

## Replies to Ichikawa, Martin and Weinberg

Timothy Williamson

Published online: 15 April 2009  
© Springer Science+Business Media B.V. 2009

### 1 Reply to Ichikawa

*The Philosophy of Philosophy* analyses the Gettier thought experiment as an argument whose major premise is a counterfactual conditional. In ‘Knowing the Intuition and Knowing the Counterfactual’, Jonathan Ichikawa objects that this misrepresents the thought experiment as more accident-prone than it really is. If the world does not cooperate, the counterfactual will fail: it will be false that if the Gettier text had been realized, there would have been justified true belief without knowledge. Even if the world cooperates enough to make the counterfactual true, it may still not cooperate enough to enable us to know it, if it could too easily have been false. Ichikawa denies that the thought experiment is such a hostage to empirical fortune. *The Philosophy of Philosophy* considers such objections, arguing that although the thought experiment is indeed not immune to misfortune, that should not drive us to scepticism. If we identify an unwanted way in which the Gettier text might well have been realized, we can easily fix it by extending the text to rule out that way (pp. 200–204). There is no need to pretend that we had already fixed the problem before we had even thought of it. Rather than repeating those arguments, I will inquire whether Ichikawa’s proposed alternative does any better.<sup>1</sup>

The main idea behind his proposal is to replace the Gettier text by the Gettier story. The latter strictly implies far more than the former does, because many things are true in a story without having been explicitly stated in the text of that story. Thus, unlike the Gettier text, the Gettier story is supposed to strictly imply that there is justified true belief without knowledge, so that the problem of unwanted realizations no longer arises. However, we are not to conceive thinkers performing

---

<sup>1</sup> He refers the reader to Ichikawa and Jarvis (2008) for more details.

---

T. Williamson (✉)  
New College, Holywell Street, Oxford OX1 3BN, U.K.  
e-mail: timothy.williamson@philosophy.ox.ac.uk

the thought experiment as conceptualizing the Gettier story to themselves *as* what is true in the fiction presented by the Gettier text. Rather, given their familiarity with the practice of fiction, reading the text puts them in a position to entertain a more specific proposition, which they might articulate to themselves as ‘Things are like *that*’.<sup>2</sup> They can then argue: it is possible for things to be like that; necessarily, if things are like that then there is justified true belief without knowledge; therefore, it is possible for there to be justified true belief without knowledge.

The significance of the proposal does not lie in the use of strict implication rather than the counterfactual conditional to formulate the major premise. For a strict implication entails the corresponding counterfactual, and the latter suffices to validate the passage from the possibility of its antecedent to the possibility of its consequent, while making an epistemically less risky claim. In fact, Ichikawa could take over exactly the formalization of the Gettier argument that the book recommends, counterfactual and all. The only difference would lie in the interpretation of the predicate **GC**. Whereas I explain it in terms of realizing the Gettier text, Ichikawa would explain it in terms of realizing the Gettier story, or verifying the proposition that things are like *that*. What matters for his proposal is that the strict implication be true, not that it be used as the major premise. Its truth is what is supposed to remove the element of epistemic luck.

An obvious danger for Ichikawa’s account is that truth in fiction may itself depend on counterfactuals that one can be wrong about if the world does not cooperate, as in David Lewis’s classic account (1978). In a world in which it is highly abnormal not to have many alternative sources of knowledge for a given belief, Gettier’s text may present a fiction in which it is not true that the protagonist has justified true belief without knowledge. If we are in an abnormal pocket of ignorance within such a world, then the Gettier story does not strictly imply that there is justified true belief without knowledge; perhaps it even strictly implies that there is no justified true belief without knowledge. Thus it is not obvious that Ichikawa’s account avoids the epistemic risks to which he objects in mine. That we articulate the proposition as ‘Things are like *that*’ rather than ‘Whatever is true in the Gettier fiction is true’ does not help us avoid the risk of error when we believe that proposition to strictly imply a case of justified true belief without knowledge.

Ichikawa does not address the challenge in exactly that form. He does tell us that when one considers the Gettier text, one ‘enriches it, considering a more-determinate scenario’. He allows that ‘divergent private fillings-out of the scenario’ are possible, but suggests that we normally avoid them because ‘we have particular conventions, grounded in our practices with fictions, that govern how to move from a weaker description to a stronger scenario in the intended way’. This suggests a psychological process, inspired by the Gettier text, of constructing our own fiction, whose additional content can be generated in part by our false beliefs. Thus, even if it is not really true in the Gettier fiction that there is justified true belief without knowledge, the proposition that *we* express by ‘Things are like *that*’ may nevertheless still imply that there is justified true belief without knowledge.

<sup>2</sup> Ichikawa and Jarvis (2008).

Is such an account psychologically plausible? Often, when we perform a thought experiment, we simply read the text and make a verdict, or hesitate to do so. We need not sit there, daydreaming, adding detail after detail to the story. Visual imagery is absent or, if present, usually irrelevant (one may imagine the protagonist of Gettier's story with black hair). Entertaining the proposition that things are like *that* is not just a matter of demonstrating a possible world that one can already see in one's mind's eye.

In Ichikawa's example, the original Gettier-like text is this:

T1 At 8:28, somebody looked at a clock to see what time it was. The clock was broken; it had stopped exactly twenty-four hours previously. The subject believed, on the basis of the clock's reading, that it was 8:28.

Ichikawa points out that T1 conjoined with the originally unintended T2 no longer works as a Gettier case:

T2 The subject knew in advance that the clock had stopped exactly twenty-four hours previously.

On Ichikawa's view, the normal reader of T1 constructs a more specific proposition than T1 expresses that is incompatible with the proposition that T2 expresses. But how, exactly? Normal readers of T1 are *taken by surprise* when the possibility of elaborating T1 with T2 is pointed out to them. They have not already gone through a psychological process of formulating T2 and elaborating T1 with its negation, still less with the negation of T3 or of T4<sup>3</sup>:

T3 The subject knew in advance that the clock had stopped an exact multiple of twenty-four hours previously.

T4 The subject knew in advance that she suffered from a rare psychological condition that rendered her incapable of reading a clock as saying 8:28 at any time other than 8:28.

If the normal reader of T1 has excluded T2, T3 and T4 before even having formulated them, it is not through a psychological process of explicitly formulating something that, combined with T1, logically implies their negations. The original exclusion, if psychological at all, is dispositional rather than occurrent. The relevant disposition might be to add the negations of T2, T3 and T4 to T1 if the question arises. On this view, what proposition the reader expresses by 'Things are like *that*' is determined by the set of sentences she is disposed to add to T1 if the question arises.

Such a view faces many problems.

First, the sentences a given reader is disposed to add to T1 if the question arises are quite likely to form an inconsistent set. This need not mean that she is disposed to continue T1 in an inconsistent way. She may rather be disposed to continue T1 in different individually consistent but jointly inconsistent ways, depending on which

<sup>3</sup> Readers of this exchange between Ichikawa and me may be abnormally on the look-out for such loopholes, but even they cannot specifically anticipate each of them.

question arises first. For example, a mildly suggestible reader may be disposed to add either T5 or T6 to T1, depending on which question arises first:

T5 The subject was not wearing a wristwatch.

T6 The subject was wearing a wristwatch, but it had stopped.

After all, each of T5 and T6 seems to be a reasonably adequate and natural way of filling out the text of T1. But once such a reader has added T5 to T1, she may lose the disposition to add T6, and vice versa. Even the dispositions to add T2, T3 and T4 may be sensitive to what else has already been added. It is by no means obvious how to construct a metaphysically possible proposition out of such a complex web of interrelated dispositions. Yet Ichikawa's rendering of the Gettier argument requires such a proposition, to verify the premise that it is possible for things to be like *that*. Indeed, he requires us to *know* that it is possible, if the argument is to give us knowledge of its conclusion, that it is possible for there to be justified true belief without knowledge. Thus Ichikawa's account makes the epistemology of the first premise of the Gettier argument problematic in ways that mine does not.

Second, an account of the content of 'Things are like *that*' in terms of the reader's psychological dispositions to fill out the original text almost guarantees that different readers will express different contents by the sentence. Ichikawa's appeal to our familiarity with the conventions of fiction merely suggests that those contents will tend not to differ too greatly from reader to reader. By hypothesis, the contents do not in general coincide with the shared story itself, on pain of undermining his implied account of the epistemology of the major premise. When two readers apparently disagree in their verdicts on a thought experiment, that already constitutes a difference in their psychological dispositions to fill out the original text, and so provides *prima facie* reason to believe that they associate different contents with 'Things are like *that*'. If so, they are not really disagreeing on Ichikawa's account. But that is not how philosophical discussion of thought experiments works. We allow the text of the thought experiment to fix a shared content, and then discuss that scenario.

Third, even if one is performing a thought experiment in isolation from other people, one's verdict on it is still in general a judgment that can be true or false independently of one's disposition to make it. Only a very foolish philosopher, on coming to a verdict on a tricky example in ethics, would think 'I can't be wrong about this; it's just how I'm envisaging the case'. Ichikawa's account has not yet shown enough distance between the specification of the scenario in a thought experiment and our dispositions to make further judgments as to what would be the case in that scenario. Providing that distance will tend to reintroduce the epistemic risks with respect to the major premise of the Gettier argument that he complains about in my account.

Whatever the details of Ichikawa's account, it exemplifies a strategy against which I argued in *The Philosophy of Philosophy*: the attempt to combat scepticism about our evidence in philosophy by psychologization. His application of the strategy is quite subtle, since what he psychologizes is not the content of the

premises of the Gettier argument but our relation to that content. Rather than grasping it by understanding the public verbal description of the scenario, our filling out of the description is supposed to acquaint us psychologically with the content in a way that not only allows us to articulate it as ‘Things are like *that*’ but somehow enables us to avoid the normal risks of error in judging what would be the case in a scenario. The strategy is no more successful here than elsewhere.

We cannot realistically expect that the method of thought experiments in philosophy will turn out to be much *more* reliable than the methods of the natural sciences. What, then, is so bad about accepting that the former is not immune from a mass of easily corrected small errors like those to which the latter are quite obviously subject?

## 2 Reply to Martin

In ‘Reupholstering a Discipline’, Mike Martin contests two themes of *The Philosophy of Philosophy*: the idea that philosophy makes progress and knowledge maximization as a principle of charity in interpretation. I will discuss his comments on each theme in turn.

1. Martin suggests that the idea of progress in a discipline, although applicable to mathematics and the natural sciences, fails to fit some of the humanities, such as history, so that we should not be too surprised if it also fails to fit philosophy. In making that suggestion, he does not intend to align himself with post-modern conceptions of inquiry as merely continuing or subverting a conversation.

Does the comparison with history help Martin’s case? Of course, historians tend to distance themselves from the idea of progress in history as the succession of human doings and sufferings. But that is not what Martin means. He is questioning whether history as a discipline makes progress. He says very little to support a negative answer. Yet significantly more is known about a vast range of historical matters than was known fifty years ago. That is not just fact-grubbing. Significantly more is understood about those matters too. For instance, the role of religious belief in the English Civil War is better understood than it was. Again, history has vastly extended the range of its inquiries, for instance into the lives of members of marginalized groups. Its methods have developed, not only through the application of advances in science, such as statistics and DNA analysis. More critical and more imaginative ways of learning from documents are available. It seems merely quixotic to deny that if such changes really have taken place, they constitute progress in history as a discipline. Of course, someone may deny that the role of religious belief in the English Civil War is really any better understood now than it was fifty years ago. But, if so, that neither would nor should be a matter of indifference to the historians concerned. It would be a failure on their own terms. Martin does not seem ready to endorse such scepticism about history as a discipline. Although the progress it makes is piecemeal, the sort of progress I suggested philosophy makes is piecemeal too. Thus the comparison with history, to the extent to which it is relevant, undermines Martin’s case against progress in philosophy.

Martin is more willing to countenance progress for individual philosophers than for the discipline as a whole. But this individualistic preference is hard to reconcile with the social nature of philosophy. Many of us see ourselves as participating in a collective enterprise, one that goes back to the ancient Greeks if not further and, we hope, will continue for millennia to come. The point of publishing is not only to advance one's career or benefit from feedback. It is to contribute, if only in a very small way, to an inquiry that will continue after one's death. For example, one tries to bring to others' attention a possibility they have missed, so that in the long run philosophers may determine whether it actually obtains. One defines one's individual goals in relation to progress in the discipline.

Even at the level of the individual, Martin is cautious about the idea of philosophical knowledge, and seems happier speaking of understanding. The thought that, individually or collectively, we might progress in understanding without progressing in knowledge is a familiar cop-out in defence of the humanities. Does it withstand scrutiny? If you don't know why Rome fell, you don't understand why Rome fell. If you do know why Rome fell, you have at least made considerable progress towards understanding why it fell. You might know that such-and-such caused Rome to fall without knowing why such-and-such caused Rome to fall, but what you lack there is more knowledge. Nor is knowing why some mysterious sort of non-propositional knowledge. What constitutes knowing why the barn caught fire is, for instance, knowing that the burning match Innocent dropped in the straw caused the fire. It is not radically different in philosophy. If you don't know why zombies are impossible, you don't understand why zombies are impossible. In such cases, the idea that understanding transcends knowledge depends on too narrow a conception of the facts to be known.<sup>4</sup> In any discipline, practitioners may be reluctant to summarize progress in a one-liner.

2. In the latter part of his comments, Martin develops a putative counterexample to knowledge maximization as a principle for determining reference. He assumes that one can simultaneously attend visually to each of two objects in different parts of one's visual field; I am quite happy to grant that assumption for the sake of argument. In the imagined case, one is simultaneously attending visually to each of two qualitatively identical pink Dolly Mixtures, righty and lefty, while attending in thought only to lefty. One thinks 'That's pink', referring only to lefty. Martin suggests that knowledge maximization cannot determine whether the reference is to righty or to lefty, because the sentence would express knowledge either way. He stipulates that the 'salient non-intentional facts' are symmetrical between righty and lefty.

Presumably, Martin is not suggesting that the symmetry between righty and lefty at the non-intentional level is perfect. If asked 'Which one?', one will answer 'The one on the left', not 'The one on the right'. That difference corresponds to some causal asymmetry at the non-intentional level. The question is whether my account can engage with that asymmetry. It can if the question is asked; as Martin notes, only the correct assignment of reference makes the thought 'That's furthest to the left' knowledgeable. But he stipulates, legitimately, that one has no such actual

<sup>4</sup> In this respect knowing why is worth comparing with knowing how; see Stanley and Williamson (2001).

thought in the given case. The book does not formulate knowledge maximization in counterfactual terms.

The account of reference determination in the book appeals to naturalness at more than one point. Martin designs the case to pre-empt a straightforward appeal, but neglects a subtler possibility. Knowledge itself should be a somewhat natural relation (p. 268). If one refers to lefty in the counterfactual case in which the question ‘Which one?’ subsequently occurs, then it is more natural for the reference to be to lefty in the actual case in which the question does not occur, for the two cases begin the same. A difference in reference between them would make knowledge less natural.

Perhaps we can go deeper. For although one is visually attending to both lefty and righty equally, it does not follow that one’s thought ‘That’s pink’ has symmetrical causal relations to righty and lefty. The simplest suggestion would be that it is causally explained by the fact that lefty is pink and not by the fact that righty is pink. That is insufficiently general, for in one version of the example I first think ‘They are both pink’, an event which does have symmetrical causal relations to righty and lefty, and then infer ‘That’s pink’ of lefty, with no further causally relevant input from the colour of lefty. In such a case, what—if anything—do we imagine the reference of ‘that’ to lefty rather than righty as consisting in?

An appealing picture is that when one starts using ‘that’ with reference to lefty, one opens some sort of mental file, if only a very temporary one, with a predominant causal connection to lefty that enables it to act as a channel for perceptual information about lefty more directly than it can act as a channel for perceptual information about righty: potentially, a channel for *knowledge* of lefty rather than righty. It does not matter whether that knowledge includes the particular item that it is pink, for the thought ‘That is pink’ still has a compositional semantics. Even if no knowledge actually happens to be gained through that channel, the naturalness of the reference relation may still keep the reference constant between the actual case and counterfactual cases in which knowledge of lefty is gained through the channel.

That picture is not the only one consistent with the knowledge maximization principle. However unappealing on other grounds, a descriptivist account on which ‘that’ somehow abbreviates a description such as ‘The one on the left’ or ‘The one actually on the left’ is also consistent with the principle. Other, more complex possibilities are consistent too. But if we suppose that *no* such story is to be told in the example, it ceases to be a clear case of reference to lefty; a verdict of reference failure becomes much more plausible. The case was not intended to be one in which the reference to lefty was puzzling to everyone; it was intended to discredit the knowledge maximization principle by providing a puzzle distinctive to it. But, as just seen, it can be accommodated within the book’s framework of knowledge maximization and naturalness.

### 3 Reply to Weinberg

Experimental results can in principle undermine the procedures of any intellectual community, by revealing patterns of variation in its members’ judgments that are



hard to reconcile with the supposition that those judgments are even moderately reliable. It does not follow that every intellectual community should suspend its procedures until the relevant experiments have actually been done and shown to have reassuring results, otherwise all inquiry would come to a halt, since the procedures for interpreting experimental results would themselves have been suspended. In 'On Doing Better, Experimental-Style', Jonathan Weinberg recognizes that the experimentalist challenge to the armchair methods of philosophy must do better than appeal to the mere sceptical possibility of seriously disquieting experimental results. Such results must actually have been obtained, if the challenge is to attain any urgency. According to Weinberg, such results have indeed been obtained, enough of them to remove any initial presumption in favour of the procedures of a well-established discipline like philosophy.

In *The Philosophy of Philosophy*, I complained about experimentalists' use of experiments on undergraduates just beginning philosophy to cast doubt on its methods as applied by highly trained practitioners (p. 191). Weinberg responds that the existence of 'real expertise' in philosophy is just another empirical hypothesis in need of experimental test. By the same token, the existence of real expertise in physics is just another empirical hypothesis in need of experimental test. As before, the question is how urgent the challenge is. Weinberg does not directly claim that any actual experimental results cast doubt on the existence of real expertise in philosophy. He does cite a paper by Shanteau (1992) as showing something about the characteristics of domains conducive to the development of real expertise.<sup>5</sup> It will be useful to consider whether Shanteau's work really does support Weinberg's argument.

One of Shanteau's main conclusions is that expert competence fares better where the stimuli to be evaluated are static, and worse where they are dynamic (as in making real-time judgments about unfolding events). A moving target is harder to hit. Armchair philosophy typically involves the evaluation of constant stimuli, such as the scenario of a thought experiment, often presented by a written description, so in that respect Shanteau's paper is encouraging. He also suggests that it is easier to achieve expert competence with stimuli that involve things, and harder with stimuli that involve human behaviour. In that respect general metaphysics *may* be better off than moral philosophy, although it is unclear whether Shanteau would count a written description of a moral predicament as a stimulus that involves things or as one that involves human behaviour.

<sup>5</sup> In Shanteau's terms, some armchair philosophers are clearly experts in philosophy, since he in effect defines an 'expert' in a field as someone generally regarded as an expert in the field by those who work in that field, and some armchair philosophers are clearly generally regarded as experts in philosophy by philosophers. Similarly, he defines 'competence' in a field as what the experts in the field generally regard as competence in the field; by that standard, there is clearly competence in armchair philosophy, since some of it is clearly generally regarded by those generally regarded by philosophers as experts in philosophy as showing competence in philosophy. Reliance on such operational definitions is widespread in the literature on expertise that Weinberg cites, because non-experts often have no expert-independent way of assessing expertise. This is one of several reasons why the bearing of that literature on the status of philosophy is much less direct than Weinberg appears to suggest. For the sake of argument, I will go along with his apparent assumption that Shanteau's conclusions about the bearing of task characteristics on competence in experts apply similarly to real expertise; unless he is assuming that, it is not obvious why he cites Shanteau's paper. I concentrate on that paper because the other work that Weinberg cites from the expertise literature, Ericsson et al. (2006), is even less relevant to his claims.



Shanteau also argues that expert competence tends to be associated with tasks of recurrent types, on which frequent feedback is possible. Of course, intellectually taxing research in any discipline is not wholly repetitive. Nevertheless, anyone with a Ph.D. from a programme in analytic philosophy is likely to have received feedback from their teachers on their evaluations of scores of thought experiments and arguments, many of them variations on recognizable themes. The feedback process continues throughout an academic career in philosophy in the form of reactions from colleagues, audiences and referees. In this respect too, armchair philosophy does not seem especially badly off.<sup>6</sup>

Another feature of experts, according to Shanteau, is that they tend to decompose complex problems. Armchair philosophy lends itself to that process. For example, in determining whether a thought experiment provides a counterexample to a complex proposed analysis, one often decomposes the task into subtasks corresponding to the sub-clauses of the *analysans*.

The literature on expertise that Weinberg cites does not constitute even a *prima facie* challenge to the natural assumption that there is real expertise in armchair philosophy. His challenge to the assumption is not urgent. More specifically, he does not identify any respect in which he has shown the challenge to be more urgent for philosophy than it is for other academic disciplines, such as physics or psychology. It is no better than routine scepticism about the results of unspecified experiments that have not been carried out. Scientists perform only a tiny fraction of all the experiments that it would be physically possible for them to perform. Virtually any scientific theory has implications for the results of experiments that will never be performed. If scientists had to remain neutral about the results of all unperformed experiments, they would have to avoid commitment to virtually any scientific theory.

The question of expertise arose as a challenge to experimental philosophers' reliance on data about beginners in philosophy to cast doubt on work by experienced, intensively trained armchair philosophers. Since Weinberg's response to the challenge in his paper provides no evidence against the assumption that there is real expertise in philosophy, it fails to legitimize his use of such data.

Is Weinberg's use of other psychological data any more convincing? He cites evidence that verdicts on thought experiments are affected by factors that vary independently of the correctness of those verdicts. For instance, the verdicts sometimes show sensitivity to the order in which the thought experiments are considered, to differences in wording between logically equivalent descriptions, and to whether they are made in a clean, tidy environment or a dirty, messy one.<sup>7</sup> The conflicting verdicts cannot all be correct.<sup>8</sup> However, in using these data Weinberg

---

<sup>6</sup> Weinberg's paradigm of an intellectual method that required abandonment rather than reform is introspectionism in psychology. The scope for feedback from others on introspective reports is notably narrower than it is on judgments in armchair philosophy.

<sup>7</sup> The importance of order effects in verdicts on thought experiments is, of course, already emphasized in Williams (1970).

<sup>8</sup> Attempts to reconcile the verdicts by contextualist hypotheses about their content would be far-fetched in most of these cases. Williamson (2005) suggests that the data used to support contextualist or subjective-sensitive invariantist hypotheses in epistemology are better explained in terms of errors induced by giving too much weight in some settings to factors that are psychologically salient in those settings.

ignores the difference between one-off individual judgments and consensus reached through the interaction of many participants in a public philosophical debate, conducted over several years in conferences and journals. In the course of such a debate, most participants are forced by their opponents to consider the thought experiments in orders and wordings more favourable to those opponents. Some participants consider them in clean, tidy environments, others in dirty, messy environments, many sometimes in one and sometimes in the other. Of course, these interpersonal and intrapersonal variations do not guarantee convergence on the right answer. Nevertheless, if the initial individual judgments are more accurate than chance, without being perfectly reliable, then the majority view has a higher probability of being right. Such social controls are as important in philosophy as they are in the natural sciences, and the social dimension of philosophy is frequently emphasized in *The Philosophy of Philosophy*, for example on the first page. Weinberg would have to work much harder to show that verdicts on thought experiments are no more accurate than chance, especially since the relevant data include uncontroversial thought experiments as well as controversial ones. Thus his appeals to such framing effects are undermined by his neglect of the psychological and social conditions of actual philosophical practice.

Weinberg repeatedly cites the use of double-blind methods as a way in which science has learned to do better. So it has, but in its modest way armchair philosophy uses double-blind methods too where appropriate, most notably in the refereeing of submissions to journals. Although many such social mechanisms in philosophy are common to most academic disciplines, they are none the worse for that.

Weinberg makes no attempt to specify the psychological or social nature of armchair philosophy. He sprays his experimental data in its general direction, as though everything in the area deserves to be hit at least by the charge of not currently deserving our confidence. Such an indiscriminate approach is peculiarly liable to shoot itself in the foot, or worse. Its targets include informal qualitative epistemological judgments, such as verdicts on Gettier cases. But Weinberg's paper is itself full of informal qualitative epistemological judgments, for example about whether we are justified in believing that armchair methods in philosophy are reliable. Nor could any current natural science proceed without such judgments. Even statistical data need to be interpreted; the judgment that they render some hypothesis untenable remains an informal, qualitative one, whatever formal and quantitative considerations it draws on. Presumably, Weinberg thinks that we are entitled to accept many informal qualitative epistemological judgments made by natural scientists, without special qualms about their reliability. He does not seem to think that we are entitled to accept many informal qualitative epistemological judgments made by armchair philosophers, without special qualms about *their* reliability. Does anything in his data justify this differential attitude?

It would not help Weinberg to say that the natural scientists' epistemological judgments are supported by empirical data while the armchair philosophers' are not. First, the relevant judgments concern the relation between data and theory, not the correctness of the data themselves. Both natural scientists and philosophers can make them in the armchair. Although natural scientists' epistemological judgments

may be informed by background knowledge, the book shows that the same is true of verdicts on Gettier cases (p. 185). In any case, there are real-life Gettier cases as well as imaginary ones; for epistemological purposes, it matters little which sort one uses (pp. 192–193). Second, Weinberg offers no evidence (experimental or otherwise) that informal qualitative epistemological judgments are more reliable about real-life cases than about imaginary ones.

In practice, Weinberg lays down the experimental challenge for armchair philosophy and simply fails to mention that it could be laid down for natural science too. Of course, experimental philosophers have tested for framing effects in informal qualitative epistemological judgments in armchair philosophy without testing for framing effects in informal qualitative epistemological judgments in natural science.<sup>9</sup> Framing effects threaten to constitute a rather general problem for human cognition, although not a wholly insuperable one. The experimental philosophers' practice of testing for them in philosophy and not elsewhere is analogous to that of a group of men who spend their time testing for framing effects in women's judgments, and find many. They never do the tests on men. They conclude that women's judgments are unreliable, and not to be trusted. When asked about men's judgments, they reply that since there is no evidence that they are unreliable, they can be trusted. Such experimental misogyny would not deceive Weinberg, but his own experimental anti-philosophy is scarcely more respectable from a scientific point of view. The experiments are not properly controlled, because the experimenter is looking for framing effects only where it suits him to find them.

Some experimental results on human judgment *are* disquieting. We really are less reliable than we thought we were. Our judgments are often influenced by irrelevant factors. We need to map out our intellectual vices, in order to manage them more effectively. Many disciplines have in effect already evolved methods that may allow them to work round some of the worst effects of the vices. We can reasonably hope that future advances in cognitive psychology will enable us to do better. *The Philosophy of Philosophy* is quite explicit that philosophy can learn from experiment (not just of the thought kind), and itself applies experimental work on the psychology of reasoning to philosophical issues (pp. 6, 102–106). But one point of the book is that any psychological kind that includes armchair philosophical judgments includes a mass of non-philosophical judgments too. In order to manage framing effects more successfully, we need to know more about what the relevant psychological kinds are. There is no reason to believe that they will be restricted to psychological processes that are dispensable in the way in which experimental philosophers may suppose armchair philosophy to be dispensable.

On the basis of the evidence that Weinberg offers, the idea that armchair philosophy is peculiarly at risk from experimental results is a bluff. The experimental critique discredits itself by confusing a scientific spirit with a scientific one. Bad science does not make good philosophy.

---

<sup>9</sup> The possibility of framing effects in natural science should hardly come as a surprise after Kuhn (1962) and much subsequent empirical work on the practice of science. Weinberg speaks of scientific practices as 'unchallenged'; they are not unchallenged in general, just by experimental philosophers.

**Acknowledgement** The author thanks the three commentators for their interesting questions, and to all the participants at the Arché workshop in St Andrews which led to this symposium for discussion, including Stephen Stich, who coauthored the paper with Jonathan Weinberg as presented there.

## References

- Ericsson, K. A., Charness, N., Feltovich, P. J., & Hoffman, R. R. (2006). *The Cambridge handbook of expertise and expert performance*. Cambridge: Cambridge University Press.
- Ichikawa, J., & Jarvis, B. (2008). Thought-experiment intuitions and truth in fiction. *Philosophical Studies*, 142, 221–246.
- Kuhn, T. S. (1962). *The structure of scientific revolutions*. Chicago: Chicago University Press.
- Lewis, D. (1978). Truth in fiction. *American Philosophical Quarterly*, 15, 37–46.
- Shanteau, J. (1992). Competence in experts: The role of task characteristics. *Organizational Behavior and Human Decision Processes*, 53, 252–266.
- Stanley, J., & Williamson, T. (2001). Knowing how. *Journal of Philosophy*, 98, 411–444.
- Williams, B. (1970). The self and the future. *The Philosophical Review*, 79, 161–180.
- Williamson, T. (2005). Contextualism, subject-sensitive invariantism and knowledge of knowledge. *The Philosophical Quarterly*, 55, 213–235.