A conspicuous art: putting Gettier to the test[*]

JOHN TURRI
john.turri@gmail.com

> *Of heaven, not me, make an experiment.*
> *I am not an impostor that proclaim*
> *Myself against the level of mine aim;*
> *But know I think and think I know most sure*
> *My art is not past power nor you past cure.*
> — *Helena (All's Well That Ends Well, 2.1.763–67)*

Abstract: Professional philosophers say it's obvious that a Gettier subject does not know. But experimental philosophers and psychologists have argued that laypeople and non-Westerners view Gettier subjects very differently, based on experiments where laypeople tend to ascribe knowledge to Gettier subjects. I argue that when effectively probed, laypeople and non-Westerners unambiguously agree that Gettier subjects do not know.

1. Introduction

A Gettier case features a protagonist who has a justified true belief, but who, according to prevailing philosophical wisdom, obviously lacks knowledge (Gettier 1963). Among professional philosophers,

---

there is a virtual consensus that a Gettiered subject does not know.[1] Call this *the Gettier intuition.*

Experimental philosophers and psychologists have recently argued that laypeople don't share the Gettier intuition. When surveyed, laypeople reveal a tendency to *ascribe knowledge* to Gettier subjects. Studies have been reported on laypeople from Western and non-Western cultures (Starmans and Friedman 2012, Weinberg, Nichols and Stich 2001, Cullen 2010).

If (a) philosophers and laypeople are talking about the same thing when they speak of 'knowledge', and (b) both groups competently assess Gettier cases, and (c) both frankly report their judgments, then the two groups will respond similarly when asked whether the Gettier subject knows. Thus, if the two groups respond differently, then either (a), (b) or (c) is false. And the experimentalists have provided evidence that the groups do respond differently. So at least one of (a), (b) or (c) is probably false. In light of the experimental findings, all three options have been explored.

Some accept that philosophers and laypeople are talking past one another. It's a mere verbal disagreement (as suggested by Sosa 2007, 2009). Others accept that one of the two groups isn't competently assessing the cases, and since philosophers are trained experts at assessing thought experiments, whereas laypeople are amateurs, the laypeople are probably at fault (Ludwig 2007, Williamson 2007, 2011). Philosophical training makes one better at noticing import-

---

[1] Sartwell 1991, Turri 2012a, and Hetherington forthcoming dissent. For a review of some recent attempts to solve the Gettier problem, see Turri 2012b.

ant but subtle details, which explains the difference in performance. Others accept that one of the groups isn't frankly reporting their judgment about the cases.[2] Since philosophers ostensibly set aside purely practical concerns and study these cases in order to reach the truth about some important question, whereas laypeople's verbal performances remain sensitive to a broader range of practical concerns, it's more likely that laypeople aren't frankly reporting their judgment.

This paper provides resources that strengthen the expertise response to the experimental results: laypeople who answer that the Gettier subject knows aren't competently enough assessing the case. Proponents of the expertise response incur an empirical commitment. If the expertise response is correct, then philosophers notice and assign proper weight to features of the cases that untutored laypeople overlook. Accordingly, we would expect that if we effectively guide participants to notice and assign proper weight to those same features, then they will respond similarly to philosophers. That is, they will say that the Gettier subject doesn't know.

To that end, I propose a simple and natural technique for effectively guiding participants to competently assess Gettier cases, and I report a series of experiments that demonstrate its effectiveness. Section 2 introduces the technique. Section 3 discusses some previous experimental results which don't use the technique. Sections 4–9 report new experimental findings gained by using the technique. Section 10 briefly concludes the discussion.

---

[2] I'm exploring this possibility in ongoing work.

2. *Conspicuous tripartite structure*

One very effective way to think about Gettier cases is *in three stages*. Start with a belief that is well enough justified to satisfy the justification condition on knowledge. All seems well. Then introduce an element of bad luck that would normally prevent the justified belief from being true. All seems ill. Then introduce a conspicuously distinct element of good luck that makes the belief true anyway (Zagzebski 1996). But not all is made well again. This is how I find myself thinking about Gettier cases (see Turri 2011), and it has proven very effective in discussing matters with undergraduates and laypeople more generally. In my experience, when the case is dramatized this way, the Gettier intuition is powerful and widely shared.

But this is mere anecdote. What will happen if we guide participants to think of the case this way too, dramatizing the conspicuously distinct elements of luck via a tripartite structure? One way to do this is to physically present the story in three distinct stages, on separate pages or screens, and ask participants to keep track of the (apparent) truth of the target proposition in light of new information at each stage. I predict that when probed this way, participants will not ascribe knowledge to the Gettier subject.

No previous experimental work on Gettier cases adopts this approach. In a word, and vividly: where my predecessors used clubs, I recommend a trident. The next section reviews two previous studies to demonstrate how my proposed technique differs.

3. Two previous studies

The most recent empirical challenge to the Gettier intuition comes from Christina Starmans and Ori Friedman (2012). In a series of impressive experiments, they observed that laypeople consistently ascribe knowledge to Gettier subjects at rates exceeding chance.[3] Here is one of their cases.

> Katie is in her locked apartment writing a letter. She puts the letter and her blue Bic pen down on her coffee table. Then she goes into the bathroom to take a shower. As Katie's shower begins, two burglars silently break into the apartment. One burglar takes Katie's blue Bic pen from the table. But the other burglar absentmindedly leaves his own identical blue Bic pen on the coffee table. Then the burglars leave. Katie is still in the shower and did not hear anything.

This is clearly a Gettier case. But it is not presented in stages, and the respective sources of bad and good luck are not conspicuously distinct. When asked whether Katie 'really knows' or 'only thinks' that there is a pen on the table, participants selected 'really knows' at rates exceeding chance (69%). Participants also reported being highly confident in their judgment. Rates of knowledge attribution

---

[3] Jackson (2011: 469) identifies several responses to surveys suggesting that the folk think that Gettier subjects know. The first response is to 'insist that the surveys . . . are defective,' on the grounds that they 'violate one or another principle of good social science polling, and maybe use Gettier cases that aren't among those found most compelling by analytic philosophers.' But this description certainly doesn't apply to Starmans and Friedman's work, which displays a level of methodological rigor typical of the best experimental social science and features paradigmatic Gettier cases.

and reported confidence didn't differ significantly between participants who read Gettier cases and participants who read non-Gettier control cases.

Participants answered a series of comprehension questions before answering the key test question about whether the Gettier subject knows. This eliminated data from participants who weren't reading the story carefully. Moreover, Starmans and Friedman claim, the questioning encouraged participants 'to closely approximate the logical steps that philosophers have deemed relevant for the attribution of knowledge'; it 'forced' them to 'reason through the relevant factors' of the case. In particular, the questions 'highlighted' the typical 'disconnect' between justification and truth, and the typical 'luck element' distinctive of Gettier cases (2012: 5–6, 10). However, Starmans and Friedman note, this 'extra guidance' didn't prevent laypeople from mostly attributing knowledge (2012: 6).[4] Nevertheless, they accept that it's important to effectively guide participants to think through the relevant details.

I replicated Starmans and Friedman's findings from this study. Participants (N = 28)[5] were tested using Starmans and Friedman's

---

[4] Starmans and Friedman don't conclude that the folk concept of knowledge is justified true belief. For they also observed that the folk *declined* to ascribe knowledge when the subject's evidence was merely 'apparent', even when the belief was true and justified. Apparent evidence is 'evidence that *appears* to be informative about reality, but is not really' (2012: 9). By contrast, 'authentic evidence' is, roughly, evidence that makes the belief true when based on it.

[5] Thirteen female, aged 18–63, M = 32.5, SD = 10.47. As with the experiments reported below, participants were recruited using Amazon Mechanical Turk and compensated $.30 for approximately 2–3 minutes of their time. Ninety-six percent reported English as their native language. They

original materials. A majority (57%) answered that Katie 'really knows' that there is a pen on the coffee table. This does not differ significantly from the original study's results (69%, binomial, $p = 0.25$),[6] though neither does it differ significantly from what would be expected by chance (binomial, $p = 0.572$). This pattern of response clearly differs from the alleged philosophical consensus. Starmans and Friedman also report results from a *weighted knowledge ascription*, which is the product of the answer to the dichotomous knowledge question (really knows = 1; only thinks = -1) and the reported confidence (1–10, low to high). Scores for this measure could range from -10 (fully confident knowledge denial) to 10 (fully confident knowledge ascription). Starmans and Friedman observed a mean weighted knowledge ascription that exceeded chance in Katie's case ($M = 3.92$, $SD = 8.21$). I observed a mean weighted knowledge ascription that did not differ significantly either from their results ($M = 1.5$, $SD = 9.31$, $t(27) = -1.375$, $p = 0.18$) or from chance ($t(27) = 0.852$, $p = 0.401$). I conducted two follow-up studies and continued to observe the same basic pattern, further reinforcing Starmans and Friedman's findings.[7]

---

filled out a brief demographic survey after testing. I excluded data from two participants who failed comprehension questions. Including data from these participants doesn't significantly affect the results reported below. Participants were not allowed to re-take any survey, and participants who had taken previous similar surveys were excluded by their AMT Worker ID. Except for the experiment discussed in section 7, all participants were located throughout the United States.

[6] All tests are two-tailed, unless otherwise noted.

[7] Following up on a couple suggestions from Frank Jackson (p.c.), I ran two slightly modified surveys involving Katie's case. In the one version (N=29), time-indexing was added to the test question itself, to make sure that parti-

In an older influential study, Jonathan Weinberg, Shaun Nichols, and Stephen Stich (2001) presented results from a survey conducted on undergraduates at Rutgers. Their sample population included participants of Indian, Pakistani and Bangladeshi descent. For purposes of analysis, Weinberg, Nichols and Stich treated these participants as a single group that could tell us something about 'the epistemic intuitions of people from the Indian sub-continent (hereafter SCs)' (2001: 443). Weinberg, Nichols and Stich had participants (*N*=23) read a typical Gettier case.

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

The case was presented all at once, not in stages. Weinberg, Nichols and Stich don't report asking any comprehension questions. Sixty-

cipants were ascribing knowledge at the relevant part of the story: as Katie is finishing up her shower, rather than, say, after she walks back out into the living room. Sixty-five percent of participants in this condition ascribed knowledge. In the other version (*N* = 22), the burglar's blue Bic pen was *not* described as 'identical' to Katie's. Forty-one percent of participants in this condition ascribed knowledge. Overall, rates of knowledge ascription didn't differ significantly across the three versions. Overall for the replication and the two minor variations (*N* = 79), 56% of participants ascribed knowledge. And although this differs significantly from 69% (binomial, *p* = 0.018), it doesn't differ significantly from chance (binomial, *p* = 0.368). What seems most important in all of this is that there's no sign here that the folk think it's obvious that the Gettier subject doesn't know.

one percent of their participants said that the Gettier subject 'really knows'. (I return to this experiment in Section 7.)

4. A staged encounter

This section reports an experiment that illustrates my alternative tripartite approach to Gettier surveys.

4.1. Method

Participants (N = 52)[8] were randomly assigned to one of two conditions: Control and Authentic Gettier.[9] All participants read a story

---

[8] Eighteen female, aged 18–59, M = 27.58, SD = 8.78. One-hundred percent reported English as their native language. I excluded data from eight participants who failed comprehension questions.

[9] '*Authentic* Gettier' because it features authentic evidence, in Starmans and Friedman's sense of that term (see n. 4 above). Starmans and Friedman observed that participants were much more likely to ascribe knowledge in Gettier cases involving authentic evidence than in Gettier cases involving apparent evidence. The cases used for SC Tripartite and Zebra feature apparent evidence. Some of the results reported below further support Starmans and Friedman's hypothesis; see n. 30.

For the record, I think that Starmans and Friedman are right that laypeople tend to view the cases as different, with more ascribing knowledge when the evidence is authentic. I take their findings to establish the following: in an identifiable (even if somewhat fuzzy) range of cases where the protagonist perceptually detects the truth and justifiably continues to believe that truth over a short period of time, the default position for a lot of people – perhaps a small majority – seems to be that the protagonist continues to know, whether or not she's Gettiered. Waxing figuratively for just a moment, we might put the point this way: at least in the short term, the epistemic inertia of initially successful perception can, to some extent, inhibit Gettierization from having the expected effect on lay judgment. A related insight finds expression in the professional literature, where many philosophers claim that fake-barn cases, which involve successful perception, are cases of knowledge and differ importantly from standard Gettier cases (see Turri 2012a for discussion and references).

In a forthcoming paper accepted after this paper was accepted for pub-

in three stages. The *first* and *third* stages were the same in both conditions; but the critical *second* stage differed. Here is the story, with the different second stages clearly marked.

Stage One: Robert recently made a purchase for a rare 1804 US silver dollar. He keeps the coin on display over the fireplace in his library. This evening Robert is having his neighbors over for dinner. He puts the coin in its display over the fireplace, shuts the library doors behind him, and hurries to greet his guests, who just arrived. He greets them and says, "Guess what? There is an 1804 US silver dollar in my library."

Stage Two – Authentic Gettier: When Robert shut the library doors, a coin thief silently entered through the library window, stole Robert's 1804 US silver dollar, and quickly escaped. Robert had only been out of the library for a few seconds and did not hear anything. The coin was already gone by the time Robert greets his guests and tells them, "There is an 1804 US silver dollar in my library."

Stage Two – Control: When Robert shut the library doors, the vibrations from the door shutting caused the silver dollar to fall from its display and land on the rug near the fireplace.

lication, Jennifer Nagel, Valerie San Juan and Raymond Mar observed different results for cases that they claim are essentially similar to Starmans and Friedman's. (The relevant experimental stimuli weren't publicly available as of this writing.) But Nagel et al. collected and analyzed responses very differently from how Starmans and Friedman did, so a direct comparison of the disparate results is fraught.

Robert had only been out of the library for a few seconds and did not hear anything. The coin had already fallen onto the library floor by the time Robert greets his guests and tells them, "There is an 1804 US silver dollar in my library."

Stage Three: Robert's house is a very old mansion. Back in the early 1800s, when the house was originally built, one of the carpenters accidentally, and without noticing, dropped an 1804 US silver dollar into the mortar mix used to make the fireplace. This lost silver dollar is still in the fireplace in the library. But no one has seen it for hundreds of years, and no one will ever see it again. It will remain hidden in Robert's library.

Each stage appeared on a different screen.[10] Participants in each condition were asked one comprehension question at each stage. It was always the same comprehension question (options in brackets):

When Robert greets his guests, is there an 1804 US silver dollar in his library? [Yes/No][11]

---

[10] NB: I don't claim that physically breaking up the story on separate pages or screens is *necessary* to effectively guide participants. I claim only that it will be, for the most part, sufficient. Effective tri-partitioning might be accomplished all at once. I am pursuing this possibility in ongoing work, but results reported in section 9 — *viz.* the rate of knowledge ascription in 1-Gettier compared to 3-Gettier — suggest that it will be very difficult.

[11] Options were rotated randomly for all questions in all experiments reported here, except for confidence measures, which were always ordered 1–10 ('not at all confident' to 'completely confident'). At the end of the second stage in the Authentic Gettier condition, participants who answered 'no' to the comprehension question passed. Even though this turns out to be false in the story, participants paying close attention will be led to this reasonable but false belief about the story at this juncture.

After the story was complete, participants were then asked the test question:

> When Robert greets his guests, he _____ that there is an 1804 US silver dollar in his library. [really knows/only thinks he knows]

Participants were then asked to report how confident they were in their answer to the test question (1–10, low to high).

The motivation for setting the case up this way is simple: it dramatizes the tripartite structure of Gettier cases, guiding participants to notice the intersection of evidence, truth and luck, and highlighting that the bad luck's source differs conspicuously from the good luck's source. This is accomplished by explicitly dividing the story into three stages and asking participants to keep track of the truth of the key proposition in light of new information.

4.2.Results and discussion

I made four predictions about the results. First, there would be an effect of condition. Second, participants in Control would say that Robert 'really knows' at rates exceeding chance. Third, participants in Authentic Gettier would strongly tend to say that Robert 'only thinks he knows'. Fourth, there would be no effect of condition on reported confidence.

All four predictions were correct. There was a dramatic effect of condition on the dichotomous test question (Fisher's exact test, $p <$ 0.001, one-tailed);[12] participants in Control overwhelmingly said

---

[12] I use one-tailed tests here because the relevant predictions were *direc-*

that Robert 'really knows' (84%, binomial, $p = 0.002$, one-tailed); participants in Authentic Gettier overwhelmingly said that Robert 'only thinks he knows' (89%, binomial, $p < 0.001$, one-tailed); and there was no effect of condition on reported confidence ($\chi^2(6) = 7.181$, $p = 0.304$).[13] (See Figure 1.)

Results from the weighted knowledge ascription tell a similar story (Figure 1).[14] A one-way analysis of variance (ANOVA) revealed a dramatic effect of condition ($M = 6.76/-6.44$, $F(1,50) = 61.66$, $p < 0.001$); weighted knowledge attribution in Control far exceeded what could be expected by chance ($t(24) = 5.348$, $p < 0.001$); and weighted knowledge denial in Authentic Gettier far exceeded what could be expected by chance ($t(26) = -5.767$, $p < 0.001$). These are exactly the results we would expect if the tripartite procedure effectively guides participants to assess Gettier cases.

The tripartite structure itself doesn't lead participants to deny knowledge, or else they would do so in the Control condition too. The occurrence of unexpected events related to the truth of the belief in question — such as the coin falling due to vibrations from slamming doors, or a long-lost rare coin of precisely the same sort being hidden in the fireplace — doesn't lead participants to deny

---

tional. However, the results are still significant even when using two-tailed tests.

[13] $M = 8.92/8.44$, $SD = 2.16/1.67$. There was no effect of condition on confidence in the other experiments reported either.

[14] Following Starmans and Friedman (2012), I define a weighted knowledge ascription as the product of the answer to the dichotomous knowledge question (really knows = 1; only thinks he knows = -1) and the reported confidence (1–10, low to high). Scores for this measure could range from -10 (fully confident knowledge denial) to 10 (fully confident knowledge ascription).

knowledge, or else they would do so in the Control condition too. Repeatedly asking about the truth of the relevant proposition doesn't lead participants to deny knowledge, or else they would do so in the Control condition too.

One final thought is that switching the contrast from 'only thinks' or 'only believes', as in the previous studies discussed in section 3, to 'only thinks he knows' might have had an effect. I ran a follow-up study to test the conjecture. Participants (N = 21)[15] in the Only Believes Gettier condition received the same story and questions as the Authentic Gettier condition, except for one difference: for the key test question, 'really knows' was contrasted with the original 'only believes' instead. But the results didn't differ significantly from Authentic Gettier: 81% selected 'only believes' (Fisher's, p = 0.683), and the mean weighted knowledge ascription was -5.71 ($F(1,46)$ = .171, p = 0.682). (See Figure 1.) [16]
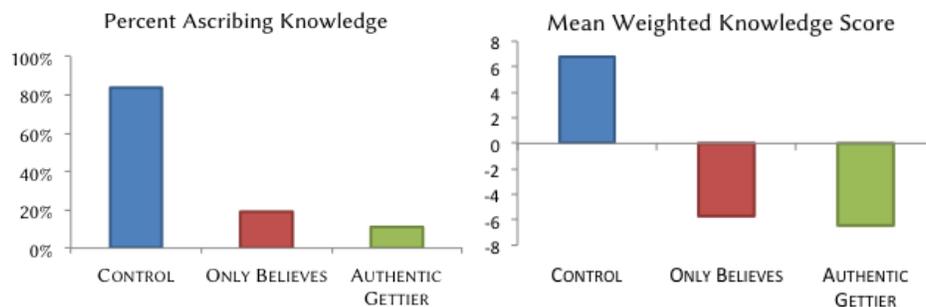

Figure 1: Left Panel: percentage of participants attributing knowledge. Right Panel: mean weighted knowledge ascription (derived by multiplying the dichotomous an-

---

[15] Ten female, aged 18–53, M = 30.9, SD = 10.8. Ninety-six percent reported English as their native language. I excluded data from nine participants who failed comprehension questions.

[16] For those worried that Control doesn't closely enough match Authentic Gettier, see section 9.

swer by reported confidence).

5. *Zebra and mule*

If I'm right that the conspicuous tripartite structure effectively guides laypeople, then participants will respond similarly to other Gettier cases presented that way, even when the cover story is very different.[17] For example, it shouldn't matter whether the protagonist forms a belief about an artifact such as a coin, or an animal such as a zebra. This section reports an experiment that tests this prediction.

5.1. *Method*

Participants (*N* = 24)[18] in the Zebra condition read the following story.

> *Stage One*: Zach has an appointment with his lawyer in an office building in New York City. As he enters the lobby on the first floor, he sees something highly unexpected: a large animal with black and white stripes under a banner that says, "Pet a zebra for children's charity." In exchange for a $10 donation to a local children's charity, you get to pet this

---

[17] Here I set aside cover stories that might trigger 'the epistemic side-effect effect' (Beebe and Buckwalter 2010, Buckwalter forthcoming, Beebe and Shea ms.; see also Turri 2012a). Pilot data from work in progress suggest that the present approach inhibits the epistemic side-effect. But responsibly addressing this issue requires more time and space than I can devote to it here.

[18] Seven female, aged 18–54, *M* = 28.42, *SD* = 10. Ninety-five percent reported English as their native language. I excluded data from six participants who failed comprehension questions.

illustrious animal. Zach quickly walks up the stairs to his lawyer's office on the second floor. He greets the receptionist and says, "Guess what? There is a zebra down on the first floor of the building."

Stage Two: The people running the charity could not afford to rent a real zebra for the charity drive. So instead they hired an artist to paint black and white stripes on a mule. The animal Zach saw in the first-floor lobby was actually a cleverly disguised mule. It looks just like a zebra, but it isn't. It's a mule.

Stage Three: The office building that Zach is in is very large. One of the companies renting space in the building is in the business of importing exotic animals. It is illegal for them to keep these animals in an office building, but they do it anyway. They recently acquired a zebra and are keeping it well-hidden in a locked, sound-proof room on the first floor of the building.

At the end of each stage, participants were asked a single comprehension question:

When Zach greets the receptionist, is there a zebra on the first floor of the building? [Yes/No].

Then they were asked the test question:

When Zach greets the receptionist, he _____ that there is a zebra on the first floor of the building. [really knows/only thinks he knows.]

Participants were asked to report how confident they were about their answer to the test question (1–10, low to high). As before, each stage was presented on a different screen. Participants couldn't go back and change answers.

5.1.1. Results and discussion

I predicted that participants in Zebra would overwhelmingly respond that Zach 'only thinks he knows'. The prediction was true: 100% of participants answered that Zach 'only thinks he knows'.[19] This supports my proposal that the conspicuous tripartite structure effectively guides participants.

6. Burglar and husband

If I'm right that the conspicuous tripartite structure effectively guides participants, then we should be able to take Gettier cases from previous experiments and transform them so that participants deny that the Gettier subject knows. For example, consider the original case involving Katie (section 3), in response to which participants tend to ascribe knowledge. We should be able to take that case and transform it stepwise in order to effectively guide participants, whereupon they will tend to deny knowledge. This section reports an experiment designed to test this prediction.

6.1.1. Method

Participants (N = 46)[20] were randomly assigned to one of two condi-

---

[19] Even including the participants who failed a comprehension question, 100% responded this way.

[20] Thirteen female, aged 18–59, M = 27.33, SD = 8.39. One hundred percent

tions: Burglar and Husband. The stories for these conditions were based on Starmans and Friedman's story involving Katie. The stories shared the *first* and *second* stages in common, but differed in the *third* stage. Here are the stories, with the different third stages clearly marked.

Stage One: Katie is 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. Then she goes into the bathroom to take a shower. It takes her fifteen minutes to finish.

Stage Two: Just after Katie started her shower, two burglars, a master and his apprentice, broke into her apartment. As they made their way around the apartment, the master burglar stole Katie's blue Bic pen from the coffee table. After five minutes, the burglars left, well before Katie finished her shower. Katie did not hear anything.

Stage Three – Burglar: Right before the burglars left Katie's apartment, the apprentice burglar started feeling a bit dizzy, so he sat down on the couch for a moment to recover. When the apprentice burglar sat down, he absentmindedly set his own blue Bic pen on the coffee table, and forgot it there. This was five minutes before Katie finished her shower.

Stage Three – Husband: Right after the burglars left, Katie's

reported English as their native language. I excluded data from thirteen participants who failed comprehension questions.

husband came home. Tired from a long journey, he put his wallet, keys and his own blue Bic pen down on the coffee table in the living room. Then he immediately lay down on the living room couch and fell asleep. This was five minutes before Katie finished her shower. Katie hasn't yet noticed that her husband is even home.

The comprehension question asked at the end of each stage was:

As Katie is in the bathroom finishing her shower, is there a blue Bic pen on her coffee table? [Yes/No]

The test question at the end was:

As Katie is in the bathroom finishing her shower, she _____ that there is a blue Bic pen on her coffee table. [really knows/only thinks she knows]

Participants were then asked to report how confident they were about their answer to the test question (1–10, low to high). Each stage was presented on a different screen. Participants couldn't go back and change answers.

The earlier replication of Starmans and Friedman's study (section 3), which used the original case of Katie, served as the Control for this experiment.

6.1.2. Results and discussion

The materials for Burglar and Husband are adapted from the original case of Katie. The goal is to see whether incrementally modifying it to fit the conspicuous tripartite structure will result in a corresponding decrease in knowledge attributions.

Katie's original story was presented all at once, and the nondescript pair of burglars was the source of both the bad and good luck: one burglar took a blue Bic pen, while the other left a pen that effectively replaced it. The story for Burglar is substantially similar, except that it is rewritten to occur in three stages and to achieve some separation between the two burglars: the master burglar steals the pen, while the apprentice burglar leaves one that effectively replaces it. Still, the burglars operate as a team, so the separation between the source of bad and good luck could be more conspicuous. The story for Husband achieves much clearer separation between the source of bad and good luck: the burglars steal the pen, but this time Katie's husband, who has nothing to do with the burglars, replaces it. All things considered, the story for Husband best fits the conspicuous tripartite structure, Control least well fits it, and Burglar falls somewhere in between.

Accordingly, I made three predictions. First, Control would see the highest rates of knowledge ascription, followed by Burglar, then Husband. Second, although I wasn't sure whether response rates in Burglar would differ significantly from those in Control, I predicted that response rates in Husband would. Third, I predicted that participants in Husband would select 'only thinks she knows' at rates greater than chance. Finally, I predicted that rates of knowledge ascription in Husband wouldn't differ significantly from those observed in Authentic Gettier (section 4).

The predictions were all true. First, the relative rates of knowledge ascription were Control 57%, Burglar 44%, and Husband

24%. Second, rates of knowledge ascription in Husband differed significantly from Control (Fisher's p = 0.04). Third, participants in Husband selected 'only thinks she knows' at rates greater than chance (76%, binomial, p = 0.028). Finally, rates of knowledge ascription in Husband didn't differ significantly from those in Authentic Gettier, for either the dichotomous (Fisher's, p = 0.272) or weighted ascription ($F(1,46)$ = .983, p = 0.327). These results provide further support for the effectiveness of conspicuous tri-partitioning.

7. *Subcontinental drift*

As mentioned earlier, previous results suggested that people from the Subcontinent ('SCs') tend to ascribe knowledge to Gettier subjects (Weinberg, Nichols and Stich 2001). If the conspicuous tripartite structure effectively guides participants, then it should do so whether the participants are from North America or the Indian Subcontinent. This section reports an experiment that tests this prediction.

7.1. *Method*

Participants (N = 27)[21] in the SC Tripartite condition were recruited

---

[21] Ten female, aged 20–60, M = 33.6, SD = 13.05. Participants listed the following native languages (with three listing two native languages): English (44%), Tamil (30%), Hindi (18.5%), Malayalam (15%), Marathi (3.7%). I excluded data from fifty-four participants who failed comprehension questions. With three dichotomous comprehension questions, each with a unique correct answer, 27 (33.3%) is also almost three times as many as would be expected by chance to pass all three checks (≈10, or 12.5%), which is significant (binomial, $p<.001$, one-tailed). Whether a participant failed a

using Amazon Mechanical Turk and located throughout India, the heart of the Subcontinent. They read a story similar to the one from Authentic Gettier, except that it had a different second stage:

> *Stage Two – SC Gettier*: What Robert doesn't realize is that the coin dealer cheated him. The coin Robert brought home from the coin shop, and which is currently on display in his library, is a very convincing forgery. It's not a real 1804 US silver dollar. Robert isn't aware of this as he greets his guests and tells them, "There is an 1804 US silver dollar in my library."

Participants in SC Tripartite were questioned in exactly the same way as participants in Authentic Gettier.

7.2. Results and discussion

I made two predictions. First, participants in SC Tripartite would ascribe knowledge at rates significantly lower than Weinberg, Nichols and Stich observed in their original study (hereafter 'SC Original'). Second, response rates in SC Tripartite would not differ significantly from those observed in Authentic Gettier.

Both predictions were true. First, only 15% said that Robert

---

comprehension question was correlated with whether they ascribed knowledge (r = -.278, N = 81, p = 0.012). Overall, 42.6% of participants who failed at least one comprehension question ascribed knowledge, compared to only 14.8% of participants who passed all comprehension questions. And 67% of all participants, including those who failed comprehension questions, said that Robert 'only thinks he knows', which far exceeds what could be expected by chance (binomial, p = 0.016); and only 33% said that Robert 'really knows', which differs significantly from what was observed in SC Original (61%, binomial, p < 0.001).

'really knows', which differs significantly from SC Original (61%, bi-nomial, $p < .001$). Second, response rates didn't differ significantly from Authentic Gettier for either the dichotomous (15% ascribed knowledge, Fisher's, $p = 1$) or the weighted knowledge ascription ($M = -5.89$, $SD = 6.14$, $F(1,52) = .117$, $p = 0.734$).

I grant that surveys conducted in participants' native languages on the Subcontinent would be preferable. Nevertheless, this survey was conducted in an official language of India, English, which is used extensively in commerce and public administration. Moreover, the results fit what was predicted and resemble the results observed in other studies reported above. Accordingly — and bearing in mind, of course, that we shouldn't draw any firm conclusions based on a single study of this size — these results should lead us to recon-sider earlier claims that the Gettier intuition isn't shared on the Subcontinent. This is certainly fitting because in the Indo-Tibetan philosophical tradition, Gettier-style cases were discussed along the banks of the Ganges well over a thousand years before Edmund Gettier published his paper from the banks of the Detroit River.[22]

8. A persistent pattern

A persistent pattern has emerged (Figure 2). Response rates across the five Gettier conditions with conspicuous tripartite structures did not differ significantly for either the dichotomous ($\chi^2(4) = 6.447$, $p = 0.168$) or the weighted knowledge ascription ($F(4,115) = 1.48$, $p =$

---

[22] See Matilal 1986: 135–7 and Stoltz 2007: 298. Stoltz doesn't think these are technically Gettier cases.

0.231). This is exactly what we would expect if the conspicuous tri-partite structure effectively guides participants.
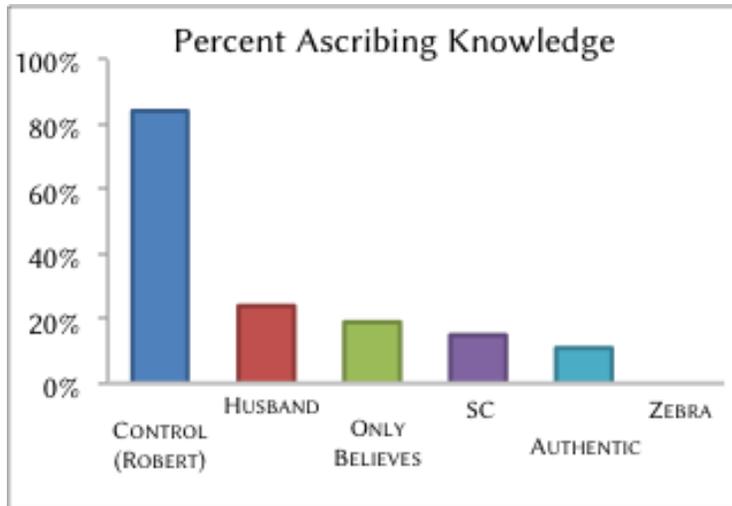

Figure 2: Comparison of the percentage of participants who ascribed knowledge across the five Gettier conditions thus far, alongside the Control (section 4).

Across all five Gettier conditions, gender had no effect on rates of dichotomous (13.5/13%, Fisher's, $p = 1$) or weighted knowledge ascriptions ($M = -6.46/-6$, $SD = 6.05/5.74$, $F(1,118) = .17$, $p = 0.681$). Age had no effect on rates of dichotomous ($\chi^2(1) = .288$, $p = 0.789$) or weighted knowledge ascriptions either ($F(1,118) = .693$, $p = 0.407$).

9. *Partition and source*

Getting participants to appreciate distinct strokes of luck is easier when the strokes emanate from conspicuously distinct sources. The stories I have used dramatize this separation by having the second stage focus on the bad luck, while the third stage focuses on the

good luck. In Husband (section 6), a thief was the source of bad luck, and the husband was the source of good luck. In Authentic Gettier (section 4), a thief was the source of bad luck, and a carpenter's mishap hundreds of years ago was the source of good luck.

But one might object that by introducing a different source of luck in, say, the second stage of Authentic Gettier as compared to Control, or by switching the source of good luck in Husband from the thieves to the husband, I have introduced elements that are, strictly speaking, *unnecessary* for Gettierization.[23] And other differences inevitably creep in as a result of changing sources. These changes arguably go beyond the minimal change needed to turn a case of knowledge into a Gettier case. More to the point, they might go beyond the minimal change needed to turn what laypeople judge is a case of knowledge into what laypeople can be counted on to recognize as a case involving Gettierization.

This objection leaves me ambivalent. On the one hand, absolutely minimally matched pairs are always good if you can get them. The more minimally matched the pair, the more persuasive the experiment, if the results turn out as predicted. On the other hand, I'm not persuaded that there is such a thing as a *minimal necessary change* to turn a case of knowledge into a Gettier case. But more to the point in the present context, even granting that there is such a thing as *minimal Gettierization*, it's an open question whether laypeople can be counted on to recognize it, and thus it's an open question what sort of guidance they need to competently assess

---

[23] Christina Starmans and Ori Friedman helpfully put this objection to me.

such a case. Nevertheless, that doesn't mean that we can't do better than what I've done thus far.

Suppose that laypeople routinely deny knowledge at significant rates in a wide range of non-minimal Gettier cases, yet they tend to ascribe knowledge in minimal Gettier cases.[24] Furthermore suppose that no plausible theoretical basis can be found to treat the two sorts of case differently, and that all the professionals think that the two sorts of case are clearly not examples of knowledge. To my mind, at that point the most plausible conclusion would be that laypeople can't be counted on to competently enough assess minimal Gettier cases. Their failure would reveal something interesting about how people ordinarily think of knowledge, and it would be very interesting to learn why they fail. But their failure would have no significant implication for a theory of knowledge.

But it might not come to that. This section reports an experiment designed to test whether the tripartite structure can effectively guide participants to assess Gettier cases that are more minimal than the ones used earlier.

9.1. Method

Participants (N=149)[25] were randomly assigned to one of five conditions: No Luck, Bad Luck, Good Luck, 1-Gettier (one-stage Gettier), and 3-Gettier (three-stage Gettier). The stories for the conditions

---

[24] I'm merely granting this for the sake of argument at this point. Below I argue that laypeople will deny knowledge in even minimal Gettier cases.

[25] Sixty-four female, aged 18–71, M = 31.8, SD = 11.9. Ninety-nine percent listed English as their native language. I eliminated data from twenty-six participants who failed comprehension questions.

were built up from narrative modules. Table 1 shows how they were built, along with the questions used. Here are the modules (participants did not see the labels):

(*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*) But the vibrations from the construction traffic caused something to happen [caused two things to happen]: ¶[26] [First,] 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*) [[But] The vibrations from the construction traffic caused something [else] to happen:] ¶ [Second,] 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

---

[26] Indicates paragraph break on the participant's screen.

the shower, she did not notice that this happened.

The comprehension questions were:

(CQ1) Right before Grace finishes her shower, she thinks that there is a blue Bic pen on the coffee table because _____. [she put one there/one fell from the ceiling]

(CQ2) Right before Grace finishes her shower, it is _____ for her to think that there is a blue Bic pen on the coffee table. [reasonable/unreasonable]

(CQ3) Right before Grace finishes her shower, is there a blue Bic pen on the coffee table? [Yes/No]

The test question was a dichotomous knowledge choice, followed by a confidence measure (1–10, low to high):

(KQ) Right before Grace finishes her shower, she _____ that there is a blue Bic pen on the coffee table. [really knows/only thinks she knows]

Table 1: List of modules composing the storylines, along with the questions, across the five conditions. Dashed lines mark page breaks.

No Luck	Bad Luck	Good Luck	1-Gettier	3-Gettier
No Luck Traffic CQ1–3 KQ Confidence	No Luck Traffic Bad Luck CQ1–3 KQ Confidence	No Luck Traffic Good Luck CQ1–3 KQ Confidence	No Luck Traffic Bad Luck Good Luck CQ1–3 KQ Confidence	No Luck CQ1–3 ------------- Traffic Bad Luck CQ3 ------------- Good Luck CQ3 KQ Confidence

Participants in 3-Gettier were asked two more questions than participants in other conditions were,[27] but as explained earlier (section 2), this is part of the tri-partitioning treatment. I sparingly made minor grammatical and narrative adjustments to ensure that the modules fit together smoothly in the different conditions (indicated by bracketed material above). For example, in 1-Gettier, instead of twice writing, 'But the vibrations from the construction traffic caused something to happen,' I once wrote, 'But the vibrations from the construction traffic caused two things to happen,' and then used ordinal introductory elements ('First' and 'Second') to begin the next two paragraphs. And in 3-Gettier, the storyline was more intelligible by including the Traffic module after the first page break.

---

[27] They answered CQ3 at the end of each stage, for a total of three times, compared to just once for participants in other conditions.

9.2.Results and discussion

Only one source of luck appears in the storyline — construction traffic — and is responsible for both the bad luck and the good luck. No other agent enters the scene: neither thief, nor husband, nor even so much as a mouse or housefly. In the Good Luck module, the previously hidden pen falls right on to the coffee table, and so is just as noticeable as the original pen. In the Bad Luck module, the pen that Grace set down merely falls on to the floor and is not described as hidden or otherwise concealed from view. The modular design allows us to compare how participants are affected by adding bad luck only, good luck only, as well as bad and good luck together. Moreover, it allows us to directly compare responses to unpartitioned versus partitioned versions of the exact same storyline.

Grace is arguably a more minimal Gettier subject than Katie is. Katie is victimized, but Grace isn't. The source of luck in Grace's story is impersonal and benign, but the source of luck in Katie's story is personal and malign (a pair of agents invade her home and rob her). Moreover, the source of luck in Grace's story is clearly unitary, whereas in Katie's story it's at least notionally binary.[28] On each comparison, Katie's story introduces complexities inessential to Gettierization: victimization, agency, malignancy, and a notionally binary source of luck. Grace's story is a more minimal Gettier case. If the tripartite method succeeds here, then it will have passed its most grueling test.

---

[28] I say 'at least notionally' because, as mentioned in section 6, the thieves are nondescript and act as a team.

I made eight predictions about the results. Knowledge ascription in No Luck would be very high, because it involves a normal case of indirect knowledge of a mundane matter of fact. Second, knowledge ascription in Bad Luck would be at floor, because the unexpected bad luck results in a false belief, whereas knowledge requires truth. Third, knowledge ascription in Good Luck would be high, and probably no different from No Luck. A bit of good luck that installs a "back-up" truth-maker, and nothing more, shouldn't hinder knowledge ascription. Fourth, knowledge ascription in 1-Gettier would not differ significantly from chance, in line with my earlier replication of Starmans and Friedman's case of Katie (section 4). Fifth, knowledge ascription in 1-Gettier would differ significantly from No Luck. Sixth, knowledge ascription in 3-Gettier would occur at rates significantly less than chance. Seventh, knowledge ascription in 3-Gettier would not differ significantly from Bad Luck. That is, tri-partitioned Gettierization will inhibit knowledge ascription at rates similar to falsehood. Finally, in light of all that, I predicted the following relative ordering of knowledge-ascription rates across the conditions (low to high): Bad Luck < 3-Gettier < 1-Gettier < Good Luck < No Luck.

All eight predictions were true. First, knowledge ascription in No Luck was very high (81%) and beyond what could be expected by chance (binomial, $p = 0.001$). Second, knowledge ascription in Bad Luck was very low (16%) and below what could be expected by chance (binomial, $p < 0.001$). Third, knowledge ascription in Good Luck was high (76%), greater than could be expected by chance (bi-

nomial, p = 0.016), and it didn't differ significantly from No Luck (Fisher's, p = 0.758). Fourth, knowledge ascription in 1-Gettier didn't differ significantly from chance (48%, binomial, p = 1) and was actually lower than I expected. Fifth, knowledge ascription in 1-Gettier differed significantly from No Luck (Fisher's, p = 0.018).[29] Sixth, knowledge ascription in 3-Gettier occurred at rates significantly less than chance (29%, binomial, p = 0.018). Seventh, knowledge ascription in 3-Gettier didn't differ significantly from Bad Luck (Fisher's, p = 0.347).[30] Finally, the relative ordering of knowledge ascription rates was as predicted (Figure 3).

Results from the weighted knowledge ascription tell a similar story. Mean weighted knowledge ascription in No Luck was high and beyond what could be expected by chance (M = 5.55, SD = 7.05, $t(30)$ = 4.38, p < 0.001). In Bad Luck it was low and below what could be expected by chance (M = -6.03, SD = 6.84, $t(31)$ = -4.99, p < 0.001). In Good Luck it was high, beyond what could be expected by chance (M = 5, SD = 7.5, $t(28)$ = 3.6, p < 0.001), and it didn't differ significantly from No Luck ($t(28)$ = -.39, p = 0.697). In 1-Gettier it didn't differ significantly from chance (M = .31, SD = 8.6, $t(28)$ = 0.195, p = 0.847) but differed significantly from No Luck ($t(28)$ = -3.3, p = 0.003). In 3-Gettier it was low, below what could be expec-

---

[29] It did not differ from 3-Gettier (Fisher's, p = 0.208). Given that knowledge ascription in Husband and Burglar didn't differ significantly, I didn't predict that it would differ significantly between 1-Gettier and 3-Gettier either.

[30] Rates of knowledge ascription in 3-Gettier, an authentic-evidence Gettier case, differed significantly from Zebra, an apparent-evidence Gettier case from section 5 (100%, binomial, p < 0.001). This supports Starmans and Friedman's point that laypeople are sensitive to the distinction and that it significantly affects rates of knowledge ascription.

ted by chance (M = -3.72, SD = 7.6, $t(27)$ = -2.583, p < 0.008), and it didn't differ significantly from Bad Luck ($t(27)$ = 1.6, p = 0.119). Finally, the relative ordering of mean weighted knowledge ascriptions across the conditions was the same as for dichotomous knowledge ascriptions (Figure 3).

As an added bonus, even though dichotomous knowledge ascription rates in 3-Gettier didn't differ from 1-Gettier (Fisher's, p = 0.104, one-tailed), weighted knowledge ascriptions did ($t(27)$ = -2.8, p = 0.018, one-tailed).

Figure 3 summarizes the results from this experiment, which further support my proposal that tripartitioning effectively guides laypeople to assess Gettier cases.
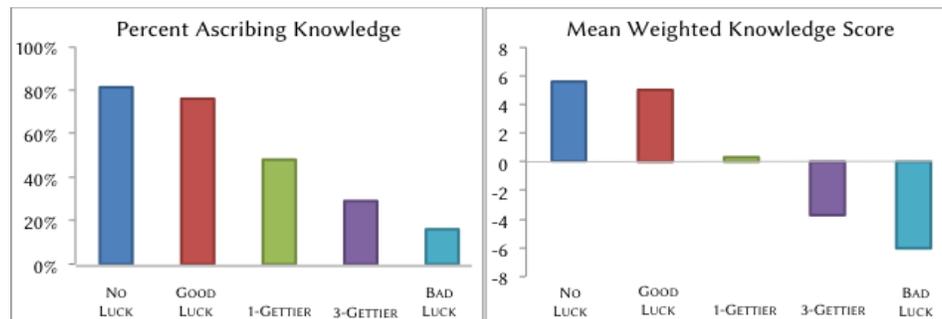
Figure 3: Left Panel: percentage of participants attributing knowledge in the five conditions. Right Panel: mean weighted knowledge ascription (derived by multiplying the dichotomous answer by reported confidence).

A couple other features of these results deserve mention. First, some might find the rate of knowledge attribution in Bad Luck surprisingly high, given that it was a clear case of false belief. One in six participants said that a false belief was real knowledge. But this

doesn't differ significantly from what other researchers have found,[31] so it seems that a small but persistent minority of people use 'really knows' non-factively. Second, some might find the rate of knowledge attribution in No Luck and Good Luck surprisingly low, given that the pen had only been out of Grace's sight for a very short time, and she is alone in a secure, locked apartment. Overall in these two conditions ($N = 60$), 78.3% ascribed knowledge. Although this far exceeds what could be expected by chance (binomial, $p <$ 0.001) and doesn't differ significantly from what was observed in Control from section 4 (84%, binomial, $p = 0.153$), it's still note-worthy that roughly one in six participants think that we don't really know such routine facts based on memory and inference. Put-ting the two observations together, roughly equal numbers of parti-cipants ascribe knowledge in ways that diverge significantly from what mainstream epistemological theories say is the true verdict: one-sixth ascribe knowledge too liberally, endorsing non-factive ap-plications,[32] and one-sixth ascribe knowledge too conservatively, endorsing skeptical abstentions.

10. Conclusion

I conclude that the conspicuous tripartite structure effectively guides laypeople to competently assess Gettier cases. When probed this way, laypeople across very different cultures, male and female,

---

[31] Starmans and Friedman (2012) report ~10% ascribing knowledge in their false belief controls, and Buckwalter (ms.) reports a similar consistent pat-tern of participants using 'knows' non-factively.

[32] Here I'm setting aside knowledge ascriptions to Gettier subjects.

young and old, reveal that they overwhelmingly share the Gettier intuition.

Further research could investigate whether any factors mediate tripartitioning's effect on participant response. For example, it's possible that tripartitioning works primarily because it causes participants to be more reflective. Relatedly, further research could investigate whether dual-process psychology (Stanovich and West 2000) explains the difference between judgments in tripartitioned and non-tripartitioned conditions. Perhaps the automatic, unreflective application of our concept of knowledge ("System 1 processing") classifies Gettier cases as knowledge, and it is only when distinct, reflective psychological capacities are triggered ("System 2 processing") that participants tend to deny knowledge in such cases. If this dual-process conjecture turns out to be true, then even though the Gettier intuition is *readily* shared, it isn't *immediately* shared. And this would be a very important fact about how we ordinarily think about knowledge.[33]

References

Alexander, Joshua and Jonathan Weinberg. 2007. Analytic epistemology and experimental philosophy. *Philosophy compass* 2.1: 56–80.

Beebe, James and Wesley Buckwalter. 2010. The epistemic side-effect effect. *Mind and language* 25.4: 474–498.

Beebe, James and Joseph Shea. Ms. Gettierized Knobe effects.

Buckwalter, Wesley. 2012. Non-traditional factors in judgments about knowledge. *Philosophy compass* 7: 278–289.

Buckwalter, Wesley. Forthcoming. Gettier made ESEE. *Philosophical psychology*.

Buckwalter, Wesley. Ms. Factive verbs and protagonist projection.

Cullen, Simon. 2010. Survey-driven romanticism. *Review of philosophy and psychology* 1: 275–296.

Gettier, Edmund. 1963. Is justified true belief knowledge? *Analysis* 23.6: 121–3.

Hetherington, Stephen. Forthcoming. There can be lucky knowledge. In *Contemporary debates in epistemology*. Ed. Matthias Steup, John Turri and Ernest Sosa. Wiley-Blackwell.

Jackson, Frank. 2011. On Gettier holdouts. *Mind & Language* 26.4: 468–481.

Ludwig, Kirk. 2007. The epistemology of thought experiments: first vs. third person approaches. *Midwest studies in philosophy* 31.1: 128–159.

Matilal, Bimal. 1986. *Perception: an essay on classical Indian theories of knowledge*. Oxford UP.

Nagel, Jennifer. Forthcoming. Intuitions and experiments. *Philosophy and phenomenological research*. http://philpapers.org/archive/NAGIAE.1.pdf. Accessed 23 July 2012.

Award.

Nagel, Jennifer, Valerie San Juan and Raymond A. Mar. Forthcoming. Lay denial of knowledge for justified true beliefs. *Cognition*.

Sartwell, Crispin. 1991. Knowledge is merely true belief. *American philosophical quarterly* 28.2: 157–165.

Sosa, Ernest. 2007. Experimental philosophy and philosophical intuitions. *Philosophical studies* 132: 99–107.

Sosa, Ernest. 2009. A defense of the use of intuitions in philosophy. *Stich and his critics*. Ed. Dominic Murphy and Michael Bishop. Blackwell.

Stanovich, Keith and Richard West. 2000. Individual differences in reasoning: implications for the rationality debate. *Behavior and brain sciences* 23: 645–665.

Starmans, Christina and Ori Friedman. 2012. The folk conception of knowledge. *Cognition*. http://dx.doi.org.proxy.lib.uwaterloo.ca/10.1016/j.cognition.2012.05.017.

Stoltz, Jonathan. 2007. Gettier and factivity in Indo-Tibetan epistemology. *Philosophical quarterly* 57.228: 394–415.

Turri, John. 2011. Manifest Failure: the Gettier problem solved. *Philosophers' imprint* 11.8: 1–11

Turri, John. 2012a. Is knowledge justified true belief? *Synthese* 184.3: 247–259.

Turri, John. 2012b. In Gettier's wake. *Epistemology: the key thinkers*. Ed. Stephen Hetherington. Continuum.

Weinberg, Jonathan, Shaun Nichols and Stephen Stich. 2001. Normativity and epistemic intuitions. *Philosophical topics* 29: 429–60

Williamson, Timothy. 2007. *The philosophy of philosophy*. Blackwell.

Williamson, Timothy. 2011. Philosophical expertise and the burden of proof. *Metaphilosophy* 42: 215–229.

Zagzebski, Linda. 1996. *Virtues of the mind*. Cambridge UP.