

Intuitions, reflective judgments, and experimental philosophy

Michael Hannon¹ 

Received: 22 July 2016 / Accepted: 20 April 2017
© Springer Science+Business Media Dordrecht 2017

Abstract Experimental philosophers are often puzzled as to why many armchair philosophers question the philosophical significance of their research. Armchair philosophers, in contrast, are often puzzled as to why experimental philosophers think their work sheds any light on traditional philosophical problems. I argue there is truth on both sides.

Keywords Intuition · Reflective judgment · Experimental philosophy · Philosophical methodology · Skepticism · Ought implies can

In their “Experimental Philosophy Manifesto”, Joshua Knobe and Shaun Nichols highlight the polarizing effect of experimental philosophy. They write,

Many find it an exciting new way to approach the basic philosophical concerns that attracted them to philosophy in the first place. But many others regard the movement as insidious—a specter haunting contemporary philosophy. (Knobe and Nichols 2008: p. 3)

Philosophers incline towards one of two opposite extremes when it comes to experimental philosophy. On the one hand, experimental philosophers are often puzzled as to why anyone could fail to find philosophical significance in the questions addressed by their research. On the other hand, many armchair philosophers are puzzled as to why experimental philosophers think their work sheds light on traditional philosophical problems. I will argue there is truth on both sides.

The most vigorous debate surrounding experimental philosophy is about the philosophical significance of empirical evidence on *intuitions*. An idea “widely endorsed by

✉ Michael Hannon
mhannon09@gmail.com

¹ Queen’s University, Kingston, Canada

experimental philosophers” is that only survey-based empirical research “can deliver the intuitions that can serve as [an] evidential basis for or against philosophical claims” (Alexander and Weinberg 2007: p. 61). Thus, some of the best-known work in experimental philosophy has involved surveys designed to probe people’s intuitions on various philosophical issues. Knobe even characterizes experimental philosophy as a “philosophical movement that proceeds by conducting systematic experimental studies of people’s ordinary intuitions” (2007: p. 119).

While some experimental studies have confirmed armchair predictions about what ‘the folk’ will find intuitive, this body of empirical work has largely contrasted what armchair philosophers have assumed, thereby challenging the idea that people think about these issues in anything like the way philosophers have claimed. For example, some evidence suggests that non-philosophers do not attribute knowledge, free will, and deontic modals in the ways that philosophers have expected.¹ As a result, an accomplishment often touted for experimental philosophy is its potential to discredit philosophical views that have gone largely unquestioned, thereby casting doubt on a standard way of doing philosophy.

The aim of my paper is to pose a constructive challenge to some experimental attacks on traditional philosophical views. However, it would obscure the dialectic to call my target ‘experimental philosophy’. This movement has diverse ambitions and it should therefore be characterized in a broad way. Following Rose and Danks (2013), we may broadly characterize experimental philosophy as any experimental inquiry with a philosophical purpose. Although discussions of experimental philosophy often focus on the practice of surveying the intuitions of non-philosophers, this is a fairly limited conception of the movement as a whole. There is simply too much variability to draw conclusions about all experimental philosophy, so there is little point in debating the merits or deficiencies of ‘experimental philosophy’. Rather, as Rose and Danks suggest, we should focus on particular studies, proposals, or uses of empirical data.²

In keeping with this suggestion, my examination will take the form of a case study of two experimental attacks on traditional philosophical views. These two studies, like a lot of work in experimental philosophy, are written against the background of a conception of philosophy according to which intuitions play an important argumentative role.³ According to the first study, folk moral judgments do not conform to the widely assumed principle that ‘ought’ implies ‘can’; thus, philosophers who claim this principle is intuitive because it reflects commonsense moral judgments are said to be deeply mistaken (Chituc et al. 2016). According to the second study, many people do not have the intuitions that are presumed to be necessary to motivate philosophical skepticism; thus, skepticism is said to be much less interesting and much less worrisome than philosophers have claimed (Nichols et al. 2003). Both of these studies challenge a traditional and widely accepted view in philosophy.

¹ See Weinberg et al. (2001), Nahmias et al. (2006), and Chituc et al. (2016).

² For similar suggestions, see Nadelhoffer and Nahmias (2007), Knobe and Nichols (2008), and Sommers (2011).

³ For a contrary view, see Deutsch (2010, 2015) and Cappelen (2012). For replies, see Bengson (2014), Weinberg (2014), and Nado (2016).

I will question the philosophical significance of these survey-based results. More specifically, I will argue that the intuitions tested in the first study are not part of the evidence for the principle that ‘ought implies can’, and the intuitions tested in the second study do not play an important evidential role in motivating philosophical skepticism. In other words, the data gathered by these studies are irrelevant to the traditional philosophical views they purportedly undermine.

Unlike many critics of experimental philosophy, however, I do not completely reject the conception of philosophy according to which the intuitions play important argumentative roles. I defend a moderate view according to which the philosophical relevance of intuitions varies from topic to topic. I will also make concrete suggestions about how to study philosophically relevant judgments empirically. Ultimately, my analysis tries to better understand the scope, significance, and limitations of both experimental and traditional philosophical methods.

1 What are ‘intuitions’?

Before we examine the philosophical significance of experimental evidence on intuitions, we must first get a grip on what intuitions are. The word ‘intuition’ is used with a variety of extensions in both ordinary language and philosophy. Even if we limit our focus to the role of intuitions in philosophical inquiries, there is little agreement about how to best characterize them. [Lewis \(1983\)](#) equates intuitions with beliefs, [van Inwagen \(1997\)](#) says they are dispositions to believe, [George Bealer \(1998\)](#) claims they are marked by special phenomenal characteristics, and [Kauppinen \(2007\)](#) argues they are mental states that have been subjected to critical examination as part of reflective participation in traditional philosophical discourse.

Fortunately, this plurality of conceptions needn’t forestall our investigation. I will focus on ‘intuitions’ in roughly the following sense: spontaneous (snap, immediate, first-off) judgments about cases. This captures what many experimental philosophers have in mind when they investigate ‘intuitions’. For example, [Shaun Nichols, Jonathan Weinberg, and Stephen Stich](#) define an ‘intuition’ as a “spontaneous judgment about the truth or falsity of a proposition” (2012: fn. 2). Intuitions in this sense can usefully be contrasted with ‘reflective judgments’, which are judgments held after some reflection or consideration. Admittedly, this is a rough distinction; but it is not without merit, and marking it will facilitate our discussion.

Thanks to the growing movement of experimental philosophy, there is now a large body of evidence on people’s intuitions about a variety of philosophical topics.⁴ It is, however, a further question whether this kind of experimental data can shed light on philosophical problems. The most promising way to answer this question is to determine whether intuitions are actually treated as an important source of evidence in philosophy. To the extent that philosophers are interested in developing theories (of knowledge, belief, responsibility, free will, etc.) that accommodate intuitions, this data will be philosophically significant. Thus, we must ask whether the intuitions tested by

⁴ For overviews of this literature, see [Knobe and Nichols \(2008, 2013\)](#), [Alexander \(2012\)](#), [Machery and O’Neill \(2014\)](#), [Lombrozo et al. \(2014\)](#), and [Sytsma and Buckwalter \(2016\)](#).

experimental philosophers really are part of the evidence for philosophical theories. To tackle this question, I will focus on two prominent studies.

2 Study one: ‘ought’ implies ‘can’

According to the principle ‘ought implies can’, someone ought to do something only if they are able to do it. This principle played a central role in the work of Immanuel Kant, and has been widely accepted since. In *Religion Within the Boundaries of Mere Reason*, Kant writes: “For if the moral law commands that we *ought* to be better human beings now, it inescapably follows that we must be *capable* of being better human beings” (1793: 6:50). In other words, our obligations cannot exceed our abilities. This idea is not only treated as a basic test of moral obligations, it has also been employed in several debates in ethics and related areas; for instance, it has been used to address the issue of free will vs. determinism, moral dilemmas, and internalist vs. externalist accounts of moral motivation. Many philosophers have argued this principle is true not only universally, but also necessarily, analytically, or conceptually (Vranas 2007: p. 171; Zimmerman 1996: p. 79). In other words, ‘ought’ is supposed to imply ‘can’ by virtue of the concepts expressed by those very words.

In a fascinating paper recently published in *Cognition*, a team of researchers at Duke explore whether there is a conceptual entailment from what someone ‘ought’ to do to what they ‘can’ do, as those concepts are ordinarily understood (Chituc et al. 2016). If ‘ought’ analytically or conceptually implies ‘can’, as most philosophers assume, then people should deny that an agent ought to do something the agent cannot do. However, some empirical data provide a reason to be skeptical of such a relationship between ‘ought’ and ‘can’ in moral judgment. Three experiments were used to test the cognitive underpinnings of these concepts, and the results indicate that folk moral judgments do not conform to the widely assumed philosophical principle that ‘ought’ implies ‘can’ (or OIC for short). More specifically, the data suggest that participants will judge that an agent ought to do something the agent can’t do when the agent is to blame for the inability. In other words, it is blame, not ability, that seems to impact moral ‘ought’ judgments for impossible actions.⁵

These findings are said to have normative significance. Chituc and his colleagues say their results

pose a serious challenge for the many philosophers who hold that ‘ought’ implies ‘can.’ Because this principle is usually taken as an analytic (e.g. Zimmerman 1996) or a conceptual (e.g. Vranas 2007) entailment, it is supposed to follow necessarily from the concepts expressed by the words ‘ought’ and ‘can.’ Our results show that it does not. (Chituc et al. 2016: p. 23)

In what follows, I will doubt the normative significance of these findings. First, I will outline the experimental study conducted by the research team at Duke, and then I will argue that this study does not pose a serious challenge to OIC.

⁵ Buckwalter and Turri (2015) also provide evidence that participants sometimes make judgments that do not accord with OIC, but they do not explore the cognitive underpinnings of these judgments.

Chituc et al. investigate ought judgments by experimentally manipulating blame across two vignettes involving an agent who is unable to keep a promise.⁶ Both versions of the case start the same way:

Adams promises to meet his friend Brown for lunch at noon today. It takes Adams thirty minutes to drive from his house to the place where they plan to eat lunch together.

The *low blame* version of the case ends this way:

Adams leaves his house at eleven thirty. However, fifteen minutes after leaving, Adams car breaks down unexpectedly. Because his car is not working at that time, Adams cannot meet his friend Brown at noon, as he promised.

The *high blame* version of the case ends like this:

Adams decides that he does not want to have lunch with Brown after all, so he stays at his house until eleven forty-five. Because of where he is at that time, Adams cannot meet his friend Brown at noon, as he promised. (Chituc et al. 2016: p. 21)

Following each vignette, participants were asked, “Do you agree or disagree with the following statement: At eleven forty-five, it is still true that Adams ought to meet Brown at noon.” Participants answered on a scale from -50 (completely disagree) to 50 (completely agree), with 0 being “neither agree nor disagree.”

What the research team found is that participants were more likely to say that an agent ought to keep a promise in the *high blame* condition than in the *low blame* condition.⁷ On the whole, 60% of subjects in the *high blame* condition gave answers above the midpoint, but in the *low blame* condition only 31% of subjects gave answers above the midpoint. This suggests that folk moral judgments do not conform to OIC. Much of the appeal of OIC, however, is supposed to be that it reflects commonsense moral judgments. Yet these findings suggest there is no conceptual entailment from ‘ought’ to ‘can’. Thus, the research team concludes that this widely assumed philosophical principle is probably false—or at least not obviously true. It seems many people think our moral obligations *can* exceed our abilities.

Do these findings pose a serious challenge to the principle that ‘ought’ implies ‘can’? For these results to have such normative significance, we must presuppose a conception of philosophy according to which the intuitive judgments tested in this study play an important argumentative role in supporting OIC.⁸ However, I believe these intuitions play no such role. The relevant judgments must avoid conceptual pitfalls, but the participants in this study are making a conceptual confusion that has gone unnoticed. Let me explain.

⁶ These cases are adapted from Sinnott-Armstrong (1984).

⁷ This was a within-subjects study, and the order in which participants read the two scenarios didn’t affect their responses.

⁸ The authors of this study do not use the word ‘intuition’ to describe their data, but as I’ve defined this notion it is clearly what they have in mind.

Following Schroer and Schroer (2013), I will say an intuition is ‘muddled’ when it is misdirected or about a notion other than the one under discussion. Not all intuitions are muddled, but if an intuition is muddled, this is enough to disqualify its purported philosophical significance. This is because the intuition would be about a notion other than the philosophically relevant one under discussion. In what follows, I will argue the intuitions elicited by Chituc et al. are likely muddled because they are about a notion of ‘ought’ that differs from the philosophically relevant one they were looking to test.

As teachers, we often find that our philosophically untutored students have muddled intuitions. To borrow an example from Schroer and Schroer (2013: pp. 1267–1268), imagine a professor preparing her opening lecture on the problem of free will. At some point during the lecture, she plans to elicit her class’ intuitions on the question of whether moral responsibility is compatible with determinism. A challenge facing this teacher is not just about how to elicit an initial intuition from her students; she must also ensure that her students have avoided various conceptual pitfalls. For example, are her students aware that fatalism is not the same thing as determinism? If not, then the intuitions she elicits might be about the compatibility between fatalism and moral responsibility. Are her students aware that we can still justify imprisoning dangerous individuals as a matter of social defense even if these people are not, strictly speaking, morally responsible? If not, then the intuitions she elicits might be about the compatibility between determinism and a specific account of punishment, not moral responsibility. It is easy for students to fall victim to these conceptual confusions. As educators, part of our role is to ‘unmuddle’ our students by clearing up these misunderstandings and conflation in their understanding of the relevant concepts.

Upon reading Chituc et al.’s study, I wondered if people’s intuitions about the relevant ‘ought’ statements were similarly muddled. More specifically, I hypothesized that participants were more likely to say that an agent ought to keep a promise in the *high blame* condition because they believed that Adams ought *to have* met Brown at noon. In other words, they were assessing the counterfactual situation. If this were correct, then participants were inclined to agree with the statement “At eleven forty-five, it is still true that Adams ought to meet Brown at noon” because they were reporting their belief that it was still true that Adams ought to have met Brown at noon. This judgment, however, is perfectly compatible with OIC.

To investigate whether people are liable to make this conceptual confusion, I first surveyed two classrooms of philosophy students at my university (approximately 50 students in total). These students were presented with the same two vignettes used in Chituc et al.’s study, and then asked: “Do you agree or disagree with the following statement: At eleven forty-five, it is still true that Adams ought to meet Brown at noon.” They each reported their answers on a piece of paper that included a scale from –50 (completely disagree) to 50 (completely agree), with 0 being “neither agree nor disagree.” What I found was similar to the results reported in Chituc et al.’s study: students tended to agree with the ‘ought’ attribution in the *high blame* version but not the *low blame* version.

After collecting their responses, the students were invited to explain why they agreed or disagreed with the relevant statement in each version of the case. Interestingly, many students who agreed with the ‘ought’ attribution in *high blame* went on to clarify their response in roughly the way I hypothesized: they said they agreed with the ‘ought’

attribution because, as one student put it, “Adams shouldn’t have stayed at his house so late—he should have left in time to meet his friend at noon”. This student was clearly responding to a situation type other than the one Chituc et al. were testing. The task was *not* to evaluate a counterfactual scenario, but rather to answer whether, at 11:45 am, it is still true that Adams ought to meet Brown at noon.

As the discussion continued, many students who initially agreed with the ‘ought’ attribution in *high blame* also agreed that, strictly speaking, it is *not* the case that Adams ought to meet Brown at noon, since it is impossible for him to do so. In other words, several people who initially *seemed* to reject the entailment from ‘ought’ to ‘can’ did not in fact reject it. Rather, they interpreted the ‘ought’ statement in a way the original study was not intended to test: they agreed with the ‘ought’ attribution because they were reporting their belief that Adams *ought to have* met Brown at noon. Further deliberation allowed them to disambiguate these two readings and then competently answer the question asked.⁹

But this is mere anecdote. To further explore the possibility that some participants in Chituc et al.’s study were providing muddled responses, I recruited 95 participants (U.S. residents; aged 18–74 years, mean age range = 35–44 years; 43 female) through Amazon Mechanical Turk to participate in an experimental study that was designed to disambiguate different meanings. Participants were compensated \$0.20 for approximately 2 min of their time, and repeat participation was prevented. Each participant was given the same *high blame* case used by the research team at Duke, and (using the same scale) participants were asked to rate whether they agreed or disagreed with the statement: “At eleven forty-five, it is still true that Adams ought to meet Brown at noon.” Once their responses were submitted, participants were taken to a second screen and asked to report which of the following sentences best described their opinion:

- (A) At eleven forty-five, it is still true that Adams ought to meet Brown at noon. It is impossible for him to do this (because it would now take too long to get there on time), but he is obligated to do what is now impossible.
- (B) At eleven forty-five, it is true that Adams ought to have met Brown at noon, even though it is now impossible for him to do this (because it would now take too long to get there on time). This is because Adams should have left earlier.
- (C) Neither (A) nor (B) adequately describes my opinion.

Options (A) and (B) were written to clarify the relevant ambiguity. If a participant were to select (A), this would constitute a clear endorsement of the idea that our obligations can exceed our abilities, thereby confirming Chituc et al.’s claim that ‘ought’ does not imply ‘can’. If a participant were to select option (B), this would indicate he or she was assessing a counterfactual situation that is not incompatible with OIC. Option (C) was

⁹ Admittedly, it is difficult to distinguish among students (a) merely clarifying an earlier statement, (b) revising their opinion of an earlier statement, and (c) changing what they say, without changing their opinion, in response to social cues. In a weak attempt to screen off (b) and (c), I chose to start the discussion by asking students to clarify their own responses, rather than by trying to interpret their responses myself and asking if they agreed. After several students indicated they meant (something like) “Adams *ought to have* been there at noon”, I tried to clarify this view and, once I did, several students who made the ‘ought’ attribution in *high blame* said (something like) “That’s exactly what I meant”. People can of course be mistaken about what they thought they meant, but I think we should take at least some of these students at their word.

included for participants who were not satisfied with either (A) or (B). Participants who selected (C) were given the opportunity to explain their answer.

Of the 93 participants included in this analysis,¹⁰ 66 people (70%) agreed with the statement: “At eleven forty-five, it is still true that Adams ought to meet Brown at noon” ($M = 12.19$, $SD = 33.40$), $t(92) = 3.52$, $p < 0.001$. This figure is slightly higher than that reported by Chituc et al., who found that 60% of subjects in the *high blame* condition gave answers above the midpoint on the agreement scale. However, the crucial follow-up question reveals something interesting: 23 of the 66 people who agreed with the initial ‘ought’ attribution selected option (A), whereas 38 of these 66 participants selected option (B). In other words, only 34% of the participants who made the initial ‘ought’ attribution were willing to select the option that clearly violated OIC, while 58% of this group indicated they were assessing a counterfactual situation that is compatible with OIC.¹¹ This was a significant difference, $X^2(1, N = 66) = 6.807$, $p < 0.01$. In the full group of 93 participants, only 23 (25%) selected the option that violated OIC, whereas 64 (69%) selected the option compatible with OIC.¹²

This suggests the responses collected by the research team at Duke were largely a result of ambiguity. The majority of people who seems to reject the principle that ‘ought’ implies ‘can’ *did not in fact reject it*. Rather, they clarified their initial response by selecting the option that is perfectly compatible with OIC: option (B). Thus, the data reported by Chituc et al. fail support the claim that folk moral judgments do not generally conform to the widely assumed OIC principle.¹³ (That said, a substantial minority (25%) were still willing to say an agent can be obligated to do the impossible, which many philosophers will find surprising.)

The worry that participants are filling in the details differently in the two statements is really nothing new. Sosa (2007) says these type of experimental results really concern in the first instance only people’s responses to certain words, and so these surveys might be revealing verbal disagreement but not any substantive disagreement.¹⁴ Unlike Sosa, however, I do not merely point to the bare possibility that participants who offer different responses to survey questions might be part of a verbal dispute. Rather, I’ve tried to show that participants in Chituc et al.’s study were interpreting the relevant question differently from how the experimenters intended, and so the data only reveal a disagreement in questions, not answers.

¹⁰ Two participants were excluded for failing a comprehension check: they both agreed and disagreed with the same statement across two questions.

¹¹ Five participants (8%) who agreed with the initial ‘ought’ attribution picked option (C), indicating that neither option adequately described their opinion.

¹² Thank you to Mike Stuart for analyzing the data in this section.

¹³ To their credit, the research team was sensitive to the possibility of ambiguity: that’s part of the motivation for how they designed their second experiment, which tested whether their findings were distorted by blame validation. However, this follow-up study was not designed to rule out the type of ambiguity I am positing.

¹⁴ Recognition that these intuitive disagreements might be verbal can be found even within the experimentalist movement itself (e.g. Nichols and Ulatowski 2007).

3 Study two: philosophical skepticism

In one of the earliest papers in contemporary experimental philosophy, Nichols et al. (2003; reprinted 2012) argue that one of the most popular and influential forms of skepticism is actually much less interesting and much less worrisome than philosophers have thought. More specifically, they present evidence that many—perhaps most—people do not have the intuitions that are required to motivate the skeptic’s argument.

Nichols, Stich, and Weinberg (henceforth NSW) focus on *Cartesian* skeptical arguments. These arguments are characterized by their use of ‘skeptical hypotheses’, which describe purportedly undetectable, cognitively debilitating states such as dreaming, hallucination, and victimization by the Cartesian demon. To illustrate, consider the following brain-in-a-vat version of the skeptic’s argument:

1. I don’t know that I am not a handless brain in a vat.
2. If I don’t know that I am not a handless brain in a vat, then I don’t know that I have hands.
3. Therefore, I don’t know that I have hands.

This argument has played an important role in the work of some leading 20th Century philosophers, and many contemporary epistemologists provide essentially the same formulation of skepticism.¹⁵ Suitably articulated, this argument will lead us to deny much of our putative knowledge of the world around us.

Presumably, skepticism is enduring and significant because many people find the skeptic’s reasoning persuasive. In the *Meditations*, for instance, Descartes appeals to the ease with which we acknowledge that the Meditator must know that he is not dreaming if he is to know that he is sitting by the fire. We empathize with the Meditator and recognize that we share his epistemic disadvantage. Skepticism gets its teeth because our willingness to ascribe knowledge allegedly decreases when skeptical considerations are raised.

Let’s call any intuitive judgment according to which one does not know the falsity of a skeptical hypothesis a ‘skeptical intuition’. For example, you have a skeptical intuition if, after considering the possibility that you are a brain in a vat, it seems to you that you don’t know that you’re not a brain in a vat. According to NSW, the skeptic’s argument relies crucially on skeptical intuitions, and these intuitions must be widely shared. If many people did not have skeptical intuitions, then skepticism would be “much less *interesting* and much less *worrisome* than philosophers have taken it to be” (NSW 2012: p. 224).

NSW present evidence that skeptical intuitions are “*far* from universal” (2012: p. 224). Using several vignettes, NSW tested whether or not experimental subjects were willing to ascribe knowledge to someone who has been confronted with a skeptical hypothesis. One of their cases involves two college roommates, neither of whom is a brain in a vat, discussing the possibility that they are brains in vats. One of the roommates, named George, appeals to certain perceptions he has in order to justify his belief that he is not a brain in a vat. The participants in this study were asked whether George “really knows” or “only believes” that he is not a brain in a vat. NSW report

¹⁵ See DeRose (1999: p. 2), Cohen (1999: p. 62), and Schiffer (1996: p. 317).

that 80% of those participants with more philosophical training denied knowledge (“only believes”), whereas only 45% of those with less philosophical training denied knowledge (2012: p. 241). This indicates that people with less exposure to philosophy are more likely to claim that George knows he is not a brain in a vat; thus, they are more willing to reject the first premise of the skeptic’s argument (outlined above). In contrast, people with more philosophical training are more likely to have skeptical intuitions.

Another vignette used to test for skeptical intuitions was based on a scenario described by [Dretske \(1970\)](#):

Pat is at the zoo with his son, and when they come to the zebra cage, Pat points to the animal and says, “that’s a zebra.” Pat is right — it is a zebra. However, given the distance the spectators are from the cage, Pat would not be able to tell the difference between a real zebra and a mule that is cleverly disguised to look like a zebra. And if the animal had really been a cleverly disguised mule, Pat still would have thought that it was a zebra. Does Pat really know that the animal is a zebra, or does he only believe that it is? (NSW 2012: p. 236)

Although the majority participants maintained that Pat “only believes” the animal is a zebra, people with a lower socio-economic status (measured by level of education) were significantly more likely to say that Pat “really knows” (31%) than people with a higher socio-economic status (11%). This suggests that skeptical intuitions vary across different socio-economic groups: those with a lower socio-economic status are more willing to deny the skeptic’s conclusion than people with a higher socio-economic status.

In a cross-cultural study, NSW presented their participants with another variant of Dretske’s zebra case. They found a significant difference between Western and Indian Subcontinent participants. More precisely, only 32% of Westerners said the zoogoer “really knows” the animal is a zebra, whereas 50% of participants from the Indian Subcontinent were willing to ascribe knowledge to the zoogoer.

This data suggests that whether or not people have skeptical intuitions partly depends on factors such as level of philosophical training, socio-economic status, and cultural background. This is bad news for skepticism, according to NSW, because the skeptic’s argument relies crucially on skeptical intuitions. If many people do not have skeptical intuitions, then many will not find the skeptic’s argument plausible. Consequently, skepticism would be much less interesting and much less worrisome than philosophers have claimed.

However, I think we should reject this conclusion on the grounds that skepticism does *not* rely on the intuitions tested by NSW.¹⁶

Imagine you are teaching Epistemology 101 and are having a class discussion about Descartes’s *Meditations*. You present the class with Descartes’s evil demon scenario, which hypothesizes the existence of a demon as clever and deceitful as he is powerful, who has directed his entire effort to misleading us. You draw your students’ attention

¹⁶ This section provides a condensed version of the argument made in ([Hannon 2017](#)), but it also makes an important revision. In my earlier work, I contrasted ‘surface intuitions’ with ‘reflective intuitions’. I now think this is a mistake. The relevant reflective judgments need not be intuitive at all, so I now contrast intuitions with reflective judgments.

to fact that this demon presents a complete illusion of an external world to the senses, when in reality no such external world exists. This includes the illusion of one's own body, such as the sensation that one has hands. After describing this scenario, you suggest the following striking conclusion: nobody knows they have hands.

In my experience, a sufficiently large classroom of undergraduates will rarely meet this skeptical argument with unanimous approval. Drawing on NSW's data, we may reasonably assume that a non-trivial number of students will not find the skeptic's conclusion immediately compelling. Of those students who do not find it compelling, imagine one of them exclaims, "Come on, Professor! I know that I am not deceived by an evil demon." According to NSW, this student does not have a skeptical intuition. Does this mean the skeptic's argument relies on a dubious assumption (for this student)?

Presumably, you will try to get this student to see the force of the skeptic's argument. You will probably do this by engaging in Socratic questioning; for instance, you could ask: "How do you know that you aren't deceived? What evidence could you provide us with? Wouldn't you think you weren't deceived even if you in fact were?" And so forth. With a little extra work, many resilient students eventually—and perhaps grudgingly—concede they don't really know that they're not deceived (or at least they come to appreciate the difficulty in maintaining knowledge in a non-dogmatic way).

For how long this type of questioning must continue will depend on a combination of the student's resilience and the inquirer's skillfulness. But it is plausible, and certainly true in my experience, that after a sufficient amount of Socratic questioning, skepticism will start to look genuinely puzzling—or at least not obviously false – to many people who did not initially have a skeptical intuition. Of course, many students will not be led to *endorse* skepticism (few of us do), but the skeptic's argument will now seem credible, puzzling, and philosophically significant.¹⁷

The Epistemology 101 example helps illustrate the central flaw in NSW's meta-skeptical argument. According to NSW, many people are willing to ascribe knowledge even when faced with a skeptical hypothesis. They take this to show that skepticism is much less interesting and worrisome than philosophers have thought. However, *the mere fact that some people are not immediately perturbed by skeptical considerations does nothing to undermine skepticism*. Why not? It is because the plausibility of skepticism doesn't crucially depend on people's immediate, intuitive judgments about hypothetical cases. Reflective judgments are what matter, not intuitions.¹⁸

While the distinction between 'intuitive' and 'reflective' judgments is rough, the force of skepticism clearly depends on more than just one's immediate reactions to scenarios involving a skeptical hypothesis. When a philosopher says competent speakers would use 'knows' in certain ways, she is not necessarily predicting that such speakers will respond in these ways without adequate reflection. The fact that a student in your

¹⁷ This is an empirical claim that must be tested. Later I will argue that the simple survey method used by NSW (and many other experimental philosophers) is an inadequate tool to test this claim.

¹⁸ This claim is similar to an idea that has been defended elsewhere, namely, that the intuitions being tested by the experimentalists are not the philosophically relevant ones (see Williamson 2005; Kornblith 2007; Ludwig 2007; Kauppinen 2007; Deutsch 2010). In Sect. 4, however, I will argue that my view is distinct from these other proposals in several important ways.

Epistemology 101 class is not immediately perturbed by a skeptical hypothesis does not threaten the significance of skepticism because no premise in the skeptic's argument relies on one's having a skeptical intuition. What makes skepticism worrying is that many people who are prodded to think reflectively about knowledge tend to find the skeptic's argument reasonable and difficult to refute. The driving force behind skepticism is our tendency to find the premises in the skeptic's argument *plausible*, but whether we find them *immediately* plausible (i.e. intuitive) is inessential.¹⁹

Interestingly, NSW seem to concede this very point in the first paragraph of their article. When discussing skepticism's significance in the history of philosophy, they claim the intuitions undergirding skeptical arguments are supposed to be "shared by everyone (or almost everyone) who thinks reflectively about knowledge" (2012: p. 224). If this is true, as I think it is, then the intuitions elicited by their experimental study are not part of the evidence for skepticism.

The threat of skepticism is not diminished simply because some students in Epistemology 101 do not find the skeptic's argument immediately tempting. Why, then, would the fact that some people are not immediately willing to deny knowledge in the context of a survey indicate that skepticism is less significant than philosophers have thought? We should not follow NSW in thinking it does. The skeptic's argument does not crucially rely on the intuitions tested by NSW. All that matters is that people can be led (e.g., via Socratic questioning) to feel the pull of the skeptic's reasoning.²⁰

This of course raises the question of whether or not people *can* generally be led to deny knowledge as a result of Socratic questioning. For simplicity, I will use the label 'Skeptical Anxiety Hypothesis' to refer to the hypothesis that people can generally be led to deny knowledge by reflecting on skeptical considerations. How confident should we be that the Skeptical Anxiety Hypothesis is correct?

At present this is an open empirical question. Earlier I drew on my experience teaching Epistemology 101 to support this hypothesis, but this is just anecdotal. Can anything else be said in favor of the Skeptical Anxiety Hypothesis? John Turri points to the wide success gained by certain sorts of sci-fi films and books that importantly turn on inducing skeptical anxiety. He writes,

Skepticism readily captures the public imagination and is a staple of popular culture... Hollywood grows rich on it. Witness the many successful films that

¹⁹ In an earlier study, NSW consider two versions of this objection (Weinberg et al. 2001; reprinted 2008: p. 39). The first version is this: they are looking at the wrong sorts of intuitions, since the right sort require at least a modicum of reflection. In reply, NSW maintain their participants did reflect "at least minimally", as evidenced by the fact that some participants wrote brief explanatory comments after their answers. But this reply is not persuasive for two reasons: first, there is no evidence that the subjects in their 2003 study on skepticism were provided with a similar opportunity to explain their answers; and second, even if subjects were given this opportunity to reflect, more than a modicum of reflection may be necessary for some people to find skepticism intuitive. NSW then consider a second version of this objection: that the right sorts of intuitions are those that emerge after an extensive period of discussion and reflection, which is the sort philosophy typically encourages. In response, they doubt that people's reflective judgments will differ from their intuitions. However, this view cannot be convincingly maintained without argument.

²⁰ I am assuming these people can be led to feel skeptical anxiety on a rational basis (rather than, say, the use of rhetoric or a drug). If most people can be led to feel the force of skepticism in this way, it would suggest there is some shared core to our epistemic judgments and that skepticism is a part of it.

put skeptical doubts front and center. Foremost among these are *The Matrix* (i.e., Do you know that you're not in The Matrix?), *Vanilla Sky* (i.e., Do you know that your life is not just a dream?), *The Truman Show* (i.e., Do you know that your life isn't just one big charade?), *Bladerunner* (i.e., Do you know that this individual isn't a replicant?), *Dark City* and *Memento* (i.e., Do you know that "your past" is what it seems to be?). If nothing else, skillful trafficking in skeptical doubt is a lucrative business. (Turri 2015: p. 308)

These films might be less popular if audience members weren't led to feel some level of skeptical anxiety as a result of skeptical possibilities. Admittedly, this evidence isn't very conclusive. However, Turri also argues that certain psychological factors render us vulnerable to skeptical doubts. Specifically, his experiments show that the skeptic will likely instill doubt in us if she gets us to focus on inferential belief (rather than perceptual belief) with negative content (e.g., "You *don't* know that you're not a brain in a vat").

More data is needed to establish the persistent and widespread appeal of skepticism. For the purpose of my argument, however, no such evidence is needed. To block NSW's conclusion, all I must show is that their data has little to no bearing on the question of skepticism's significance. I have already argued for this.

Still, an interesting philosophical question now arises: do less reflective and more reflective judgments behave in pretty much the same ways? If they are produced by the same underlying cognitive mechanism (see Wysocki 2016), then evidence about our intuitive judgments of the sort yielded by experimental philosophy surveys might provide good (albeit indirect) evidence about the content of our more reflective judgments. I plan to explore this intriguing question in future work. In lieu of carrying out this daunting project here, I'll simply assert that anyone who thinks intuitive and reflective judgments are symmetrical in this way carries an empirical burden of showing this symmetry to be true.

4 Intuitions, reflective judgments, and experimental philosophy

In the background of these studies is a conception of philosophy according to which intuitions play important argumentative roles. Some critics of experimental philosophy argue that intuitions play no such roles,²¹ but I do not share their radical skepticism (I'll elaborate on this point below). However, I do think that many *particular* studies, such as those on which I'm focusing, are in the unhappy position of relying on intuitive judgments that are philosophically irrelevant.

So how do we determine which judgments are philosophically relevant? Some scholars claim the relevant judgments must be grounded in adequate reflection.²² Let's call this the 'reflective judgment defense'. Ludwig, for example, says: "We should not expect that in every case in which we are called on to make a judgment we are at the outset equipped to make correct judgments without adequate reflection" (2007: p.

²¹ Williamson (2007), Deutsch (2010), and Cappelen (2012).

²² See Kauppinen (2007) and Ludwig (2007).

149). [Kauppinen \(2007\)](#) largely concurs, and I also find this idea *prima facie* plausible. It seems clear, for instance, that even if our immediate, intuitive judgment is that the protagonist in a skeptical scenario has knowledge, this is not sufficient to block the skeptic's argument (as argued in the previous section).

However, the reflective judgment defense faces two related obstacles. First, it is difficult to define the difference between intuitive and reflective judgments so they can be measured or expressed quantitatively.²³ Second, the extent to which reflection is a precondition for philosophically relevant judgments will largely depend on how "reflection" and "intuition" are characterized. In Sect. 2, for example, I claimed that the data reported by [Chituc et al. \(2016\)](#) are philosophically irrelevant because their survey respondents were prone to a conceptual confusion. (In particular, participants associated different conceptual contents with certain key terms in the description, and so many participants were not answering the intended question.) To resolve this misunderstanding, I proposed a simple and natural technique for effectively guiding participants to competently assess the cases. However, it is not obvious that the 'ought' judgments elicited by my survey were a result of careful, nuanced, conceptually rigorous reflection on the part of the participant. While my experiment was designed to clarify how participants understood the relevant question, it did little to induce reflection in any strong sense.

Thus, contrary to what [Kauppinen \(2007\)](#) and [Ludwig \(2007\)](#) suggest, not all judgments must be reflective to qualify as philosophically relevant. Still, the relevant judgments must be sufficiently 'robust'. That is, they must exemplify competence, avoid certain conceptual pitfalls, and not involve performance errors. Further, we often think that reflection can lead to clarification, and without clarification people tend to interpret cases in a variety of unintended ways.²⁴ Timothy Williamson, for instance, claims that philosophers are able to "apply general concepts to specific examples with careful attention to the relevant subtleties" ([Williamson 2007](#): p. 191). In this way, reflection (and philosophical training) might make one better at noticing important but subtle details, which might explain differences in performance between philosophers and non-philosophers.²⁵

It is a well-known objection that experimental philosophers are not studying the relevant judgments. [Kauppinen \(2007\)](#), for instance, argues that when philosophers talk about 'intuitions', they are not actually talking about immediate responses that can be collected via surveys. Rather, he claims that philosophers are interested in mental states that have been subjected to critical examination as part of reflective par-

²³ [Weinberg et al. \(2013\)](#) use two measures to operationalize the difference between intuitions and reflective judgments: need for cognition ([Cacioppo and Petty 1982](#)) and cognitive reflection ([Frederick 2005](#)). However, they acknowledge there might be better ways to operationalize this distinction, and [Kauppinen \(2007\)](#) doubts that we can adequately operationalize these notions without giving up the non-participatory social scientific method used by experimental philosophers. Kauppinen says the only way to make sure that a particular response genuinely reflects the respondent's concept is to abandon experimental philosophy and the engage in dialogue with the subject.

²⁴ However, some people doubt that reflection improves our reliability (see [Kornblith 2012](#)).

²⁵ Two studies that indicate a significant difference between lay people and professional philosophers are [Sytsma and Machery \(2010\)](#) and [Horvath and Wiegmann \(2016\)](#). However, some studies do not report any difference between philosophers and lay people, such as [Schwitzgebel and Cushman \(2012\)](#).

ticipation in traditional philosophical discourse. Similarly, Ludwig (2010) claims that philosophers engaged in conceptual analysis base their analyses on how reflective, competent speakers are using a philosophically relevant term, not on people's immediate, unreflective responses. He therefore concludes that the data collected by many experimentalists are irrelevant to philosophical debates.²⁶

While my argument might resemble the views expressed by some critics of experimental philosophy, it departs from them in several important ways. First, my argument does not depend on a conception of experimental philosophy that is unjustifiably narrow. As I'll explain below, Kauppinen and other critics assume that experimental philosophers must limit themselves to studying the intuitive judgments we have been discussing throughout this paper, but this is simply not true. Second, I do not claim that philosophically relevant judgments reflect *expertise* (see Nado 2014 for an overview). Ordinary folk are required to grasp the relevant aspects of the hypothetical case considered, but as long as they understand the nature of the task, they do not require any special expertise.²⁷ Third, Kauppinen and other critics mistakenly think experimental philosophy does not have the resources to get at the philosophically relevant judgments, but I will argue that even reflective judgments do not lie forever beyond the reach of experimental philosophy. I'll now elaborate on each of these points.

Critics of experimental philosophy have made sweeping claims about the value of this research program. Kauppinen (2007), for example, argues that recent experimental work has contributed nothing to philosophy, and he believes this movement cannot resolve the central questions posed by 'the analytic project'. Cappelen (2014) thinks experimental philosophy is a completely misconceived endeavor. Deutsch (2015) leaves room for experimental philosophy to study people's beliefs, but he nevertheless takes this movement to be of no help in solving purely philosophical problems.

These objections presuppose a narrow conception of experimental philosophy by restricting the movement to the study of folk intuitions. As I've suggested, however, experimental philosophy has diverse ambitions. Not all its practitioners are engaged in the project of using surveys to investigate folk intuitions. Other types of empirical data that are relevant to philosophy include work on cognitive representations of causal structures in the world (e.g. Bonawitz et al. 2010; Danks 2007), behavioral lab experiments (e.g., Bloom 2013), and brain imaging research (e.g., Soon et al. 2008). Experimental data can also shed light on the internal psychological processes that lead people to have the intuitions they do, thereby increasing our understanding of how our minds work when we think about philosophical issues (e.g., Knobe 2003). Given this variability, we should resist the urge to draw conclusions about the merits or deficiencies of experimental philosophy as a unitary whole.

I doubt critics will resist this broader characterization of experimental philosophy. Still, they could easily modify their criticism in the following way. They might admit

²⁶ Likewise, Cappelen (2012) argues that philosophers do not rely upon intuitions as evidence when theorizing, though they may occasionally slip up with talk of 'intuitions'.

²⁷ One might define 'expertise' as the having ability to "grasp the relevant aspects of the hypothetical case considered" (Grundmann 2010: p. 500), but that would make the need for expertise a fairly trivial requirement.

that their objections only target a narrow conception of experimental philosophy and yet argue that much experimental work centers on the practice of doing surveys to test the intuitions of non-philosophers. As illustrated by the two experimental studies discussed above, empirical data is often collected against the background of a conception of philosophy according to which folk intuitions play important argumentative roles. If we reject this conception of philosophy, however, then the value of much work in experimental philosophy will be undermined.

I take this complaint to be at the heart of the arguments put forward by critics like Kauppinen (2007), Ludwig (2007), Williamson (2007), Cappelen (2012), and Deutsch (2015). While I have some sympathy with this view, I nevertheless believe it is mistaken. Even if we restrict our investigation of experimental philosophy to the study of folk intuitions, we should not completely reject the idea that intuitions play important argumentative roles in philosophy. I defend a moderate view: the philosophical relevance of folk intuitions varies from topic to topic.

Consider, for example, Gettier's (1963) challenge to the justified true belief account of knowledge. In his article, Gettier describes two cases in which a person seems to have a justified true belief that does not amount to knowledge. Upon reading these cases, Gettier expected his readers to conclude that a person could have a justified true belief without possessing knowledge. In other words, the intuitions elicited from Gettier's famous cases were treated as evidence against the justified true belief theory of knowledge.²⁸ If nobody found it intuitive to deny knowledge to the characters in his examples (or similar ones), then the argument against knowledge as justified true belief would have been weakened.²⁹

As a second example, consider intuitions about the footbridge version of a trolley case (Thomson 1985). We are supposed to just *see* that it is impermissible to push a fat man off a footbridge to stop a train bearing down on five people, killing the fat man but saving five equally innocent others. This intuitive judgment is supposed to count as defeasible evidence against the utilitarian claim that we ought to sacrifice one life to save five. So if this intuition were to show wide variability, the case against utilitarianism might be weakened.

These considerations suggest that intuitions *do* play an important role in some philosophical debates, contrary to what critics of experimental philosophy often claim. This creates room for intuition-driven surveys to get traction. Philosophical practice sometimes involves an appeal to intuitions, so the survey methodology can (and should) be used to examine this evidence. Such empirical work can both challenge and corrob-

²⁸ Williamson claims that Gettier's argument, when properly understood, does not rely on intuitions (2007: pp. 184–186). For Williamson, our evidence that knowledge isn't justified true belief doesn't consist of our intuitions that it is possible for someone to have a justified true belief that *p* without knowing that *p*. Rather, it consists of the fact that it is possible for someone to have a justified true belief that *p* without knowing that *p*. This does not appeal to intuitions as evidence but rather to facts about the world. However, the 'thin' conception of intuitions I'm working with leaves room for intuitions to be simple counterfactual judgments about contingent matters of fact, as Williamson contends. Further, Alexander (2010: pp. 382–383) convincingly argues that even if Gettier's argument can be reformulated to not mention psychological facts about us (as Deutsch claims), this argument must still indirectly appeal to intuitions as evidence.

²⁹ Some data suggests that Gettier intuitions display more variability than philosophers have thought (Weinberg et al. 2001), but this data has been heavily criticized (Machery et al. 2015).

orate the claims of traditional philosophy: sometimes the results will undermine the predictions made by armchair philosophers about what ‘we’ find intuitive; sometimes they will empirically verify the claims made by practitioners of armchair philosophy. To the extent that intuitions are treated as evidence in philosophical theorizing, philosophy will be hostage to the data gathered by experimental surveys.

Further, my argument does not depend on the contentious claim that ‘expert’ intuitions are what correctly come into play in the practice of doing philosophy. By taking a stand on the issue of what counts as expert usage of a philosophical concept, these critics run into a variety of objections. For example, [Weinberg et al. \(2010\)](#) provide empirical evidence that challenges the claim that training in philosophy increases the reliability of one’s intuitions. Similarly, Nichols et al. (Reprinted 2012: p. 242) argue there is no more reason to think one set of intuitions tracks the truth any better than another set. [Weinberg and Alexander \(2014\)](#), [Alexander \(2012\)](#), [Kornblith \(2002\)](#), and [Schwitzgebel and Cushman \(2012\)](#) have made similar criticisms. Further, [Machery et al. \(2004\)](#) and [Knobe and Nichols \(2008\)](#) argue that expert intuitions might be *more biased*. Proponents of the ‘expertise’ defense have struggled to provide a non-question-begging account for why philosophers’ intuitions are more reliable than the pre-theoretical intuitions of laypersons.

In contrast, my view throws into question the evidential value of some experimental studies without needing to assume the far more contentious (and difficult to establish) claim that philosophically relevant judgments are the product of expertise. A muddled intuition is enough to disqualify its purported significance, but being unmuddled does not, by itself, establish the reliability of one’s judgments. A philosopher might be unmuddled and yet biased, for example. In fact, it is perfectly consistent with my argument that reflective judgments—whether skeptic-friendly ones or those in favor of OIC—are defective or imperfect.

So far I’ve discussed three ways to distinguish my view from similar criticisms. First, many critics presuppose a narrow conception of experimental philosophy. Second, these critics maintain that intuitions are never (or ought never be) treated as evidence in philosophy. Third, they sometimes argue that philosophically relevant judgments must reflect expertise. Each of these claims is dubious. Now I want to distinguish my proposal in a fourth and final way. Specifically, I will challenge the claim that experimental philosophers do not have the resources to experimentally test the type of reflective judgments that often matter in philosophy.

Kauppinen argues that the only philosophically relevant judgments are those elicited through a participatory method he calls ‘dialogue and reflection’ (2007: p. 98). To create the conditions needed to elicit reflective judgments, he asks us to imagine a researcher going through a survey participant’s answers together with her, asking for the reasons why she answered one way rather than another, making sure she did correctly understand the scenario described, varying examples, teasing out implications, pointing out similarities and disanalogies with other cases, and trying to get the subject to reflect on whether her initial response really is what she wants to say. Kauppinen argues that anyone engaged in this sort of process is no longer doing experimental philosophy (2007: p. 106). Rather, one is returning to the Socratic method of engaging in serious dialogue with test subjects.

Will any amount of experimental ingenuity allow us to design a study to test people's reflective judgments? Kauppinen thinks experimental philosophy simply does not have the resources to get at the philosophically interesting judgments, so he concludes that philosophers must retreat to the armchair in defeat.

This conclusion is overly pessimistic. While the simple survey method is a doubtful tool for eliciting reflective judgments, we should not assume there are no other methodological tools available that could conceivably be used in future studies. Even if getting at these judgments requires more dialogue and reflection on the part of participants, this method can be implemented in one of two ways: either informally or in a controlled and systematic manner. In the informal way, it suffices for the purposes of doing conceptual analysis that philosophers engage in discussions with their colleagues, students, friends, and family. The problem with this approach is that it runs into familiar objections that experimentalists have made against armchair philosophy (see [Nadelhoffer and Nahmias 2007](#): p. 131). In particular, one might worry about various sample biases and experimenter biases.³⁰

But nothing rules out the possibility of setting up controlled and systematic experiments to find out what people's reflective judgments about a given topic really are. As far as I know, [Nadelhoffer and Nahmias \(2007\)](#) were the first to argue that experimental philosophers can implement the method of dialogue and reflection in this way. As they point out, a benefit of this approach is that we could collect data much faster than we could by relying exclusively on informal conversations. It would also allow us to implement the model of dialogue and reflection in a way that is less methodologically problematic than the informal method Kauppinen advocates. However, Nadelhoffer and Nahmias only highlight the possibility of studying reflective judgments in a controlled and systematic way—they do not explain how we might implement this proposal. Let's now consider some ways to implement the method of dialogue and reflection more formally and systematically.

One way to proceed would be to stage the right sort of Socratic intervention, record it, and play it for the subjects, then ask the subjects to evaluate the participants in the dialogue. To test reflective judgments about knowledge, for example, the subjects might listen to a conversation between two interlocutors, Claudia and Brett. Claudia might try to convince Brett that neither of them knows that they have hands, since neither of them can prove they are not a brain in a vat. At first, Brett might resist Claudia's reasoning; he might say, "But it's *so* unlikely that we are brains in vats!" Claudia might reply in the following way: "However unlikely it might be, you don't *know* that we aren't. After all, winning the lottery is highly unlikely, but you don't *know* that you'll lose." After a variety of considerations have been discussed, the dialogue could end in one of two ways: either Brett concludes that he doesn't know that he has hands (scenario 1) or he maintains that he knows that he has hands (scenario 2). Participants who are given scenario 1 would be asked, "Do you think Brett doesn't know he has hands?" Participants given scenario 2 would be asked, "Do you think Brett knows he has hands?"

³⁰ Stich and Weinberg (2001: p. 642) raise this objection against [Jackson \(1998\)](#).

The same technique could be used to elicit reflective judgments about whether ‘ought’ implies ‘can’. Participants might listen to a recorded conversation between two interlocutors, Emanuel and Sophia, who are discussing whether our moral obligations can exceed our abilities. Emanuel might try to convince Sophia that it is unreasonable to think anyone is ever obligated to do something impossible, while Sophia might argue that obligations can remain in place even when they become impossible to fulfill. The dialogue could also flag potential ambiguities and clarify the relevant concepts; for instance, Emanuel could ask Sophia to clarify whether she means that Adams ought *to have* met Brown at noon or whether she means that Adams is *still* obligated to meet Brown at noon, even though it is now impossible for him to do so.

Another way to proceed is to hypothesize at greater length as to just what it is that subjects who offer ‘surprising’ judgments might be failing to grasp, such that once they do grasp it, their judgment will change. In the case of skepticism, for example, some people might not grasp the epistemic closure principle that has been touted as a linchpin in the most compelling skeptical arguments (e.g., Dretske 1970; Stroud 1984). This principle states that, necessarily, if you know that p , and you know that if p then q , then you also know that q (or at least you are in a position to know this). Perhaps some participants do not realize that if they claim to know they have hands, then, plausibly, they must also know they are not a brain in a vat. Others might claim knowledge because they take themselves to have something close enough to knowledge such that there is no point in refraining from a knowledge attribution (see BonJour 2010: p. 73). To these people we might point out that an envatted brain is *wildly* wrong about their world, not merely basically-right-but-short-of-metaphysical-certainty. Once we generate a set of “missing links” that might bring the unworried to feel skeptical anxiety, it would be easy enough to design an experimental instrument for online delivery that would take subjects through those links.

Additional possibilities are available to empirically test my hypothesis about the ambiguity of ‘ought’ statements. For example, the experimentalist could attempt to unuddle the participant before she answers the survey. The experimenter might do this by giving the subject a miniature lecture, or by framing the vignette in a way that helps the participant avoid various confusions. This process runs into the danger of leading the witness, so one might prefer to retroactively separate subjects that are muddled from those that are not. Instead of actively trying to unuddle subjects, one could use control questions to determine whether a misunderstanding has taken place. The reports of subjects who confuse these concepts could then be disqualified.³¹ A nice feature of this second approach is that, if done correctly, it does not run the risk of prejudicing participants by ‘contaminating’ their philosophically untutored intuitions.

One could also use tasks like the Cognitive Reflection Test (Frederick 2005) to see whether untutored-but-basically-more-reflective subjects have different responses than untutored-but-basically-unreflective subjects. Indeed, some philosophers have already made good use of this tool. For example, Pinillos et al. (2011) demonstrate that people who score higher on the Cognitive Reflection Test (CRT) are less likely to display the Knobe Effect (see Knobe 2003). On this basis, they

³¹ In an earlier study, Weinberg et al. (2001) use this approach to identify subjects who treat ‘knowledge’ as synonymous with ‘subjective certainty’.

argue that reflective individuals are in a stronger epistemic position and thus their judgments are trustworthier. Assuming the CRT is a good measure of whether a person will provide a ‘surface’ rather than ‘robust’ judgment, this data provides some evidence that less reflective subjects will have different responses than more reflective subjects. It should be noted, however, that several experimentalists have used the CRT but often the hypothesized effect is not observed (see Weinberg et al. 2012, Gerken and Beebe 2014, and Colaço et al. unpublished).

5 Conclusion

There is much controversy over whether the discoveries by experimental philosophers have any philosophical significance, and the most vigorous debate centers on the philosophical significance of empirical evidence on intuitions. On the one hand, an accomplishment often touted for experimental philosophy is its potential to discredit philosophical views that have gone largely unquestioned. On the other hand, critics of experimental philosophy often claim that experimentalists cannot get at the philosophically relevant data, so the practice of philosophy can continue pretty much as it has always been done—from the armchair. Both sides of this debate have tended to overreach.

This paper casts doubt on some recent experimental work, but my overall aim is constructive. I have tried to explore concrete, tangible implications and applications that support the experimentalist’s aim to study philosophical judgments empirically. Experimental philosophers have both the philosophical arguments and the methodological tools to allay many of the concerns and criticisms that have been raised by their critics. To this end, I have considered a wider range of possibilities for testing reflective judgments, but much of the difficult empirical work remains to be done.

Acknowledgements Thank you to Elizabeth Edenberg, Mike Stuart, and two anonymous referees for very helpful feedback.

References

- Alexander, J. (2010). Is experimental philosophy philosophically significant? *Philosophical Psychology*, 23(3), 377–389.
- Alexander, J. (2012). *Experimental philosophy*. Cambridge: Polity Press.
- Alexander, J., & Weinberg, J. (2007). Analytic epistemology and experimental philosophy. *Philosophy Compass*, 2(1), 56–80.
- Bealer, G. (1998). Intuition and the autonomy of philosophy. In M. DePaul & W. Ramsey (Eds.), *Rethinking intuition: The psychology of intuition and its role in philosophical inquiry*. Lanham: Rowman and Littlefield.
- Bengson, J. (2014). How philosophers use intuition and ‘intuition’. *Philosophical Studies*, 171(3), 555–576.
- Bloom, P. (2013). *Just babies: The origins of good and evil*. New York: Crown Publishers.
- Bonawitz, E. B., Ferranti, D., Saxe, R., Gopnik, A., Meltzoff, A. N., Woodward, J., et al. (2010). Just do it? Investigating the gap between prediction and action in toddlers’ causal inferences. *Cognition*, 115, 104–117.
- BonJour, L. (2010). The myth of knowledge. *Philosophical Perspectives*, 24(1), 57–83.
- Buckwalter, W., & Turri, J. (2015). Inability and obligation in moral judgment. *PLoS ONE*, 10(8).

- Cacioppo, J., & Petty, R. (1982). The need for cognition. *Journal of Personality and Social Psychology*, 42(1), 116–131.
- Cappelen, H. (2012). *Philosophy without intuitions*. Oxford: Oxford University Press.
- Cappelen, H. (2014). Replies to Weatherson, Chalmers, Weinberg, and Bengson. *Philosophical Studies*, 171(3), 577–600.
- Chituc, V., Henne, P., Sinnott-Armstrong, W., & De Brigard, F. (2016). Blame, not ability, impacts moral 'ought' judgments for impossible actions: Toward an empirical refutation of 'ought' implies 'can'. *Cognition*, 150(C), 20–25.
- Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, 13, 57–89.
- Colaço, D., Kneer, M., Alexander, J., & Machery, E. (unpublished). On second thought: A refutation of the reflection defense. Presented at the 2016 Buffalo Experimental Philosophy Conference.
- Danks, D. (2007). Causal learning from observations and manipulations. In M. C. Lovett & P. Shah (Eds.), *Thinking with data*. New York: Lawrence Erlbaum.
- DeRose, K. (1999). Introduction. In K. DeRose & T. Warfield (Eds.), *Skepticism: A contemporary reader*. Oxford: Oxford University Press.
- Deutsch, M. (2010). Intuitions, counter-examples, and experimental philosophy. *Review of Philosophy and Psychology*, 1(3), 447–460.
- Deutsch, M. (2015). *The myth of the intuitive*. Cambridge: MIT Press.
- Dretske, F. (1970). Epistemic operators. *Journal of Philosophy*, 67, 1007–1023.
- Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19(4), 25–42.
- Gettier, E. (1963). Is justified true belief knowledge? *Analysis*, 23(6), 121–123.
- Grundmann, T. (2010). Some hope for intuitions: A reply to weinberg. *Philosophical Psychology*, 23, 481–509.
- Hannon, M. (2017). Skepticism about Meta-skepticism: meditations on experimental philosophy. *Episteme*, 14(2), 213–231.
- Horvath, J., & Wiegmann, A. (2016). Intuitive expertise and intuitions about knowledge. *Philosophical Studies*, 10, 2701–2726.
- Jackson, F. (1998). *From metaphysics to ethics: A defense of conceptual analysis*. Oxford: Oxford University Press.
- Kant, I. (1793/1996). *Religion within the boundaries of mere reason* (trans. George di Giovanni, in Immanuel Kant). *Religion and rational theology* (trans. and ed. Allen W. Wood and George di Giovanni). Cambridge: Cambridge University Press.
- Kauppinen, A. (2007). The rise and fall of experimental philosophy. *Philosophical Explorations*, 10(2), 95–118.
- Knobe, J. (2003). Intentional action and side effects in ordinary language. *Analysis*, 63(3), 190–194.
- Knobe, J. (2007). Experimental philosophy. *Philosophy Compass*, 2(1), 81–92.
- Knobe, J., & Nichols, S. (2008). An experimental philosophy manifesto. In J. Knobe & S. Nichols (Eds.), *Experimental philosophy*. Oxford: Oxford University Press.
- Knobe, J., & Nichols, S. (2013). *Experimental philosophy* (Vol. 2). Oxford: Oxford University Press.
- Kornblith, H. (2002). *Knowledge and its place in nature*. Oxford: Oxford University Press.
- Kornblith, H. (2007). Grazer philosophische studien. *Naturalism and Intuitions*, 74(1), 27–49.
- Kornblith, H. (2012). *On reflection*. Oxford: Oxford University Press.
- Lewis, D. (1983). *Philosophical papers* (Vol. I). Oxford: Oxford University Press.
- Lombrozo, T., Nichols, S., & Knobe, J. (2014). *Oxford studies in experimental philosophy*. Oxford: Oxford University Press.
- Ludwig, K. (2007). The epistemology of thought experiments: First person versus third person approaches. In P. French & H. Wettstein (Eds.), *Midwest studies in philosophy*. Oxford: Blackwell.
- Ludwig, K. (2010). Intuitions and relativity. *Philosophical Psychology*, 23(4), 427–445.
- Machery, E., Mallon, R., Nichols, S., & Stich, S. (2004). Semantics, cross-cultural style. *Cognition*, 92(3), 1–12.
- Machery, E., & O'Neill, E. (2014). *Current controversies in experimental philosophy*. Abingdon: Routledge.
- Machery, E., Stich, S., Rose, D., Chatterjee, A., Karasawa, K., Struchiner, N., Sirker, S., Usui, N., & Hashimoto, T. (2015). Gettier across cultures. *Noûs*, 50(1).
- Nadelhoffer, T., & Nahmias, E. (2007). The past and future of experimental philosophy. *Philosophical Explorations*, 10(2), 123–149.

- Nado, J. (2014). Philosophical expertise. *Philosophy Compass*, 9, 631–641.
- Nado, J. (2016). The intuition deniers. *Philosophical Studies*, 173(3), 781–800.
- Nahmias, E., Morris, S., Nadelhoffer, T., & Turner, J. (2006). Is incompatibilism intuitive? *Philosophy and Phenomenological Research*, 73(1), 28–53.
- Nichols, S., Stich, S., & Weinberg, J. (2003). Meta-skepticism: Meditations in ethno-epistemology. *The skeptics*, ed. S. Luper. Ashgate. (Reprinted 2012 in *Knowledge, rationality, and morality, 1978–2010*, ed. S. Stich). Oxford: Oxford University Press.
- Nichols, S., & Ulatowski, J. (2007). Intuitions and individual differences: The knobe effect revisited. *Mind and Language*, 22(4), 346–365.
- Pinillos, N., Smith, N., Nair, G., Marchetto, P., & Mun, C. (2011). Philosophy's new challenge: Experiments and intentional action. *Mind and Language*, 26, 115–139.
- Rose, D., & Danks, D. (2013). In defense of a broad conception of experimental philosophy. *Metaphilosophy*, 44(4), 512–532.
- Schiffer, S. (1996). Contextualist solutions to scepticism. *Proceedings of the Aristotelian Society*, 96, 317–333.
- Schroer, J., & Schroer, R. (2013). Two potential problems with philosophical intuitions: Muddled intuitions and biased intuitions. *Philosophia*, 41(4), 1263–1281.
- Schwitzgebel, E., & Cushman, F. (2012). Expertise in moral reasoning? Order effects on moral judgment in professional philosophers and non-philosophers. *Mind and Language*, 27(2), 135–153.
- Sinnott-Armstrong, W. (1984). 'Ought' conversationally implies 'can'. *Philosophical Review*, 93(2), 249–261.
- Sommers, T. (2011). In memoriam: The X phi debate. *The Philosophers' Magazine*, 52, 89–93.
- Soon, C., Brass, M., Heinze, H., & Haynes, J. (2008). Unconscious determinants of free decisions in the human brain. *Nature Neuroscience*, 11, 543–545.
- Sosa, E. (2007). Experimental philosophy and philosophical intuition. *Philosophical Studies*, 132(1), 99–107.
- Stich, S., & Weinberg, J. (2001). Jackson's empirical assumptions. *Philosophy and Phenomenological Research*, 62(3), 637–643.
- Stroud, B. (1984). *The significance of philosophical skepticism*. Oxford University Press.
- Sytsma, J., & Buckwalter, W. (2016). *A companion to experimental philosophy*. New York: Wiley.
- Sytsma, J., & Machery, E. (2010). Two conceptions of subjective experience. *Philosophical Studies*, 151(2), 299–327.
- Thomson, J. (1985). The trolley problem. *The Yale Law Journal*, 94(6), 1395–1415.
- Turri, J. (2015). Skeptical appeal: The source-content bias. *Cognitive Science*, 39, 307–324.
- van Inwagen, P. (1997). Materialism and the psychological-continuity account of personal identity. *Philosophical Perspectives*, 11, 305–319.
- Vranas, P. (2007). I ought, therefore I can. *Philosophical Studies*, 136(2), 167–216.
- Weinberg, J. (2014). Cappelen between rock and a hard place. *Philosophical Studies*, 171(3), 545–553.
- Weinberg, J., & Alexander, J. (2014). The challenge of sticking with intuitions through thick and thin. In A. Booth, & D. Rowbottom (Eds.), *Intuitions*. Oxford: Oxford University Press.
- Weinberg, J., Alexander, J., Gonnerman, C., & Reuter, S. (2013). Restrictionism and reflection. *The Monist*, 95(2), 200–222.
- Weinberg, J., Gonnerman, C., Buckner, C., & Alexander, J. (2010). Are philosophers expert intuiters? *Philosophical Psychology*, 23(3), 331–355.
- Weinberg, J., Nichols, S., & Stich, S. (2001). Normativity and epistemic intuitions. *Philosophical Topics*, 29(1–2), 429–460.
- Williamson, T. (2005). Armchair philosophy, metaphysical modality and counterfactual thinking. *Proceedings of the Aristotelian Society*, 105(1), 1–23.
- Williamson, T. (2007). *The philosophy of philosophy*. Oxford: Blackwell.
- Wysocki, T. (2016). Arguments over intuitions? *Review of Philosophy and Psychology*, 1–23.
- Zimmerman, M. (1996). *The concept of moral obligation*. Cambridge: Cambridge University Press.