a significant amount of theoretical work has been done regarding the best philosophical response to Gettier problems, very little has been said about the psychological origins of the intuition itself (Lycan, 2006).[5] More often, debates in epistemology regarding Gettier proceed from the assumption that the intuition exists, to the conclusion that JTB analyses of knowledge must be false. Presumably though, an investigation of the mechanisms underlying these kinds of intuitions may help us gain insight into the Gettier intuition's evidentiary status, and in turn, deeper philosophical significance. So, perhaps the study of Gettier cases can also benefit from the experimental techniques being used by experimental epistemologists to discover the factors involved in knowledge attribution.

Returning now to our previous discussion concerning the moral component of knowledge attribution, might normative evaluations be playing a role in typical Gettier case intuitions? In a recent article, Turri (2012) considers that precise possibility. Citing the basic psychological effect found in ESEE, Turri suggests that one might be more willing to attribute knowledge that *p* in Gettier conditions if the moral valence of *p* is manipulated. Specifically, the suggestion is that people's prior moral judgments may lead to attributions of knowledge when Gettier agents are judged to be morally bad. This, of course, would be very shocking, since it would mean that the moral component of knowledge attribution completely overturns the widely accepted verdict in epistemology that Gettier agents do not know that *p*. To put Turri's suggestion to the test, the following three experiments were conducted to investigate this question: can the epistemic side-effect effect be Gettier-ized?

4. GESEE Experiments

The following three experiments are designed to provide initial evidence that the moral valence of a particular Gettier situation significantly influences ordinary knowledge attribution to Gettier subjects.

4.1. Experiment 1 (Pump Case)

To investigate the relationship between the moral component of knowledge attribution and Gettier intuitions let us begin by revisiting the above pump example. In a between-subjects experimental design, participants in experiment 1 ($N = 86$, mean age = 36, 52 female, 87% Caucasian) were provided with one of two vignettes closely resembling cases originally used by Ulatowski (2012) in the study of action individualization.[6] The two resulting cases are shown below:

> Sam's job is to pump water into the cistern, which then supplies the water to the farms owned by several families in the community. One day, as Sam operates the pump, he hears a broadcast on the radio. The radio report says that local officials suspect a new chemical from a nearby factory, chemical X, may have found its way into the local reservoir, and that there is a chance it will be very [*beneficial/poisonous*] to all the local townspeople's crops. Sam thinks to himself, "I don't care about their crops; I just want to earn my pay," and continues pumping the water. Sure enough, the crops started [*thriving/dying*]. It turned out that the local officials were completely wrong about the chemical in the water. After analyzing the water, they found no trace of chemical X. Scientific reports later confirmed that the crops were all [*thriving/dying*] because of a fungus that had been secretly growing inside Sam's pump.

These stimuli describe the conditions of an agent who is indifferent towards bringing about either a good or a bad result. Typical of standard Gettier cases, the agent believes that the result will transpire, and that belief is supported by good evidence. While the belief is in fact true, the way the agent formed and justified the belief is only connected to a state of affairs in the world that makes the belief true by luck. Thus, the standard prediction among epistemologists is that no knowledge will be ascribed in either case.

After seeing either the good result (the crops thrive) or bad result (the crops die) conditions of the vignette above, participants were asked to make judgments about the Gettier agent's knowledge. Specifically, participants were asked whether they agreed or disagreed with the statement, "Sam knew that by pumping the water, the townspeople's crops would [*thrive/die*]." Responses were collected on a seven-item scale anchored by positive and negative agreement terms (1 = strongly disagree, 4 = neither agree nor disagree, 7 = strongly agree). For the hypothesis that knowledge attribution in Gettier conditions is sensitive to moral factors to be confirmed, we would expect people to make asymmetrical judgments depending on case valence.

This is exactly what was found. When pumping the water produced a good result (Help, $M = 3.05$, $SD = 1.59$), participants tended to give the standard philosophical answer that the agent did not have knowledge. However, participants said that Sam did know that his actions would bring about the particular result when that result happened to be bad (Harm, $M = 4.86$, $SD = 1.7$). In order to test the hypothesis that moral judgment affects knowledge attribution in Gettier cases, an independent samples t-test was conducted. As predicted, there is a significant difference between knowledge judgments made between these two groups, $t(84) = 5.04$, $p < 0.01$, $d = 1.10$. While in both Gettier cases, the agent's epistemic position is the same

(in that the relevant justified belief is only true by luck), the valence of the relevant action has an important impact on what Gettier agents are said to know. Such findings begin to suggest that moral judgments have a strong impact on people's intuitions about attributing knowledge in Gettier cases.

4.2. Experiment 2 (Mayor Case)

Of course, we would not want to draw conclusions about the general effect morality might have on Gettier judgments based only on a single experiment. After all, there may be a number of concerns about the specific stimuli used in experiment 1. For instance, environmental harm is a very sensitive issue. So, perhaps participants' personal views concerning the environment are somehow impacting their answers in this particular case, and do not reflect their actual Gettier judgments when moral factors become salient more generally.

So, experiment 2 sought to replicate the same asymmetry found in the first experiment using a completely different case that did not involve mention of the environment. Participants ($N = 78$, mean age = 36, 46 women, 76% Caucasian) were presented with vignettes about a mayor of a small town who has formed the belief that taking a certain action will either create or cut jobs for members of the local community:

> The mayor of a small town is trying to decide whether or not to sign a new contract with a local corporation. The math is all very complex, but all his economic strategists think that there's a relatively good chance that one outcome is that it will [*create/cut*] jobs for workers in the community. The mayor says, "all I really care about is campaign contributions, not people's jobs, and I am sure to get millions from the corporation if I agree." So, he decides to sign their contract. The corporation, however, didn't take any chances. They secretly switched the contract with a totally different one right before the mayor signed it. By changing all the fine print, in some cases the opposite of what the mayor thought he was signing, the corporation could be sure it got what it wanted. Sure enough, shortly after the mayor signed the contract, a number of members of the community [*got/lost*] jobs, and the mayor received a huge donation to his reelection campaign.

Just like the pump case, the mayor has a strong justification for true beliefs. Then, Gettier information is added to the story so that the result that actually transpires has nothing to do with the mayor's evidence. As in experiment 1, participants were given either the helping or harming version of this scenario and then asked if they agreed or disagreed that "the mayor knew that by signing the contract he would [create/cut jobs]." Responses were again collected on the same seven-item scale anchored by positive and negative agreement terms.

Though the means were slightly higher in experiment 2, this result replicated the same basic finding as the pump cases.[7] Participants were more likely to say that the epistemic subject knew that the certain result would take place when that result was bad (Harm, $M = 6.05$, $SD = 0.94$) than when it was good (Help, $M = 4.11$, $SD = 1.86$). An independent samples t-test again shows this difference to be significant, $t(76) = 5.92$, $p < 0.01$, $d = 1.32$. These results suggest that the effect of

moral valence on participants' epistemic judgments is not restricted to cases about the environment. The same pattern of responses appears when participants are presented with the mayor case.

4.3. Experiment 3 (Third Person Mayor Case)

So far, we have seen that the moral component of knowledge attribution plays a role in shaping intuitions about Gettier cases. But one immediate objection to these findings is that perhaps people only attribute knowledge in the pump and mayor vignettes because they are looking for a way to hold someone responsible for the outcome of these cases, and not because they believe those agents really have knowledge when the results are bad.

This idea is perhaps best illustrated with an old example from Kripke.[8] Here, Kripke asks us to imagine that we have just heard the news about President Nixon's possible involvement in Watergate, and everyone is trying to figure out what Nixon knew about the cover-up. A trained epistemologist might tell us that technically Nixon did not have knowledge of the crimes because he *merely had a true belief* about the wrongdoings. Nevertheless, few consider this response satisfactory. Instead, we find ourselves wanting to attribute the relevant knowledge to Nixon. But our knowledge attribution to Nixon in this case would be driven—perhaps not by our properly attending to the justification Nixon had for his beliefs—but rather, only as a means to justify our desire to hold him responsible for the Watergate scandal.

So the objection—what we might call the *Kripkean objection* to experiments 1 and 2—is that the data concerning the moral component only appear to reveal participants' real views about knowledge attribution. But instead, people are saying that the agents had knowledge only as a means to justify their desire to hold those agents morally responsible for the bad results. In fact, as we saw in section 2, researchers have argued that the effect that morality has been shown to have on the application of different non-moral concepts is not primarily due to people's underlying understanding of those concepts, but rather due to this kind of distortion regarding the desire to blame (Alicke, 2008).

While such explanations remain controversial, it might be thought that something similar could provide a straightforward explanation of the present data concerning knowledge ascription in Gettier cases seen in experiments 1 and 2. Thus, experiment 3 was designed to put this idea to the test. What we would like to know is whether or not people are attributing knowledge only as a means to hold agents responsible, and not because their underlying concept of knowledge is actually sensitive to the relevant moral evaluation.

To test whether people are genuinely attributing knowledge to Gettier agents independently of their desire to justify their blame judgments, participants in experiment 3 ($N = 85$, mean age = 34, 50 women, 81% Caucasian) were presented with vignettes that very closely resembled the mayor cases from experiment 2. The only difference was that they both included an additional character such that the epistemic agent in the story was not the same agent whose actions brought about the

good or bad result. This was accomplished by adding the following two sentences directly after those describing the mayor's indifference towards the community, but right before the Gettier information was presented in the previous versions:

> James the office secretary overheard everything, and is appalled by what the mayor said. Nonetheless, the mayor decides to sign the contract.

Given these third person mayor cases, participants were then asked, on the same seven-item scale as before, whether they agreed or disagreed that "James the office secretary knew that members of the local community would [get/lose] jobs."[9]

As predicted, when asked about James' knowledge, the same basic asymmetric pattern of knowledge attribution from experiment 1 and experiment 2 was found. Participants exhibited higher levels of agreement regarding James the office secretary's knowledge when the result of the case was bad (Harm, $M = 4.98$, $SD = 1.72$) than when it was good (Help, $M = 3.95$, $SD = 1.48$). But importantly, the mayor, and not James the office secretary, is to blame for causing the bad result.[10] We may even find ourselves empathizing with James as a victim of the mayor's appalling actions. Nonetheless, the effect of moral judgment on knowledge attribution in Gettier cases appears to persist even when the agent in question is not straightforwardly blameworthy. An independent t-test reveals a significant difference between these two groups, $t(83) = 2.94$, $p < 0.01$, $d = 0.64$.

Thus, it is unlikely that the effect reported across the three experiments above could be entirely mediated by the Kripke objection. Third person mayor cases display a similar moral effect on knowledge attribution, even when participants are asked about the knowledge of a completely different agent than the one who brings about the bad results. But if there is little reason to hold James responsible for the layoffs, then it is doubtful whether or not a sole desire to justify blame judgments can entirely explain this effect. Of course, experiment 3 does not conclusively prove that this

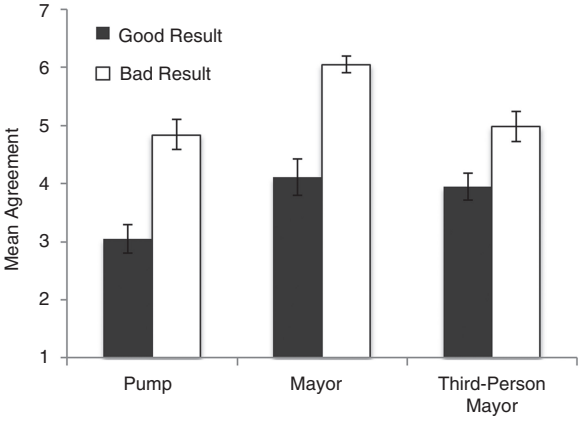

Figure 1 Participants' Mean Ratings of Agreement with Knowledge Attribution to Gettier Agents in Experiments 1–3 (±SE) by Outcome Valence. All scales ran 1–7.

distortion in epistemic concepts plays absolutely no role in people's judgments.[11] Nor does it rule out the possibility that some other kind of distortion theory can account for the present set of data. However, these results do suggest that further experimental evidence is required before dismissing the knowledge ascription found in Gettier cases out of hand. These results across experiments 1–3 are represented in Figure 1.

5. Discussion

Across three different types of vignettes, we find this same basic effect. When Gettier cases are made ESEE, morality plays an important role in knowledge attribution to Gettier agents. The effect was found in cases that involved different contexts (the economy and the environment), different kinds of evidence (a radio report, a subordinate's testimony, and eavesdropping), and even persisted where epistemic evaluations were made about completely different agents (responsible or not responsible for bringing about the outcomes in question).

Expanding ESEE to GESEE (or the *Gettier-ized* epistemic side-effect effect) also demonstrates something important about the original effect. Again, consider the chairman case. One deflationary hypothesis is that the reason why people attribute knowledge in the harm case but not the help case is because of the chairman's evidence. A general assumption often made about corporations is that they are more likely to cause environmental harm than good. So the same evidence that makes true beliefs count as knowledge in the harm condition may be insufficient in the help condition. Perhaps, then, ESEE judgments occur not because knowledge attribution has a moral component, but simply because participants require agents to possess a greater amount of justification in the good cases in order to know. Yet, the presence of the same effect in Gettier cases seems to rule out this hypothesis. It's true that in GESEE, Gettier agents in help and harm conditions across experiments 1–3 have the same evidence. But in each case, this evidence is completely disconnected from the state of affairs in the world that makes their beliefs true. It is unlikely, then, that the effect in question arises because participants think that knowing requires different amounts of justification across the various conditions. After all, according to the standard philosophical intuition, the very structure of a Gettier case is supposed to show that *no amount of justification* will lead a Getter agent to possess knowledge.

So, expanding ESEE to Gettier case intuitions in the experiments above, and perhaps even in professional philosophers' own reflections on these cases, gives us even more evidence for the moral component of knowledge attribution.[12] The deeper question of explaining why moral evaluations are having this particular effect on knowledge attribution now remains. As mentioned in section 2, there is extensive debate among researchers as to whether competence- or performance-style explanations provide the best interpretations of the current findings. These are very difficult questions. As this research continues, no assumption is made presently concerning either the correctness of implementing competence or performance

distinctions in experimental philosophy, or classifying the data in hand according to one of these two frameworks. However, supposing that these explanations might be extended to also account for knowledge judgments in experiments 1–3, let's examine how the correctness of either approach might reveal important epistemological implications when it comes to Gettier intuitions.

Beginning with distortion-style explanations, suppose that people are attributing knowledge in Gettier cases because moral valence is somehow biasing or distorting the way they ordinarily make knowledge judgments.[13] Under this kind of explanation, the attributions in the cases above can be seen as precisely the kind of performance error that properly controlled experimental methods should weed out. The fact that people attribute knowledge in some Gettier cases but not others leads us to simply discard intuitions about knowledge in the vignettes with "bad" results. Then, even if these cases do tap into a heuristic we use concerning conditions of harm, one can still reject the content of the resulting intuition as informing philosophical theories about what one actually knows.

But such a finding also provides evidence that Gettier intuitions are *case-dependent*. And this result is not without precedent in epistemology. For instance, Gendler and Hawthorne (2005) argue for a similar conclusion regarding the purportedly standard philosophical intuitions in fake barn cases (Goldman, 1976). By appraising a series of different fanciful barn case variations, their suggestion is that "the concept of knowledge, prior to its being fashioned and molded by certain philosophical traditions, never offered any stable negative verdict in the original fake barn case" (Goldman, 1976).[14] Likewise, it looks as though GESEE provides the beginnings of some evidence that intuitions in Gettier cases may also vary from case to case. This time, however, it is because of the distorting effect that *moral considerations* have on epistemic judgments.

But suppose that the experiments above really do point to a factor that systematically biases people's judgments in Gettier cases. If one thinks that GESEE results are due to some kind of distortion or error in performance, one ought to decrease credence in the default position that intuitions in Gettier cases really count as good evidence for or against traditional theories of knowledge. Not only are widely accepted Gettier intuitions easily overturned when moral considerations are added to the vignette, but this effect prompts the following question: just what other sorts of tacit factors *are* responsible for our judgments about knowledge ascription in these theoretically crucial cases? Without further research into this question, as well as the other potential biasing effects that may be present in Gettier judgments, these data begin to raise questions about the practice of using Gettier intuitions as decisive evidence for or against theories of knowledge.[15]

Alternatively, suppose that a conceptual competence-style explanation is correct. In this case, the presence of moral valence would not give rise to any additional cognitive process that biases people's judgments. Instead, knowledge attribution would be said to genuinely fluctuate in the Gettier cases above due to the underling concept of knowledge that people tacitly hold. Furthermore, suppose again that one endorses the popular position in epistemology that judgments of ordinary knowledge

ascription should count as good evidence for a correct theory of knowledge. Data strongly suggest that the *salience of moral factors* can be crucial to judgments people make about what Gettier agents know. Pace the standard philosophical intuition, people do think that Gettier agents have knowledge under these conditions.

These results appear to indicate that, just like practical interests or the salience of error possibilities heavily discussed in connection with pragmatist or contextualist theories in epistemology, morality is an equally important factor that influences ordinary ascription. Thus, if one thinks that theories of knowledge should include a correct account of the underlying ordinary concept, then one should adapt extant theories to do moral judgment justice. In other words, if a conceptual competence explanation turns out to be correct, then epistemologists attempting to capture the way knowledge is ordinarily ascribed can no longer afford to neglect the moral component of knowledge attribution in their epistemic theorizing.

As we have seen, the correctness of a competence- or performance-style explanation leads to different recommendations regarding the moral component of knowledge in epistemic theory building. But now, continue to suppose that ordinary intuitions about 'knows' really do count as important sources of evidence about *knowledge*. Notice that by accepting this principle, either explanation yields a similar result about the typical use of the Gettier case intuition. Namely, the data suggest that one should be *much less confident* about what Gettier agents do and do not actually know. If it's distortion, one should be less confident simply because the evidence demonstrates the ease to which Gettier judgments can be manipulated. If it's competence, one might argue that if Gettier agents in harm conditions have knowledge, and Gettier agents in help conditions have the exact same evidence, then Gettier agents in help conditions must also have knowledge.[16] Of course, neither of these arguments conclusively shows that Gettier agents have knowledge. However, they should decrease credence in the widely accepted view that Gettier intuitions provide insurmountable evidence against JTB analyses of knowledge.

Lastly, the current findings may point to a potential avenue for future research in experimental philosophy regarding the competence and performance distinction. It has been argued that GESEE results have important implications for epistemology. But presumably, the correct psychological explanation for the moral component of knowledge attribution is likely to also apply to the moral effects cognitive scientists have found for other non-moral concepts like intentionality. Yet, unlike intentionality, attributions of knowledge are not limited to first person cases. Questions regarding what an agent intended always involve judgments about the agents who are in some way causally responsible for an action (e.g., the agent *who intentionally harms* the environment or *unintentionally creates* jobs). As we saw in experiment 3, knowledge is different. Questions about what an agent *knows* can be asked independently of causal responsibility. Give that GESEE has been shown to persist in these kinds of third person cases, pursuing experiments regarding the influence of morality in the epistemic domain may provide a promising new way to test competing explanations for these effects.

6. Conclusion

We began with a discussion of a popular new approach in epistemology that focuses on ordinary language practices. We saw that experimental philosophers have discovered that ordinary knowledge attribution is sensitive to some surprising factors. These factors include the influence of people's prior moral evaluations on their epistemic judgments. But current research has also shown that the moral component of knowledge attribution does not only arise in the isolated cases of experimental philosophy, but are also extended to influence standard philosophical intuitions in Gettier cases. And, Gettier case intuitions have been instrumental in philosophical argumentation. We have also seen that there has been debate about the best interpretation of the kind of effect shown in GESEE. This debate questions whether distortion- or competence-based accounts better explain morality's influence on knowledge attributions. It has been argued that this debate is extremely relevant to epistemologists interested in providing epistemic theories that are in accord with the way people ordinarily attribute knowledge.

Epistemologists are faced with two equally important outcomes. If, as Gettier case data suggest, moral judgment does play an important role in people's underlying conceptual competences of knowledge attribution, then this feature of folk judgments should be accounted for by epistemic views claiming to capture the way people ordinarily attribute knowledge. Conversely, if moral judgment is shown to systematically distort ordinary knowledge judgments in Gettier cases, then perhaps the intuition that Gettier agents do not have knowledge does not provide as much evidential support for a particular analysis of knowledge as epistemologists have assumed. Either way, one should generally be much less confident when invoking Gettier intuitions as evidence. Lastly, the possibility of third person testing in the epistemic domain may provide a novel approach for future experiments regarding competence and performance explanations.

Notes

[1] For a debate about the best way to interpret these data regarding stakes see Buckwalter (2010, forthcoming); Buckwalter and Schaffer (unpublished manuscript); Feltz and Zarpentine (2010); and May, Sinnott-Armstrong, Hull, and Zimmerman (2010).

[2] Also see effects for culture (Weinberg, Nichols, & Stich, 2001), gender (Buckwalter & Stich, forthcoming), native language (Vaesen & Peterson, unpublished manuscript), and order bias (Swain, Alexander, & Weinberg, 2008). For a review of the latest research in experimental epistemology, see Buckwalter (2012).

[3] For evidence against these three accounts, see Young, Cushman, Adolphs, Tranel, and Hauser (2006); Knobe (2004); and Guglielmo and Malle (2009), respectively.

[4] Though, see evidence that this intuition may be culturally local (Buckwalter & Stich, 2011; Weinberg et al., 2001).

[5] In fact, see Zagzebski (1994) for the ''inescapable nature'' of Gettier problems.

[6] All studies reported were run online using Amazon Mturk and Qualtrics. Because Weinberg et al. (2001) found cross-cultural differences in epistemic judgments in prior work, the country of origin for participants was restricted to the United States.

[7] As opposed to experiment 1, mean judgment in the help cases of experiment 2 were not significantly above or below the midpoint, suggesting that participants were largely unsure whether or not the Gettier agent knew.

[8] I am grateful to Keith DeRose and Michael McGlone for this reference to an unpublished 1985 lecture given by Saul Kripke at Princeton University (see DeRose, 2002).

[9] We might note one important epistemic difference between this case and experiment 2: perhaps James has less evidence than the mayor for either the good or the bad result.

[10] These findings do not rule out blame entirely. Nonetheless, this case still replicates the target effect when, presumably, the desire to blame is greatly minimized. These results are also consistent with previous research by Machery (2008) and Uttich and Lombrozo (2010) involving similar cases in which intentionality is ascribed to agents whose actions are not explicitly blameworthy.

[11] In fact, the smaller effect size found in experiment 3 might suggest that the desire to blame was playing at least some (though not exclusive) role in experiments 1–2.

[12] It remains an open empirical question whether or not professional philosophers will display the same sensitivities to moral features when attributing knowledge. However, see Weinberg, Gonnerman, Buckner, and Alexander (2010) for evidence that philosophers are often just as sensitive as non-philosophers to many psychological effects when reflecting on thought experiments.

[13] While experiment 3 is somewhat suggestive, it has not, of course, conclusively ruled out this hypothesis.

[14] Also see experimental evidence suggesting that the standard philosophical intuition in barn cases is not unanimous (Colaço, Buckwalter, & Stich, unpublished manuscript).

[15] Instead, the moral component of knowledge would join a chorus of ongoing research in experimental philosophy suggesting that intuitions are highly susceptible to a host of factors irrelevant to the content of the judgment in question, such as order (Swain et al., 2008), framing (Uhlman, Pizzaro, Tannenbaum, & Ditto, 2009), and environmental effects (Helzer & Pizarro, 2011).

[16] For a detailed account of this kind of argument, see Turri (2012).

References

Adams, F., & Steadman, A. (2004). Intentional action in ordinary language: Core concept or pragmatic understanding? *Analysis*, *64*, 173–181.

Alexander, J., Mallon, R., & Weinberg, J. (2010a). Accentuate the negative. *Review of Philosophy and Psychology*, *1*, 297–314.

Alexander, J., Mallon, R., & Weinberg, J. (2010b). Competence: What's in? What's out? Who knows? *Behavioral and Brain Sciences*, *33*, 329–330.

Alicke, M.D. (2008). Blaming badly. *Journal of Cognition and Culture*, *8*, 179–186.

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

Beebe, J., & Jensen, R. (2012). Surprising connections between knowledge and action: The robustness of the epistemic side-effect effect. *Philosophical Psychology*, *25*, 689–715.

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

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

Buckwalter, W. (forthcoming). The mystery of stakes and error in ascriber intuitions. In J. Beebe (Ed.), *Advances in experimental epistemology*. New York, NY: Continuum.

Buckwalter, W., & Schaffer, J. (unpublished manuscript). Knowledge, stakes, and mistakes.

Buckwalter, W., & Stich, S. (2011, April). *Epistemology and demography*. Paper presented at the 2011 meeting of the American Philosophical Association Pacific Division, San Diego, CA.

Buckwalter, W., & Stich, S. (forthcoming). Gender and philosophical intuition. In J. Knobe & S. Nichols (Eds.), *Experimental philosophy*. Oxford: Oxford University Press.

Cohen, S. (2004). Knowledge, assertion, and practical reasoning. *Philosophical Issues*, *14*, 482–491.

Colaço, D., Buckwalter, W., & Stich, S. (unpublished manuscript). Epistemic intuitions in fake-barn thought experiments. Rutgers University.

DeRose, K. (2002). Assertion, knowledge, and context. *The Philosophical Review*, *111*, 167–203.

DeRose, K. (2009). *The case for contextualism*. Oxford: Oxford University Press.

DeRose, K. (2011). Contextualism, contrastivism, and X-Phi surveys. *Philosophical Studies*, *156*, 81–110.

Fantl, J., & McGrath, M. (2010). *Knowledge in an uncertain world*. Oxford: Oxford University Press.

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

Gendler, T.S., & Hawthorne, J. (2005). A real guide to fake barns: A catalogue of gifts for your epistemic enemies. *Philosophical Studies*, *124*, 331–352.

Gettier, E. (1963). Is justified true belief knowledge? *Analysis*, *23*, 121–123.

Goldman, A. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, *73*, 771–791.

Guglielmo, S., & Malle, B.F. (2009, June). *The timing of blame and intentionality: Testing the moral bias hypothesis*. Poster presented at the annual meeting of the Society for Philosophy and Psychology, Bloomington, IN.

Halpern, J., & Hitchcock, C. (2010). Actual causation and the art of modeling. In R. Dechter, H. Geffner, & J. Halpern (Eds.), *Heuristics, probability and causality. A tribute to Judea Pearl* (pp. 383–406). London: College Publications.

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

Helzer, E., & Pizarro, D.A. (2011). Dirty liberals!: Reminders of cleanliness promote conservative political and moral attitudes. *Psychological Science*, *22*, 517–522.

Knobe, J. (2003). Intentional action and side effects in ordinary language. *Analysis*, *63*, 190–194.

Knobe, J. (2004). Intention, intentional action and moral considerations. *Analysis*, *64*, 181–187.

Knobe, J. (2010). Person as scientist, person as moralist. *Behavioral and Brain Sciences*, *33*, 315–329.

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

Lycan, W. (2006). On the Gettier problem problem. In S. Hetherington (Ed.), *Epistemology futures* (pp. 148–168). Oxford: Oxford University Press.

Machery, E. (2008). Understanding the folk concept of intentional action: Philosophical and experimental issues. *Mind & Language*, *23*, 165–189.

Malle, B.F., & Nelson, S.E. (2003). Judging mens rea: The tension between folk concepts and legal concepts of intentionality. *Behavioral Sciences & the Law*, *21*, 563–580.

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.

Nadelhoffer, T. (2006). Bad acts, blameworthy agents, and intentional actions: Some problems for jury impartiality. *Philosophical Explorations*, *9*, 203–220.

Nagel, J. (forthcoming). Intuitions and experiments: A defense of the case method in epistemology. *Philosophy and Phenomenological Research*.

Nichols, S., & Ulatowski, J. (2007). Intuitions and individual differences: The Knobe effect revisited. *Mind and Language*, *22*, 346–365.

Phelan, M. (2011). The intentional action factory. *The Philosopher's Magazine*, *52*, 72–77.

Pinillos, N.Á. (2012). Knowledge, experiments and practical interests. In J. Brown & M. Gerken (Eds.), *New essays on knowledge ascriptions* (pp. 192–211). Oxford: Oxford University Press.

Pinillos, N.Á., Smith, N., Nair, G. S., Marchetto, P., & Mun, C. (2011). Philosophy's new challenge: Experiments and intentional action. *Mind & Language*, *26*, 115–139.

Schaffer, J., & Knobe, J. (2010). Contrastive Knowledge Surveyed. *Nous*. Advance online publication. doi:10.1111/j.1468-0068.2010.00795.x.

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

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

Swain, S., Alexander, J., & Weinberg, J. (2008). The instability of philosophical intuitions: Running hot and cold on Truetemp. *Philosophy and Phenomenological Research*, *76*, 138–155.

Turri, J. (2012). Is knowledge justified true belief? *Synthese*, *184*, 247–259.

Uhlman, E., Pizzaro, D., Tannenbaum, D., & Ditto, P. (2009). The motivated use of moral principles. *Judgment and Decision Making*, *4*, 476–491.

Ulatowski, J. (2012). Act individuation: An experimental approach. *Review of Philosophy and Psychology*, 3, 249–262.

Uttich, K., & Lombrozo, T. (2010). Norms inform mental state ascriptions: A rational explanation for the side-effect effect. *Cognition*, *116*, 87–100.

Vaesen, K., & Peterson, M. (unpublished manuscript). The reliability of armchair intuitions. Eindhoven University of Technology.

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

Weinberg, J., Nichols, S., & Stich, S. (2001). Normativity and epistemic intuitions. *Philosophical Topics*, *29*, 429–460.

Young, L., Cushman, F., Adolphs, R., Tranel, T., & Hauser, M. (2006). Does emotion mediate the effect of an action's moral status on its intentional status? Neuropsychological evidence. *Journal of Cognition and Culture*, *6*, 291–304.

Zagzebski, L. (1994). The inescapability of Gettier problems. *The Philosophical Quarterly*, *44*, 65–73.