

Arguments over intuitions?

Forthcoming in *the Review of Philosophy and Psychology*

You come up with a theory about a certain philosophical phenomenon, and enthusiastically present it to your colleague. To your disappointment, it doesn't take her much time to come up with a counterexample in the form of a hypothetical scenario. Although your theory predicts that her scenario describes a case of the phenomenon, both you and your colleague agree that this prediction is wrong. Your theory doesn't seem tenable anymore. You retreat to your room, hoping to cope with the counterexample.

On the standard metaphilosophical picture, your colleague's scenario triggers an intuition that is inconsistent with your theory. You take the intuition to constitute evidence against your theory, and thus you see that you either need to fix the theory, or to explain the intuition away—show why the intuition shouldn't be trusted after all (Bealer, 2000; Bonjour, 1998; Goldman, 2010; Pust, 2000; Sosa, 2007; Stich, 1988). Although particular theories disagree on the details of what intuitions are, these theories usually take the process of forming an intuition to be opaque to the intuiter (Foley, 1998; Gopnik & Schwitzgebel, 1998; Huemer, 2005; Nagel, 2007; Talbot, 2013). Recently, however, some philosophers have questioned this methodological picture (Deutsch, 2009; 2010; 2015; Cappelen, 2012). According to them, philosophers use arguments to secure their responses to thought experiments, and therefore the intuitiveness of a particular response is philosophically irrelevant—what matters is how good the supporting arguments are. If this alternative construal of thought experiments is right, it would mean trouble for many philosophical agendas: there would be little need for theories of philosophical intuition, for experimental philosophers running their studies, and for investigating psychological mechanisms underlying particular intuitions (Greene, 2003; Scholl, 2007).¹

Acknowledgements: I would especially like to thank Julia Staffel for her extensive advice and feedback. I also want to thank for very helpful discussion John Heil, Todd Jones, Ron Mallon, Jennifer Nagel, Howard Nye, Gillian Russell, Stephen Stich, Brian Talbot, two anonymous referees from this journal, and the audience members at the 2015 Pacific APA conference in Vancouver and the 2015 Buffalo Experimental Philosophy Conference. Finally, I also thank Dominik Dzedzic, Adrianna Senczyszyn, and Artur Tanona for help with the Polish translation, and Daniel Márquez and Louis Quero for help with the Spanish translation.

¹ Cappelen (2014) thinks that experimental philosophy is a completely misconceived endeavor. Deutsch is more congenial, leaving some room for experimental philosophy as the study of people's beliefs: "such data [gathered by experimental philosophers] are relevant to how we should treat others and how, more fundamentally, we should understand the social practices of different groups of people. Data on the philosophical intuitions of people from different cultural groups serve these broadly ethical goals and has real potential for fostering cross-cultural understanding and respect" (2015, p. 160). Neither author, however, takes experimental philosophy to be of any help in solving purely philosophical problems.

Now, contrasting arguments with intuitions requires that the former deliver verdicts independently of the latter—regardless of whether it’s intuitive, the right response to a thought experiment is the one that is supported by best arguments. However, this is a strong assumption, and it’s possible that the contrast between intuitions and arguments is overblown. So, for instance, it’s possible that intuiting a response to a particular scenario is correlated with being persuaded by an argument, because both stem from an application of the same cognitive resource. In such a case, there would be no principal difference between intuitions and arguments. My aim here is to explore this possibility.

In what follows, I focus on Deutsch’s (2010) attack on the claim that intuitions are used as evidence in philosophy.² In the first section, I present Deutsch’s view, illustrating it with an instance of the Gettier case. Then, I discuss empirical results that contradict a prediction that Deutsch seems to be making in his argument: that people whose intuition disagrees with Gettier’s would change their mind after being exposed to relevant arguments. In the third section, I focus on a possible reply to these results—that even if non-Gettier intuiters don’t change their evaluation of the scenario after being exposed to the arguments, they *should*. In this section, I also offer an account of the relationship between intuitions and the arguments that, according to Deutsch, are used to evaluate thought experiments. On this account, intuitions are produced by roughly the same mechanism that is used to evaluate the arguments, and so whether one finds an argument persuasive correlates with whether one has the intuition already. Finally, the fourth section is devoted to discussing two objections to my interpretation of the results: that non-Gettier responses might be due to performance errors, and that the cognitive resources responsible for evaluating scenarios aren’t the ones responsible for apprehending arguments.

1. Philosophers don’t intuit—they argue

Consider the following Gettier case:

John walks through the market square, and wonders what time it is. He looks at the clock on the town hall tower, and sees that the clock shows 12pm. And indeed, it is 12pm. However, John doesn’t realize that the clock stopped exactly twelve hours ago, and this is why it shows the correct time.

Does John really know that it’s 12pm or does he only believe it?

According to the theory of knowledge as justified true beliefs—the *JTB theory*, for short—John knows that it’s 12pm, because his belief is true and justified. Contrary to this evaluation, philosophers’ standard response is that John’s belief doesn’t qualify as knowledge. Since the JTB theory doesn’t capture this response, the theory is deemed incorrect.

Where does this standard response to this hypothetical case come from? According to a popular view on how such cases are evaluated, when you read the question at the end of the scenario, you have an intuition: no, John doesn’t know that it’s 12pm. You don’t infer this response from any

² This attack has been influential; for instance, Cappelen admits that his book was partially inspired by Deutsch (Cappelen, 2012, p. 20). Hence, much of what will be said pertains to the core part of Cappelen’s argument as well.

further premises, and you don't look for an argument to support it—you just have an intuition, period. The process of arriving at this answer is opaque to you, yet you still trust your answer. Then, this answer lets you deem the JTB theory false, for the theory's prediction that John knows that it's 12pm disagrees with your intuition.

However, according to Deutsch (2010), refuting the JTB theory using a Gettier case doesn't require intuitions. To refute the theory, it's enough to secure that Gettier's counterexamples are genuine—that the agents described in the scenarios really lack knowledge. A question immediately arises: where would a philosopher get her response to a thought experiment if not by intuiting it? Deutsch notices that instead of invoking intuitions, a philosopher can construct arguments to defend her evaluation of a thought experiment. In the case of a Gettier scenario, Deutsch indicates three arguments that emerged in the literature on Gettier cases. He points out that these arguments, which he dubs *G-grounds*, don't 'concern who intuits what' (p. 454), and therefore don't depend on intuition.³ Here they are:

The argument from causal disconnection. '*S*'s justified true belief that *p* might fail to be knowledge if there is a disconnection between: (i) what causes *S* to believe *p*, and (ii) what makes *S*'s belief that *p* true' (p. 455). As applied to the tower-clock case, the argument says that John doesn't know that the current time is 12pm, because his belief is caused by what is shown by a defunct clock rather than a working one.

The argument from defeaters. '*S*'s justified true belief that *p* may fail to qualify as knowledge if there are "epistemic defeaters" to *S*'s justification. Roughly, *q* is defeater to *S*'s justification for *p*, if *q* is compelling evidence against *p* about which *S* is unaware' (p. 455). In the tower-clock case, John doesn't know that it's 12pm, because if he realized that the clock wasn't working, he would suspend his belief.

The argument from epistemic luck. '[A]gents in (at least some) Gettier cases are only luckily correct in believing what they do, and that the presence of this sort of 'epistemic luck' explains why such agents lack knowledge' (p. 455). So, John was merely lucky to have looked at the tower clock exactly when it was showing correct time; had he done that five minutes later, he wouldn't have formed a true belief.

According to Deutsch, since these three arguments are good arguments, they indicate the correct solution to the Gettier case—they 'provide a very strong case for thinking that the [Gettier] cases are genuine counterexamples to the [...] epistemological theories they are meant to challenge' (p. 459). If this is right, how should philosophers treat the fact that, after reading the Gettier case, some people don't share its author's original intuition? According to Deutsch:

We can try to convince those who disagree with us by providing justifications for our intuitions. If our justifications are good enough, we can say [...] that we are right and they

³ Page numbers in parentheses refer to (Deutsch, 2010).

are wrong [...]. If [some people] intuit otherwise, they ought, on reflection, and after being exposed to the [...] G-grounds, to capitulate (pp. 458-459).

In this passage, Deutsch commits himself to two claims: a general one, that arguments point to the right solutions of thought experiments, and a specific one, that the three G-grounds indicate the right answer to the Gettier case. Yet, Deutsch is ambiguous when he claims that a person having a non-Gettier intuition *ought to capitulate* after being exposed to a G-ground. On one—descriptive—reading, the claim means that exposing such a person to a G-ground *would* change her response to the Gettier case. So construed, the claim is a forthright empirical prediction; to foreshadow the next section, let me say that the data gathered so far doesn't support this prediction.

On the other—normative—reading, the claim means that a non-Gettier intuiter *should* change her answer, even if after being exposed to the G-grounds she sticks to her original response. That is, if she doesn't change her mind, she'll act epistemically wrong, since she was offered good reasons for giving the Gettier response. Ostensibly, the normative claim is independent of any empirical results. Yet, in the last section, I'll argue that the same data undermines this claim as well.

2. *Would non-Gettier intuiters capitulate?*

Deutsch claims that a non-Gettier intuiter *ought to capitulate* after being exposed to a G-ground; as said, on one interpretation of this claim, he makes a prediction that non-Gettier intuiters will change their response to the hypothetical scenario when presented with the G-grounds. But there's a more important reason for why, on Deutsch's account, it's natural to expect the G-grounds to influence people's evaluations of the case. This reason is that, *prima facie*, a good argument, if understood, should persuade.

Of course, not all good arguments persuade. There are arguments so complex that only people with specific training and background knowledge, and only after reflecting on these arguments for a long time, could understand them and recognize them as good arguments, and thus whether laypersons are persuaded by such arguments is inconsequential. An extreme example would be the proof of the four color theorem, which was produced by a computer. Although the theorem is now taken to be proven—that is, its proof is believed to be a good argument for theorem—the proof is too long and complex for anyone to understand it, and thus to be persuaded by it (*rather*, what seems to persuade is the mere fact that such a proof has been produced).⁴ Yet—and I'll expand on this thought in the next two sections—the G-grounds are far too simple to maintain that if a person isn't persuaded by them, it's because she failed to understand them. If this is right, then Deutsch's theory that the G-grounds are good arguments does predict that people will tend to be persuaded by these arguments.

In this section, I present empirical data suggesting that this prediction is false. Here's one way to test the prediction: participants in one experimental condition read a Gettier scenario alone, and

⁴ The theorem says that, given any partition of a plane into contiguous regions, at most four colors are required to color the regions in a way that no two regions sharing a border have the same color. In lay terms: any map can be colored with at most four colors so that no countries that share a border have the same color. Interestingly, some philosophers of mathematics deny that the proof is a genuine proof, precisely because no mathematician can understand it, and thus be convinced by it (Tymoczko, 1979)—which adds fuel to the claim that good arguments should persuade.

then respond to the question whether the agent's belief counts as knowledge. Participants in another condition read the scenario followed immediately by a G-ground, and then they answer the question. Let's say that in the first condition a substantial minority of participants gives the non-Gettier response. Deutsch's prediction would be corroborated if this minority turned out to be significantly smaller in the other condition, for this would mean that the G-ground was persuasive—that it influenced the responses of the participants who otherwise would evaluate the scenario as a case of knowledge. If no such clear change is observed, Deutsch's prediction isn't supported.⁵

Moreover, I have also examined whether exposing people to an argument against Gettier's solution to the case would influence their judgments. This required adding a third condition, similar to the second one, but where the participants were primed with an argument for a JTB solution to the case. This condition would help in discriminating between Deutsch's general claim and the specific one. If the participants in this condition were persuaded by this argument, but the participants in the other condition weren't persuaded by a G-ground, it would still be consistent with Deutsch's general claim that arguments give solutions to hypothetical cases. Yet, his specific claim—that the right solution to the Gettier case is the one preferred by philosophers—would be unsupported, for the participants would be persuaded by an argument that doesn't persuade philosophers, but unimpressed by philosophers' own argument for the Gettier case.⁶ To corroborate Deutsch's stance, the participants should, ideally, be persuaded by a G-ground, but not by an argument for the opposite solution.⁷

2.1 The methodology

To elicit participants' responses, the tower-clock scenario was used. Two arguments were examined: one G-ground—specifically, the argument from epistemic luck—and one argument supporting the non-Gettier response. As applied to the tower-clock case, the argument from epistemic luck—let me call it the *only-believes* argument—reads:

⁵ In my argument, I assume that what participants report in a questionnaire is their intuition. This assumption isn't uncontroversial—many attacks on experimental philosophy argue against this assumption. For instance, Bengson (2013) points out that participants' answers can express many different mental states, like guesses or inferences, not only intuitions; Kauppinen (2007) and Cullen (2010) argue that participants' responses are, most likely, influenced by both pragmatic and semantic features of scenarios, where genuine intuitions should respond to semantic features alone (Deutsch (2009, pp. 461-462) makes a similar point); Sosa (2009) notices that participants might misunderstand the scenarios by filling in details that aren't explicitly stated.

These are viable worries. However, first, responses to them have been offered (Horvath, 2010; Nadelhoffer & Nahmias, 2007; Alexander & Weinberg, 2007). Second, answering these objections is beyond the scope of this paper—Deutsch, in the argument that I focus on here, grants that participants' responses to questionnaires count as genuine intuitions, and thus I feel justified in making this assumption as well.

⁶ In such a case, to save his conclusion, Deutsch could claim that arguments that convince philosophers but not laypersons are superior to arguments that convince laypersons but not philosophers. To do so, he'd have to explain laypersons' reactions away—for instance with a lack of background knowledge, with poor logical reasoning, or with performance errors. As I argue in the last section, this move is less plausible than it seems.

⁷ One should proceed with caution here. It still would be consistent with Deutsch's account if the participants were persuaded in both conditions (i.e., by both arguments), but gave Gettier's response if exposed to both arguments at once. This experiment doesn't test for this possibility. However, given the experiment's results, this methodological shortcoming turns out to be inconsequential.

John doesn't know that it's 12pm, because it was an accident that he checked the time exactly when the clock was showing the correct time. For example, if John looked at the clock 5 minutes later, he would be mistaken about the time.

Again, if there's no statistically significant difference between the ratio of Gettier intuiters in the no-argument condition and in the *only-believes* condition, then Deutsch's claim that non-Gettier intuiters would capitulate after being exposed to a G-ground is unjustified, at least in the case of the argument from epistemic luck.

The other argument—the *really-knows* argument—isn't discussed by Deutsch, but has been devised for the purposes of the experiment. It reads:

John knows that it's 12pm, because he believes that it's 12pm, and indeed it's 12pm. Moreover, it might be useful for him that he thinks it's 12pm. For example, imagine that John wants to catch a train at 12:10pm. He looks at the clock, thinks that it's 12pm, hurries to the railway station, and catches the train.

This argument is supposed to corroborate the non-Gettier response to the case by appealing to the fact that the belief is true, and to the usefulness of having a true belief. Hence, if the participants get persuaded by this argument, it poses a problem for Deutsch, for he maintains that arguments, not intuitions, indicate the right answer to hypothetical cases. If this argument turns out to be as persuasive as the *only-believes* one, and if it being persuasive suggests that it's a good argument, then Deutsch would have to explain away such results. Specifically, he would have to offer some criteria to show why the *really-knows* argument is bad and the *only-believes* is good, regardless of the arguments' persuasiveness.

The participant was randomly assigned to one of three conditions. In the first condition, she read the hypothetical case, and then answered the question 'Does John really know that it's 12pm or does he only believe it?' The possible responses were 'really knows,' and 'only believes.'⁸ In this condition, the participant wasn't presented with any arguments. In the second condition, the participant first read the scenario. After she confirmed that she had finished reading, the questionnaire displayed a prompt 'please, consider the following argument for at least 25 seconds,' together with a progress bar showing how much time had passed. After 25 seconds, the question whether John knows that it's 12pm was displayed below the story and the argument; the story and the argument didn't disappear though, and the participant could keep reflecting on them. In this second condition, the participant was presented with the *really-knows* argument. The third condition was similar to the second one, except that the *only-believes* argument was used.

The 25-seconds constraint was a methodological compromise. It was intended as the longest time in which, plausibly, the participant wouldn't just get bored and change the tab in the web browser, instead of analyzing the argument. Yet, after the 25 seconds passed and the question

⁸ These two responses were chosen, because they were also used in the study that Deutsch focuses on in his paper: Weinberg et al. (2001).

appeared on the form, the scenario and argument were still visible, and the participant could continue to reflect on them.

The questionnaire came in three language versions: English, Polish, and Spanish. The English version was run on participants from the U.S. and from India; both groups were recruited via the Mechanical Turk. The Polish and the Spanish version were run on native speakers recruited using chain sampling; the groups consisted mostly of college students from Poland and Spain, respectively.

2.2 The results

Figure 1 presents the ratios of non-Gettier responses—that is, the ratio of participants who responded that John really knows that it’s 12pm. For every country, the first column represents the responses in the first condition (the scenario alone), the second column represents the responses of participants primed with the *only-believes* argument, and the last column represents the responses of participants primed with the *really-knows* argument.

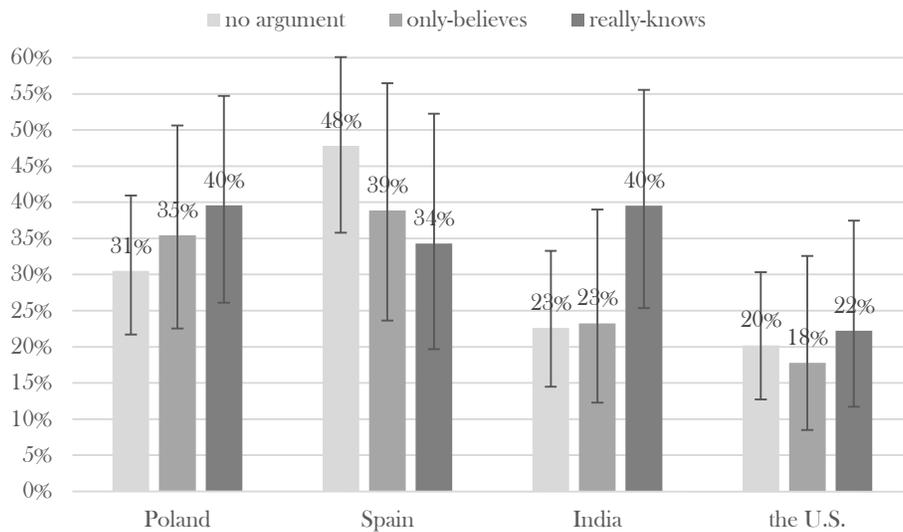


Figure 1. The proportion and 95%-confidence intervals for non-Gettier responses to the tower-clock case (i.e., John really knows that it’s 12pm).

Additional statistics of the sample, including the p-value of the test of equal proportions, are presented in table 1:

Table 1. The results of the study

	ratio of <i>really-knows</i> resp.			N			95%-confidence interval			p-value
	<i>no arg.</i>	<i>believes</i>	<i>knows</i>	<i>no arg.</i>	<i>believes</i>	<i>knows</i>	<i>no arg.</i>	<i>believes</i>	<i>knows</i>	
Poland	31%	35%	40%	95	48	48	(22%, 41%)	(23%, 51%)	(26%, 55%)	0.54
Spain	48%	39%	34%	69	36	35	(36%, 60%)	(24%, 56%)	(20%, 52%)	0.376
India	23%	23%	40%	84	43	43	(14%, 33%)	(12%, 39%)	(25%, 56%)	0.103
U.S.	20%	18%	22%	89	45	45	(13%, 30%)	(9%, 33%)	(12%, 37%)	0.87

No arg., *knows*, and *believes* denote the first, second, and third experimental condition, respectively. The p-values were obtained with the test of equal proportions, where the alternative hypothesis states that not all three ratios are equal. *N* denotes the number of participants.⁹

As both table 1 and figure 1 show, no statistically significant differences in the ratios of responses were found—no influence of arguments on responses was observed in any of the four groups examined. The ratios of *really-knows* responses in the three conditions don't differ significantly.¹⁰ This means that the experiment didn't confirm Deutsch's prediction: exposing participants either to the argument from epistemic luck, or to the *really-knows* argument, doesn't significantly influence their responses.¹¹ As always, non-significant results are hard to interpret: one cannot simply state that there is no effect, if no effect was found in a particular study. Yet, because the sample sizes were relatively big, one can expect that the effect of these two arguments on responses to the tower-clock case, if it exists, is small.¹²

Moreover, for the sake of argument, let's assume that future studies would yield similar, yet statistically significant results. Such results would still pose a problem for Deutsch, because the effect of the examined arguments is ambiguous. For instance, both in Poland and in India, the ratio of non-Gettier responses increased in both the sample primed with the *only-believes* argument and with the *really-knows* one. An opposite effect was observed in the Spanish sample, where both the arguments decreased the ratio of non-Gettier responses. The only participants who behaved unsurprisingly were Americans, since the ratio of non-Gettier responses decreased in the *only-believes* condition, but increased in the *really-knows* one. So, even if the results were statistically significant, no clear pattern would emerge: the arguments would work as intended in some countries, in other countries they would have the opposite effect; these effects would be bigger in some

⁹ Raw data and the Spanish and Polish translations of the scenario can be downloaded from <http://tomwyssocki.net/aoi.html>.

¹⁰ For all four groups, the test of equal proportions yields p-value > 0.1.

¹¹ Statistically significant difference in epistemic intuitions was found only between the Spanish sample and the Indian one (the test of proportions, $p = 0.00191$), and the Spanish sample and the American one ($p = 0.00045$). The latter result will be significant even after the Bonferroni correction for $\alpha = .1$ and ten comparisons (four tests presented in table 1 and six comparisons of unprimed responses between groups). Therefore, although the study didn't replicate the difference between the American and the Indian participants reported by (Weinberg, Nichols, & Stich, 2001), contrary to (Nagel, 2012; Nagel, San Juan, & Mar, 2013; Machery, et al., 2015; Seyedsayamdost, forth.), my results provide some evidence for the cross-cultural variability of epistemic intuitions.

¹² For each country, I have also estimated a logistic regression model with the condition as predictors and the response as the dependent variable. This analysis yielded similar, insignificant results.

countries, and smaller in other. Such a messy picture would disconfirm the prediction that non-Gettier intuiters would capitulate when presented with a G-ground.

Table 2. The differences in the ratios of non-Gettier responses in different conditions

	<i>only-believes</i>	<i>really-knows</i>
Poland	5%	9%
Spain	-9%	-14%
India	1%	17%
U.S.	-3%	2%

The numbers denote the differences in the ratios of non-Gettier responses in the unprimed and primed condition: e.g., 5% for Poland/*only-believes* means that the ratio of non-Gettier responses in the Polish sample primed with the *only-believes* argument was 5 p.p. higher than the ratio in the unprimed sample.

Were these results statistically significant, Deutsch should be worried by them for one more reason. As table 2 shows, exposing American participants to the *really-knows* argument shifts them toward the non-Gettier response no less than exposing them to the *only-believes* argument shifts them toward the standard response. That is, were statistical significance reached, the *really-knows* argument would prove no less persuasive than the *only-believes* one. In the case of the subjects from Poland and India, the *really-knows* argument would be persuasive, whereas the *only-believes* one wouldn't (in fact, it would pull people away from the *only-believes* response). Therefore, Deutsch would have to explain why people find the *really-knows* argument no less persuasive than the *only-believes* argument. According to him, the epistemic luck argument points to the right answer to the Gettier case. However, if how appealing an argument is indicates, to some extent, how well the argument justifies its conclusion, then such results would suggest that the *really-knows* argument rather than the epistemic luck argument indicates the correct response to the tower-clock case.

An obvious objection at this point in my argument would be that the experiment examined only one G-ground, and it's possible that using the other two would yield different results. This is a viable worry. Still, let me offer a response: if a person isn't persuaded by one G-ground, there's a significant chance that she wouldn't be persuaded by the other two. Here's why. As the next section shows, one could expect that an argument that shares the simple structure of the *only-believes* argument would be persuasive if and only if one is already disposed to intuit their conclusion. If this expectation is plausible, it's likely that similar results would be obtained using either of the remaining G-grounds, because, as I'll show in the following section, they share this simple structure, and they point to the same conclusion—that the Gettier case isn't a case of knowledge. Yet, again, this response to the worry succeeds only if the argument presented in the next section is valid.

3. *Should non-Gettier intuiters capitulate?*

How could Deutsch respond to the experiment's results? The most promising move would be to endorse the normative reading: even if non-Gettier intuiters aren't compelled by the epistemic luck argument, they *should* be. So, it's these intuiters' problem that they can't recognize a good argument, rather than a fault of the argument that it doesn't persuade them. If they don't change their evaluation

of the tower-clock case, they are doing something epistemically wrong, for they have been presented with a good reason to adopt the correct answer. Let me offer two responses to this objection.

First, a short response. Consider two main positions in the epistemology of disagreement: the *equal weight view* (Elga, 2007; Feldman, 2006) and the *symmetrical no independent weight view* (Kelly, 2010). On the first view, if two epistemic peers disagree about a claim—where your epistemic peer is similar to you in terms of possessed information and cognitive abilities—both should suspend their judgments. The view, as applied to the situation in question, reads: if your epistemic peers aren't persuaded by your argument, you should distrust your assessment of the argument's soundness—the argument isn't a good one. On the second position on disagreement, “both parties [...] might be perfectly reasonable even if neither gives any weight at all to the opinion of the other party” (Kelly, 2010, p. 186). According to this view, it's reasonable for people who were persuaded by the epistemic luck argument to think that the agent lacks knowledge. Yet, the same holds for the ones who weren't persuaded—it's reasonable for them to think that the agent's belief counts as knowledge. Therefore, neither of these views on disagreement entails that non-Gettier intuiters *should* adopt Gettier's intuition.

Deutsch could respond that people unpersuaded by the G-grounds aren't his epistemic peers, yet the onus of showing these people's epistemic inferiority is on him. Of course, the reason “they are epistemically inferior because they aren't persuaded by a G-ground” won't hold up, because it already assumes the conclusion: that the argument is a good one. Deutsch would have to find some other, independent grounds for showing that philosophers, in their capacity of entertaining the G-grounds, are epistemically superior to the non-Gettier respondents.

There is a simple reason to take philosophers to be significantly better at evaluating the G-grounds: philosophers are trained in dealing with arguments, and are motivated to reflect on how strong an argument is (after all, that's what philosophers do for a living). For instance, laypersons often are persuaded by arguments that contain logical fallacies, whereas philosophers are taught early in their training how to identify such mistakes. Or, when an argument to be evaluated is complex, a philosopher would decompose it into premises and evaluate them separately, whereas a layperson would often give up the task. Hence, whatever non-philosophers considering arguments in an online questionnaire say, it's inconsequential for philosophy.

Yet, professional philosophers' unique skills seem to play no role in evaluating the arguments used in the experiment. Just recall the *only-believes* argument: ‘John doesn't know that it's 12pm, *because* it was an accident that he checked the time exactly when the clock was showing the correct time.’ Its logical structure is extremely simple, no technical terms are used, and no sophisticated background knowledge is presupposed. The argument had only two elements that the participant was supposed to apprehend: first, that John's belief was formed by accident, and second, that the fact that the belief was formed by accident entails that it doesn't count as knowledge. Moreover, the argument included an example clarifying the intended meaning of *an accident*: ‘if John looked at the clock 5 minutes later, he would be mistaken about the time.’ Therefore, it's hard to point out the skills that only a philosopher, and not a layperson, would possess that are necessary for evaluating

the *only-believes* argument. Also, reflecting on this argument doesn't require much time: plausibly, 25 seconds are enough to evaluate whether the fact that John was lucky is relevant for whether his belief counts as knowledge. A similar point could be made about the *really-knows* argument.¹³

Now, for the second response to the objection that even if people aren't convinced by the G-grounds, they should be. The previous response relied on the fact that the distribution of intuitions didn't change as a result of presenting the participants with the two arguments. However, this response didn't draw any connection between what one intuits and what arguments one is sensitive to—for a proponent of this response, the fact that there was no significant difference in the distributions of primed and unprimed responses could be a sheer coincidence. In turn, on the argument to be discussed now, this similarity in distributions is no surprise. The reason is, as I'll argue, that exhibiting an intuition to a Gettier case and evaluating a G-ground uses the same cognitive resource—the concept of knowledge—and therefore whether one intuits that the scenario exemplifies knowledge is correlated with whether one would be persuaded by the G-ground. If this argument succeeds, the consequence for Deutsch's view is that collecting responses to hypothetical cases and responses to arguments are likely to give the same results, and therefore, in the case of arguments structured similarly to the G-grounds, experimental philosophers shouldn't worry too much that they measure intuitive responses rather than reactions to such arguments.

To begin this second response to the objection that non-Gettier intuiters should be convinced by the G-grounds, let me consider a particular theory of intuitions—one on which intuitions are, roughly, applications of concepts.¹⁴ According to this view, when you're asked whether John knows that it's 12pm, you unconsciously run through the details provided by the hypothetical scenario, and check whether John's belief satisfies your concept of knowledge. Here, what it means to *satisfy a concept* depends on what the correct psychological theory of concepts is. For instance, on exemplar theory, you (unconsciously) compare John's belief with what you previously judged as instances of knowledge. If his belief is sufficiently similar to these instances, you have the *really-knows* intuition; otherwise, you give the *only-believes* response. On prototype theory, on the other hand, your concept of knowledge is structured as a weighted list of features: you'll judge that John knows that it's 12pm only if his belief has the features that you associate with knowledge.¹⁵ Or, on the theory of concepts, your concept of knowledge is encoded as a tacit or implicit theory. So, you'll say that John's belief counts as knowledge, if the belief satisfies certain law-like relations; for instance, if

¹³ There are, of course, arguments so complex that philosophical training is required to see their validity or invalidity, or arguments embraced or rejected by laypersons just because they agree or disagree with their conclusions, arguments for the theory of evolution being a salient example. However, the G-grounds and the *really-knows* argument aren't of this sort.

¹⁴ Some version of this theory is endorsed, among others, by (Goldman, 2010; Goldman & Pust, 1998; Hintikka, 1999; Sosa, 2007), and, to some extent, also by (Devitt, 2006; Kornblith, 1998). In this context, by a concept of X I mean the part or capacity of the mind that lets you classify something as an X, rather than an abstract entity, like, say, the Fregean sense of X. This part of the mind stores the information required to classify an object as an X; the information is retrieved and used when the classification occurs. This notion of concepts and classificatory judgments is further explored in the last section.

¹⁵ More precisely: you'll Judge John's belief as a case of knowledge only if the aggregated weight of the features of his belief reaches a certain threshold, where the weights come from your concept—i.e., your prototype—of knowledge.

your implicit theory encodes a rule “if a belief is luckily acquired, it doesn’t count as knowledge,” John won’t know that it’s 12pm on your concept. Again, on any psychological theory of concepts, the exact process of applying the concept is usually opaque to you; all you’re aware of is only the final categorization.¹⁶

The theory of intuitions as applications of concepts, as sketched above, can be now used to explain the results of the experiment. First, notice that if some people have different intuitions, one of the possible explanations is that their concepts differ. Under the assumption that the concepts differ, if there is a substantial minority of *really-knows* responses in the sample, there should be a substantial minority of people having non-Gettier concepts—i.e., concepts on which Gettierized beliefs count as knowledge.¹⁷

Second, recall that Deutsch’s three G-grounds and my *really-knows* argument share a simple structure: a belief does/doesn’t count as knowledge, *because* it has a certain feature or features (e.g., “John doesn’t know that it’s 12pm, *because* it was an accident that...”). The claim in the *because*-clause doesn’t appeal to any evidence that can be gathered independently of what you already judge as knowledge. Rather, it seems that in order for you to treat this claim as a reason to deny John’s belief the status of knowledge, your concept of knowledge already has to encode the features invoked in the *because*-clause. For instance, if you are to be persuaded by the epistemic luck argument, your concept should already encode that a luckily-acquired belief doesn’t count as knowledge.

Now, let me combine the two observations. On the view just offered, your concept of knowledge underlies both your epistemic intuitions and your judgments whether the G-grounds are good arguments. Hence, if a group of people don’t exhibit Gettier intuitions, they aren’t likely to be impressed by Deutsch’s G-grounds either; and indeed, the experiment’s results confirm this conclusion. If this story is accurate, what does that mean for Deutsch’s claim that the G-grounds indicate the correct answer to the Gettier case?

“If our justifications are good enough, we can say [...] that we are right and they are wrong,” says Deutsch (2010, p. 459). On the view that I am proposing, Deutsch finds the G-grounds good enough *just because* his concept of knowledge resonates with the case’s features highlighted by these arguments. Had he a different concept, he would evaluate the G-grounds differently (as well as he would have non-Gettier intuitions).

¹⁶ For a presentation of prototype theory, see (Rosch, 1973; Rosch, 1978; Rosch & Mervis, 1975). For exemplar theory, see (Medin & Schaffer, 1978; Nosofsky, 1988). For the theory theory, see (Carey, 2009; Keil, 1994; Murphy & Medin, 1985).

¹⁷ One could think that showing that people have different concepts of knowledge doesn’t say anything about the subject matter of epistemology, which many epistemologists take to be non-conceptual and non-linguistic (Brown, 2012). However, since on the standard methodological picture epistemic intuitions are treated as evidence for what knowledge is, if our intuitions differ, this means that we disagree about what this philosophically interesting phenomenon, knowledge, really is. It’s possible that our best theory ends up telling us that there is no such thing as a universally shared concept of knowledge. If this is the case, then studying knowledge requires that we make decisions about what the subject of inquiry should be. For example, we might try to focus on a notion of knowledge that fits well with the normative postulates associated with knowledge. A proper discussion of this point can’t be given here.

Although Deutsch can still insist that the G-grounds provide the right answer to the Gettier case, this claim boils down to the question whose concept—his or a non-Gettier intuiter’s—is *the* correct one. Now, Deutsch can hold that philosopher’s concept is superior to laypersons’, but this move is already familiar to experimental philosophers. It goes by the name of the expertise defense, and discussing it is beyond the scope of the paper.¹⁸ Here, it suffices to notice that Deutsch’s proposal to move the focus from intuitions to arguments doesn’t address the ultimate problem—is there anything like *the* correct concept of knowledge, and, if so, how can philosophers identify it? Neither invoking intuitions nor appealing to the G-grounds to discover the right answer to the Gettier case will work, unless you already buy into the assumption that philosophers’ concept of knowledge is the correct one. However, if you do think that philosophers’ concept is accurate, then intuitions suffice to secure the proper answer to a Gettier case, for they are applications of *the* accurate concept of knowledge.

In the above answer to the objection that non-Gettier participants ought to change their responses, I have assumed the theory of intuitions as concept applications. However, this response to Deutsch can be easily reused by a meta-philosopher who doesn’t subscribe to this particular theory. It’s enough if the philosopher takes the source of epistemic intuitions to play a crucial role in evaluating the G-grounds. For instance, say that, according to the theory of intuitions that you find most plausible, a person who reads a Gettier scenario has an intuition in virtue of apprehending the platonic form of knowledge: instead of classifying the agent’s belief in virtue of the information stored in the intuiter’s mind, the intuiter directly grasps the relevant platonic form when the classification occurs.¹⁹ Moreover, say that you also think that evaluating the argument from epistemic luck requires apprehending the relevant platonic form. On such a picture, this platonic form will determine, modulo performance errors, both the person’s intuitions and whether she finds the argument persuasive, and so examining her epistemic intuitions would suffice to reveal whether she would be impressed by the G-grounds.²⁰

To wrap up the discussion of the postulate that non-Gettier intuiters *should* be persuaded by a G-ground (even if they aren’t), let me offer a picture of Deutsch’s G-grounds. For Deutsch, the G-grounds are arguments that point to the right solution of the Gettier case regardless of whether you find this solution intuitive or not. In the light of my argument, the G-grounds are explanations of the

¹⁸ For instance, such arguments have been offered by (Devitt, 2006, pp. 104-105; Hales, 2006, p. 171; Horvath, 2010, pp. 464-472; Ludwig, 2007). In response, experimental philosophers both theorized that philosophical training doesn’t improve one’s concepts (Weinberg, Gonnerman, Buckner, & Alexander, 2010), as well as conducted empirical studies suggesting that philosophers’ intuitions are susceptible to the similar manipulations as laypersons’ (Schulz, Cokely, Feltz, & Adam, 2011; Schwitzgebel & Cushman, 2012; Tobia, Buckwalter, & Stich, 2013; Tobia, Chapman, & Stich, 2013).

¹⁹ The view can be attributed not only to Plato, but also, as Goldman and Pust (1998) point out, to (Steup, 1996; Plantinga, 1993; Katz, 1981). One notable difference between the concept-application theory and the form-apprehension theory is that, on the first one, having an intuition doesn’t require any link between the capacity producing the intuition and any extra-mental entity, whereas, on the second one, intuiting requires such a link (Goldman, 2007).

²⁰ Notice, however, that the proponent of this response either would have to explain the variability of intuitions with an error theory, or hypothesize that people having different intuitions grasp different platonic forms. Without any further argument, the latter claims leads to a conclusion similar to Weinberg et al.’s (2001): that one can’t ascertain which platonic form is the right one, or even that there is a unique right one.

standard intuition, identifying the features of the case that make one intuit that the belief doesn't count as knowledge.

Now, under certain conditions such an explanation can function as an argument too—that is, it can persuade to change one's initial response. Such a change in the intuiter's response is possible when she makes a performance error when evaluating the case, not noticing a certain detail that is relevant for evaluating the scenario. Subsequently, when she reads the explanation-argument, she notices that she omitted a relevant detail, reevaluates the scenario (now taking the detail into account), and gives a different response. For instance, say that a person reads the tower-clock scenario, but doesn't realize that John was lucky to look at the clock at the right time, and therefore judges it to be a case of knowledge. Yet, when she's presented with the epistemic luck argument, she realizes that she missed an important detail of the scenario, and reapplies her concept of knowledge to the case. This time, however, she keeps in mind that John was lucky, and, in result, she intuits that John only believes that it's 12pm.

Again, the intuiter would be persuaded only if, on her concept of knowledge, this detail is indeed relevant—e.g., only if her concept of knowledge already encoded that epistemic luck is incompatible with knowledge. If on her concept of knowledge this detail is irrelevant, she wouldn't be persuaded by the epistemic luck argument.

4. Performance errors and multiple concepts

My preferred interpretation of the result is that non-Gettier respondents have a slightly different concept of knowledge than the ones who share Gettier's evaluation of the case. There are at least two problems to this interpretation. The first one is that non-Gettier responses could be explained away with performance errors, which would mean that the responses don't express the respondents' actual concept. The second problem is that people might be using different concepts when evaluating a case and when evaluating an argument, and thus it doesn't make sense to talk about *the* concept of knowledge of a particular person or group.

Let me begin the discussion of the objection from performance errors with an observation. In a usual experiment on intuitions, where only participants' unprimed responses to cases are gathered, the objection from performance error typically points out how participants could misunderstand the scenario or the intuition prompt. In the case of the current study, formulating the challenge is more complex, for the proponent of the challenge would have to explain:

1. why the ratios of non-Gettier responses in the no-argument condition were so high (at least one-fourth of the answers);
2. the fact that the arguments didn't affect the responses;
3. *or*, were the results significant, (a) why the arguments were sometimes counterproductive, and (b) why the *really-knows* argument worked in some groups.

First, let me note that high ratios of non-Gettier responses are consistent with several recent studies on Gettier cases. For instance, most of the participants in Starmans and Friedman's (2012)

study attributed knowledge to agents in three different Gettier cases.²¹ They have used three Gettier scenarios, and, for each scenario, the ratio of participants attributing knowledge were higher than 60%. Similarly, Turri, Buckwalter, and Blouw (2015) report that in response to three Gettier cases, more than 40% of the participants attribute knowledge to the protagonist, whereas Nagel reports that, in an experiment using a scenario very similar to my tower-clock case, 44% of her participants had a non-Gettier intuition (Nagel, 2012, p. 175). Moreover, Nagel, San Juan, Mar (2013) asked their participants, who read a Gettier scenario, whether the agent knows the relevant proposition. Then, the participants who initially said that the agent knows the proposition were asked a subsequent question: does the agent know the proposition, or does she merely feel like she knows it? Even after being asked this auxiliary question, 35% of all participants ascribed knowledge to the agent.²² Since in all these experiments the ratio of non-Gettier intuitions was significant, to explain away the *really-knows* responses in my study, one should ideally appeal to general features of Gettier cases rather than to particular features of the tower-clock scenario.²³

One such possible source of error is that some participants, including the non-Gettier respondents, could have failed to apprehend the scenario's details. Hence, the non-Gettier intuitions don't reflect the respondents' concept of knowledge—had these participants paid attention to the details, they would have given the Gettier response. If such performance errors are rife, then one shouldn't take non-Gettier respondents to possess concepts, on which Gettier cases count as knowledge.

There are at least three responses that one could offer to this objection. The first response was already given: the case and the arguments were short and simple, they didn't use any technical terms or require any sophisticated background knowledge. Admittedly, on the other hand, since the scenario and the argument were short, the participants could easily omit an important detail. Still, together with other results showing that there's a significant minority ascribing knowledge in a diverse array of Gettier cases—including ones that are much longer than the tower-clock scenario—it seems that a widespread error in apprehending either the scenario or the argument due to omitted details would be rather unlikely. Second, since the *only-believes* argument highlights the scenario's particular detail—John being lucky—it's reasonable to expect that reading the argument made the participants pay attention to this detail. Therefore, because being exposed to the argument didn't significantly influence the distribution of responses, it seems plausible that for non-Gettier respondents luckily acquired true beliefs can sometimes really count as knowledge. It is, of course, also possible that the participants indeed missed some relevant details of the scenario, for instance, precisely because the stories were so short. But to explain away the most important results of the

²¹ These results were replicated (Turri, 2013).

²² In these studies, control conditions included also cases of justified false beliefs. In these cases, participants were significantly less likely to make knowledge ascriptions than in Gettier cases, which means that non-Gettier intuitions can't be explained by stipulating that by "really knows" the participants meant, for instance, "is convinced."

²³ However, in a cross-cultural study by Machery et al. (2015), the ratios of participants ascribing knowledge in Gettier cases were lower. For instance, one of the cases, called the trip case, was evaluated as follows: 25% (U.S.), 35% (India), 32% (Brazil), and 18% (Japan) said that the protagonist in the scenario knows the proposition; the other available response read that the agent feels like knows that proposition, but she doesn't actually know it.

study—that arguments don't influence participants responses, or influence them in a surprising way—the objection from performance errors needs to be stronger. Namely, one needs to hypothesize not only that non-Gettier respondents in the *no-argument* condition didn't comprehend the story, but also that non-Gettier respondents in the *only-believes* condition failed to comprehend the argument. What's more, the respondents would have to miss only the details of the story that, according to epistemologists, make the story a case of non-knowledge—e.g., if a participant thought that John has a false belief, or if she was skeptically-minded and thought that glancing quickly at tower-clocks never yields knowledge, she would give the standard response to the case. These two failures—to comprehend the relevant details *and* the G-ground—would have to happen at the same time, if a person with a standard concept of knowledge were to give a non-Gettier response in the *only-believes* condition. This possibility doesn't seem very likely.

Lastly, virtually any similar study—be it in social psychology, linguistics, or experimental philosophy—can be undermined by the objection from performance error. Yet, the results of such experiments in social psychology and linguistics are usually trusted, unless one identifies some salient source of performance errors found in the experiment's methodology. Hence, to undermine this experiment's results, one would have to identify such a plausible source of performance errors.

Now, in the literature on epistemic intuitions, some such sources of errors have been already indicated. For instance, Turri (2013) presented his participants with Gettier scenarios, where the structure of the case was made conspicuous, and observed that most of his participants judged that the agent lacks knowledge. Let me now discuss his experiment using one of these scenarios (names of the modules in parentheses weren't visible to the participants):

(No Luck) Grace is alone in the living room of her locked apartment, writing a letter with a blue Bic pen. She puts the letter and the blue Bic pen down on her coffee table and goes into the bathroom to take a quick shower. While she is in the shower, the apartment remains locked and nobody enters.

(Traffic) Just as Grace began her shower, several construction vehicles drove by her apartment building. Because Grace was in the shower, she didn't notice the traffic.

(Bad Luck) The vibrations from the construction traffic caused something to happen. The vibrations caused the blue Bic pen, which she had set down, to shift position, fall off the coffee table, and land on the floor. Because Grace was still in the shower, she didn't notice that this happened.

(Good Luck) The vibrations from the construction traffic caused something else to happen: For a very long time, there has been a blue Bic pen hidden in the light fixture in Grace's living room, right above the coffee table. No one, not even Grace, has ever noticed this hidden pen. The vibrations caused this hidden blue Bic pen to fall down and land on the coffee table. Because Grace was still in the shower, she did not notice that this happened (Turri, 2013, p. 12).

The intuition prompt read: "Right before Grace finishes her shower, she _____ that there is a blue Bic pen on the coffee table. [really knows/only thinks she knows]" (Turri, 2013, p. 12).

The experiment had five conditions. The first scenario consisted of the *no-luck* module alone, and thus described a regular case of knowledge. The second scenario consisted of the *no-luck*, *traffic*, and *good-luck* modules, also describing a regular case of knowledge. The third scenario consisted of the *no-luck*, *traffic*, and *bad-luck* modules, and therefore described a case of justified false belief. Finally, the two crucial scenarios included all four modules; these scenarios described Gettier cases. In one of the conditions, *one-stage Gettier*, the participants read the whole scenario at once, and then responded to the intuition prompt. In the other Gettier condition, *three-stage Gettier*, the participants were presented with each module on a different screen; after each module, the participants were asked whether there's a blue pen on the table. The purpose of asking these questions was to make the participants track what happens to the truth-maker of the agent's belief. In this condition, the intuition prompt was displayed after the last module.

In this experiment, the ratios of knowledge ascriptions were high for the regular cases of knowledge (81% and 76%, no significant difference), and low for the justified false belief case (16%). Crucially, 48% of the participants ascribed knowledge in *one-stage Gettier*, but only 29% did that in *three-stage Gettier*. Moreover, the ratio in *three-stage Gettier* doesn't significantly differ from the one of false justified belief. From these results, Turri concludes "that the conspicuous tripartite structure effectively guides laypeople to competently assess Gettier cases. When probed this way, laypeople [...] reveal that they share the Gettier intuition" (Turri, 2013, p. 15). Thus, according to Turri, non-Gettier respondents don't pay enough attention to what's happening with the truth-maker of the protagonist's belief; if they did, they would agree with Gettier. If that's true, then my preferred interpretation of my results—that there's a minority of people for whom Gettier cases really do exemplify knowledge—can't be right.

There are two ways to respond to this worry. The first way is to argue that the conclusion that Turri draws from his results is too strong. Notice, for instance, that 29% of his participants still ascribed knowledge to the protagonist, when the scenario was divided into the three modules. Equally well, then, one could conclude that even when the tripartite structure of the case is made conspicuous, almost one in three persons ascribes knowledge to the agent. Moreover, Turri doesn't report whether the crucial difference in ratios between the two Gettier scenarios was statistically significant. Lastly, it's possible that presenting the modules separately actually introduced a performance error. For instance, since the intuition prompt was displayed with the last module, a participant reading the scenario might focus on the last module when answering the scenario. Imagine that the concept of knowledge is stored as, say, a prototype. Since the last module in *three-stage Gettier* described a luckily event, when responding to the intuition prompt, the participants might focus on this feature of the scenario, rather than, say, the belief being true or justified. Therefore, in this particular context, the weight of one of the features—the belief being luckily true—might be temporarily increased. If that's how Turri's manipulation works, the participants' responses wouldn't accurately reflect their prototype of knowledge.

The other, quite different way to respond to the worry from Turri's results is to agree to his conclusion, and consider the consequences for Deutsch's thesis. My results suggest that either arguments don't influence responses to Gettier cases, or this influence is ambiguous. On the other

hand, Turri's results suggest that framing a case in a certain way makes the participants more likely to give a response that, according to most epistemologists, is right. Therefore, if you think that "only believes" is the proper answer to a Gettier case, it seems that framing the scenario in a specific way is more effective a mean to persuade a non-Gettier intuiter than the argument from epistemic luck. Keeping that in mind, recall the condition for good arguments: a good argument—i.e., one that points to the right solutions of thought experiments—should, *prima facie*, persuade. If Turri is right and making the structure of a scenario conspicuous "effectively guides laypeople to competently assess Gettier cases," then the correct solutions to hypothetical cases aren't indicated by arguments, but by intuitions triggered with carefully prepared thought experiments.

Nagel, San Juan, and Mar (2013) also offer explanations of non-Gettier responses in terms of performance errors. The authors notice that "[t]hese cases are by nature complex, involving actual or hypothetical oddities in evidence, and it may be hard for a reader of a Gettier case to keep track of a more naïve perspective on a potentially misleading sequence of events [...]" (Nagel, San Juan, & Mar, 2013, p. 659). This is, again, a live option. However, first, as already indicated, Turri's results discussed above show that making participants keep track of the peculiar sequence of events still leaves a big minority of them ascribing knowledge to the agent. Second, my clock-tower scenario, unlike many Gettier cases, didn't describe any shift of the truth-makers of the agent's belief. For instance, in Turri's scenario, Grace has a belief that there's a pen on the table, because she left a pen on the table. Then, this pen falls off the table, a similar pen falls on the table, and thus the belief's truth-maker changes. In contrast, John's belief was caused by what the broken clock was indicating, and thus what caused the belief was never what was making it true. There were no changes in truth-makers to be tracked.²⁴

Let me sum up the worry from performance errors. In most experiments with Gettier cases, there's a significant minority of participants who ascribe knowledge to the agent. The ratio of these respondents doesn't shrink to less than one-fifth, even when the participants are exposed to arguments, guided through the changes in the belief's truth-maker, or asked a follow-up question whether the agent really knows or just thinks she knows. On my preferred interpretation, on these participants' concepts the agent's belief counts as knowledge. However, it's also possible that non-Gettier responses are due to a failure in apprehending the scenario and the arguments. As usual in such situations, more work is needed to evaluate this explanation.

There's a second worry about my interpretation of my results: there might be multiple categorization processes underlying classifications to a category. Following Mallon (forth.) and

²⁴ For the same reason, a different explanation of non-Gettier intuitions, proposed by Nagel et al. in the same paper, doesn't apply to the tower-clock case: "Another possible difficulty has to do with an ambiguity about when the mental state of the agent is to be measured. In many Gettier cases, the agent does know the key proposition at some point in the story but then knowledge is lost due to some unforeseen course of events. Professional epistemologists typically understand that they are being asked to gauge the mental state of the agent at the very end of the case, taking everything into account. A lay reader may not share this assumption, or may simply lose interest early in the story rather than following it through every twist" (Nagel, San Juan, & Mar, 2013, p. 659).

Goldman (2010), I take concepts to be parts or capacities of the mind used in categorization, or, in an even more inclusive manner, capacities used in higher cognitive processes:

Within cognitive science, a concept of x is a body of information about x that is stored in long-term memory and that is used by default in the processes underlying most, if not all, higher cognitive competences when they result in judgments about x , [like] categorization, induction, linguistic understanding, metaphor building, planning, and perhaps other competences (Machery, 2010, p. 401).

Again, on the theory of intuitions as concept applications, an individual who reads a scenario has an intuition when she's uses using the relevant concept to classify the phenomenon described in the scenario. If I tell you the scenario about John, and ask if he knows that it's 12pm, you apply your concept of knowledge, and report whether you think that John mental state counts as knowledge. However, the process producing your answer to my question might be more complex than that. There's a growing consensus in the psychology of concepts that there might be multiple distinct bodies of information underlying higher cognitive competences concerning a given category; what body of information—i.e., concept—is used in a cognitive process depends on the very process and on the context (Machery, 2009; Medin, Lynch, & Solomon, 2000; Sloman, 1996; Smith & Sloman, 1994; Piccini & Scott, 2006; Weiskopf, 2009). For instance, some psychologists argue that, when you categorize a skunk unreflectively, you're using your skunk exemplars (or the prototype), but when the categorization is reflective, you use your implicit theory of skunks. Thus, unreflectively, you'll categorize a black, middle-sized animal with a white stripe on its back as a skunk, for the animal looks similar enough to the skunk exemplars stored in your long-term memory. However, if you learn that the animal is just a raccoon whose back was painted white, you won't call it a skunk, since one of the requirements of your tacit theory of skunks is that a skunk should have skunk parents (Keil, 1989, pp. 179-183).²⁵ Now, it seems possible that introducing an argument before the intuition prompt would make the participants use a different concept than when the story was presented without an argument. So, when the participant reads the story alone, she categorizes John's belief using her exemplars of knowledge; for instance, a non-Gettier intuiter might take this belief to be knowledge, because it's similar to her exemplars of knowledge—it's true and it's justified. However, when primed with an argument, a participant will be more reflective, and thus use her tacit theory of knowledge. And, for instance, this theory can have a rule saying that if a belief is true by accident, it can't count as knowledge. If this is what's really happening, responses to the scenario alone don't reveal the concept that is used when one is considering arguments. Deutsch could then object that the concept relevant for philosophy is the one used in entertaining arguments, and thus studies on intuitions are philosophically inconsequential.

Although it's possible that responses in the no-argument condition were produced by a different cognitive resource than the responses in the two remaining conditions, one would need more evidence for this hypothesis. Since no differences between the conditions were observed, positing a

²⁵ For other studies showing that different categorization mechanisms are activated in different contexts, see (Sloman, 1996; Smith & Sloman, 1994).

shift in concepts between experimental conditions needs more support; it's enough to posit a single concept producing the three responses in one individual to explain the data. What's more, all the arguments used in the study seem to be doing is to highlight a feature of the case that might be relevant for evaluating it. It seems unlikely that such highlighting alone suffices to activate an entirely different categorization mechanism. Moreover, in all three conditions, the participants were asked to ponder the story for the same amount of time, and thus it's not clear why in the first experimental condition they should be less reflective than in the other two conditions. Second, this line of reasoning comes with a risk. If the philosophically relevant concept is the one used in evaluating arguments, but not in producing intuitive responses, then the analysis of knowledge doesn't seem to target what laypersons usually mean when they say that someone knows something. While it might not be a problem for Deutsch, many epistemologists wouldn't be satisfied with this corollary.

Finally, let me repeat the last move made in response to the previous worry—let's assume that different concepts are used when evaluating a case, and that only the one used in evaluating arguments is the one that is relevant for philosophical purposes. Under this assumption, I would agree with Deutsch that studies on intuitions, including experimental philosophy as it's currently done, are much less relevant to standard philosophical problems than their practitioners take them to be. However, it wouldn't yet follow from this conclusion that experimental philosophy is deemed to be philosophically inconsequential. If evaluating arguments uses different concepts than forming intuitions, it doesn't mean that these concepts can't be studied empirically—for instance, the experiment presented here is a (simple and incipient) example of such a study. And its results suggest that, if there's really such a deeper, philosophically relevant concept of knowledge, it may differ in different people, and in a significant minority of them this concept allows for some Gettierized beliefs to count as knowledge.

5. Conclusion

The aim of the study presented here was to show that there is little advantage in invoking arguments rather than intuitions to account for how thought experiments are evaluated. As the empirical results suggest, being exposed to a G-ground doesn't significantly influence a person's evaluation of a case. One plausible explanation of these results is that both intuitions and the persuasiveness of the G-grounds depend on the person's concepts, and therefore a person with a concept of knowledge on which a Gettierized belief counts as knowledge won't be persuaded by arguments that merely highlight the fact that the belief is Gettierized. Therefore, if a difference in epistemic intuitions reveals something about the concept(s) of knowledge underlying these intuitions, and evaluating the G-grounds boils down to applying these concepts, then people whose epistemic intuitions differ from philosophers' won't be, most likely, persuaded by the G-grounds.

The attack focused on Deutsch's defense of the claim that arguments, not intuitions, point to right answers to thought experiments. Yet, a claim similar to Deutsch's has been also defended, at length, by Cappelen (2012). Although Cappelen's argument is more complex than Deutsch's, its crucial part is similar: philosophers use arguments, not intuitions, to evaluate thought experiments. In his book, Cappelen attempts to show that the authors of certain classic thought experiments

evaluate these experiments using arguments. Yet, many of these purported arguments have a structure similar to Deutsch's G-grounds—they point to details of the hypothetical scenario that its author finds relevant for evaluating it. If this is the right way to construe how these authors deal with their thought experiments, the argument presented in this paper poses a challenge to Cappelen's line of reasoning as well. Evaluating thought experiments still remains a matter of intuition, and all his alleged examples of arguments do is to ensure that the intuiter pays attention to certain details of the scenario when intuiting the response to it.

6. References

- Alexander, J., & Weinberg, J. (2007). Analytic epistemology and experimental philosophy. *Philosophy Compass*, 2(1), 56-80.
- Bealer, G. (2000). Theory of the A Priori. *Pacific Philosophical Quarterly*, 81, 1-30.
- Bengson, J. (2013). Experimental Attacks on Intuitions and Answers. *Philosophy and Phenomenological Research*, LXXXVI(3), 495-532.
- BonJour, L. (1998). *In Defense of Pure Reason*. London: Cambridge University Press.
- Brown, J. (2012). Words, Concepts and Epistemology. In J. Brown, & M. Gerken, *Knowledge Ascriptions* (pp. 31-54). Oxford: Oxford University Press.
- Cappelen, H. (2012). *Philosophy without intuitions*. Oxford University Press.
- Cappelen, H. (2014). X-Phi Without Intuitions? In A. Booth, & D. Rowbottom, *Intuitions*. Oxford: Oxford University Press.
- Carey, S. (2009). *The Origins of Concepts*. Oxford: Oxford University Press.
- Cullen, S. (2010). Survey-Driven Romanticism. *Review of Philosophy and Psychology*, 1(2), 275-296.
- Deutsch, M. (2009). Experimental Philosophy and the Theory of Reference. *Mind & Language*, 24(4), 445-466.
- Deutsch, M. (2010). Intuitions, Counter-Examples, and Experimental Philosophy. *The Review of Philosophy and Psychology*, 1, 447-460. doi:10.1007/s13164-010-0033-0
- Deutsch, M. (2015). *The myth of the intuitive*. Cambridge: MIT Press.
- Devitt, M. (2006). *Ignorance of Language*. Oxford University Press.
- Devitt, M. (2006). *Ignorance of Language*. Oxford: Clarendon Press.
- Elga, A. (2007). Reflection and Disagreement. *Nous*, XLI(3), 478-502.
- Feldman, R. (2006). Epistemological Puzzles About Disagreement. In S. H. (edd.), *Epistemology Futures* (pp. 216-236). Oxford: Oxford University Press.
- Foley, R. (1998). Rationality and intellectual self-trust. In M. DePaul, & W. Ramsey, *Rethinking Intuitions*. Lanham, MA: Rowman & Littlefield Publishers.

- Goldman, A. (2007). Philosophical Intuitions: Their Target, Their Source, and Their Epistemic Status. *Grazer Philosophische Studien*, 74, 1-26.
- Goldman, A. (2010). Philosophical Naturalism and Intuitional Methodology. *Proceedings and Addresses of the American Philosophical Association*, 115-150.
- Goldman, A., & Pust, J. (1998). Philosophical theory and intuitional evidence. In M. DePaul, & W. Ramsey, *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*.
- Gopnik, A., & Schwitzgebel, E. (1998). Whose concepts are they, anyway? The Role of Philosophical Intuition in Empirical Psychology. In M. DePaul, & W. Ramsey, *Rethinking Intuition* (pp. 75-91). Lanham.
- Greene, J. (2003). From neural “is” to moral “ought”: What are the moral implications of neuroscientific moral psychology? *Nature Reviews Neuroscience*, 4, 847-850.
- Hales, S. (2006). *Relativism and the foundations of philosophy*. Cambridge, MA: MIT Press.
- Hintikka, J. (1999). The Emperor's New Intuitions. *Journal of Philosophy*, 96(3), 127-147.
- Horvath, J. (2010). How (Not) to React to Experimental Philosophy. *Philosophical Psychology*, 23(4), 447-480.
- Huemer, M. (2005). *Ethical Intuitionism*. Palgrave Macmillan.
- Katz, J. (1981). *Language and Other Abstract Objects*. Totowa, New Jersey: Rowman and Littlefield.
- Kauppinen, A. (2007). The Rise and Fall of Experimental Philosophy. *Philosophical Explorations*, 10(2), 95-118.
- Keil, F. (1989). *Concepts, kinds, and cognitive development*. Cambridge, MA: MIT Press.
- Keil, F. (1994). Explanation, Association, and the Acquisition of Word Meaning. In L. Gleitman, & B. Landau, *The Acquisition of the Lexicon*. Cambridge, MA: MIT Press.
- Kelly, T. (2010). Peer Disagreement and Higher Order Evidence. In A. Goldman, & D. Whitcomb, *Social Epistemology: Essential Readings* (pp. 183-217). Oxford: Oxford University Press.
- Kornblith, H. (1998). The role of intuition in philosophical inquiry: an account with no unnatural ingredients. In M. DePaul, & W. Ramsey, *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*. Rowman and Littlefield.
- Ludwig, K. (2007). The epistemology of thought experiments: First versus third person approaches. *Midwest Studies in Philosophy*, 31, 128-159.
- Machery, E. (2009). *Doing without Concepts*. Oxford: Oxford University Press.
- Machery, E. (2010). Précis of Doing without Concepts. *Philosophical Studies*, 149(3), 401-410.
- Machery, E., Stich, S., Rose, D., Chatterjee, A., Karasawa, K., Struchiner, N., . . . Hashimoto, T. (2015). Gettier Across Cultures. *Noûs*.

- Mallon, R. (forth.). Experimental Philosophy. In H. Cappelen, T. S. Gendler, & J. Hawthorne, *The Oxford Handbook of Philosophical Methodology*. Oxford University Press.
- Medin, D., & Schaffer, M. (1978). Context Theory of Classification Learning. *Psychological Review*, 85(3), 207-238.
- Medin, D., Lynch, E., & Solomon, K. (2000). Are there kinds of concepts? *Annual Review of Psychology*, 51, 121-147.
- Murphy, G., & Medin, D. (1985). The Role of Theories in Conceptual Coherence. *Psychological Review*, 92(3), 289-316.
- Nadelhoffer, T., & Nahmias, E. (2007). The past and future of experimental philosophy. *Philosophical Explorations*, 10(2), 123-149.
- Nagel, J. (2007). Epistemic Intuitions. *Philosophy Compass*, 2(6), 792-819.
- Nagel, J. (2012). Intuitions and Experiments: A Defense of the Case Method in Epistemology. *Philosophy and Phenomenological Research*, 85(3), 495-527.
- Nagel, J. (2012). Mindreading in Gettier Cases and Skeptical Pressure Cases. In J. Brown, & M. Gerken, *Knowledge Ascriptions* (pp. 171-191). Oxford: Oxford University Press.
- Nagel, J., San Juan, V., & Mar, R. (2013). Lay Denial of Knowledge for Justified True Beliefs. *Cognition*, 129(3), 652-661.
- Nosofsky, R. (1988). Exemplar-Based Accounts of Relations Between Classification, Recognition, and Typicality. *Journal of Experimental Psychology*, 14(4), 700-708.
- Piccini, G., & Scott, S. (2006). Splitting concepts. *Philosophy of Science*, 73, 390-409.
- Plantinga, A. (1993). *Warrant and Proper Function*. New York: Oxford University Press.
- Pust, J. (2000). *Intuitions as Evidence*. New York: Routledge.
- Rosch, E. (1973). Natural categories. *Cognitive Psychology*, 4, 328-350.
- Rosch, E. (1978). Principles of categorization. In E. Rosch, & B. Lloyd, *Cognition and Categorization* (pp. 189-206). Hillsdale, NJ: Lawrence Erlbaum.
- Rosch, E., & Mervis, C. (1975). Family resemblances: studies in the internal structure of categories. *Cognitive Psychology*, 8: 382-439, 8, 382-439.
- Scholl, B. (2007). Object Persistence in Philosophy and Psychology. *Mind & Language*, 22(5).
- Schulz, E., Cokely, E., Feltz, & Adam. (2011). Persistent Bias in Expert Judgments About Free Will and Moral Responsibility: A Test of the Expertise Defense. *Consciousness and Cognition*, 20, 1722-31.
- Schwitzgebel, E., & Cushman, F. (2012). Expertise in Moral Reasoning? Order Effects on Moral Judgment in Professional Philosophers and Non-philosophers. *Mind and Language*, 27, 135-53.

- Seyedsayamdost, H. (forth.). On Normativity and Epistemic Intuitions: Failure of Replication. *Episteme*. Retrieved from ssrn.com/abstract=2538611
- Sloman, S. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, 119(1), 3-22.
- Smith, E., & Sloman, S. (1994). Similarity-versus rule-based categorization. *Memory & Cognition*, 22(4), 377-386.
- Sosa, E. (2007). Intuitions: Their Nature and Epistemic Efficacy. In C. Beyer, & A. Burri, *Grazer Philosophische Studien. Philosophical Knowledge* (pp. 51-67).
- Sosa, E. (2009). A Defense of the Use of Intuitions in Philosophy. In M. Bishop, & D. Murphy, *Stich and His Critics* (pp. 101-112).
- Starmans, C., & Friedman, O. (2012). The folk conception of knowledge. *Cognition*, 124(3), 272-283.
- Steup, M. (1996). *An Introduction to Contemporary Epistemology*. Upper Saddle River, New Jersey: Prentice-Hall.
- Stich, S. (1988). Reflective equilibrium, analytic epistemology and the problem of cognitive diversity. *Synthese*, 74(3), 391-413.
- Talbot, B. (2013). Reforming intuition pumps: when are the old ways the best? *Philosophical Studies*, 165(2), 315-334.
- Tobia, K., Buckwalter, W., & Stich, S. (2013). Moral Intuitions: Are Philosophers Experts? *Philosophical Psychology*, 26, 629-38.
- Tobia, K., Chapman, G., & Stich, S. (2013). Cleanliness Is Next to Morality, Even for Philosophers. *Journal of Consciousness Studies*, 20, 195-204.
- Turri, J. (2013). A conspicuous art: putting Gettier to the test. *Philosophers' Imprint*, 13(10).
- Turri, J., Buckwalter, W., & Blouw, P. (2015). Knowledge and Luck. *Psychonomic Bulletin and Review*, 22(2), 378-390.
- Tymoczko, T. (1979). The Four-Color Problem and Its Philosophical Significance. *The Journal of Philosophy*, 76(2), 57-83.
- 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), 429-60.
- Weiskopf, D. (2009). The plurality of concepts. *Synthese*, 169(1), 145-173.

