

February 21, 2013

Forthcoming in *Cognition*

Lay Denial of Knowledge for Justified True Beliefs

Jennifer Nagel,^a Valerie San Juan,^b and Raymond A. Mar^c

^aUniversity of Toronto, Philosophy

170 St. George St., Toronto, ON, Canada M5R 2M8

jennifer.nagel@utoronto.ca

^bOntario Institute of Studies in Education, Human Development and Applied Psychology

252 Bloor St. W., Toronto, ON, Canada M5S 1V6

valerie.sanjuan@utoronto.ca

^cYork University, Psychology

4700 Keele St. W. Toronto, ON, Canada M3J1P3

mar@yorku.ca

Corresponding Author: Jennifer Nagel

P: +1 416-978-3311 | F: 416-978-8703.

Abstract

Intuitively, there is a difference between knowledge and mere belief. Contemporary philosophical work on the nature of this difference has focused on scenarios known as “Gettier cases.” Designed as counterexamples to the classical theory that knowledge is justified true belief, these cases feature agents who arrive at true beliefs in ways which seem reasonable or justified, while nevertheless seeming to lack knowledge. Prior empirical investigation of these cases has raised questions about whether lay people generally share philosophers’ intuitions about these cases, or whether lay intuitions vary depending on individual factors (e.g. ethnicity) or factors related to specific types of Gettier cases (e.g. cases that include apparent evidence). We report an experiment on lay attributions of knowledge and justification for a wide range of Gettier Cases and for a related class of controversial cases known as Skeptical Pressure cases, which are also thought by philosophers to elicit intuitive denials of knowledge. Although participants rated true beliefs in Gettier and Skeptical Pressure cases as being justified, they were significantly less likely to attribute knowledge for these cases than for matched true belief cases. This pattern of response was consistent across different variations of Gettier cases and did not vary by ethnicity or gender, although attributions of justification were found to be positively related to measures of empathy. These findings therefore suggest that across demographic groups, laypeople share similar epistemic concepts with philosophers, recognizing a difference between knowledge and justified true belief.

Keywords: Justified True Belief; Knowledge Ascription; Gettier Cases; Mental State Inference; Individual Differences; Lay Beliefs

1. Introduction

1.1. Justified true belief and Gettier Cases

When do we naturally see others as knowing, rather than simply believing, that something is the case? In tackling this question, both philosophers and psychologists of mental state ascription have been struck by certain core features of knowledge. According to a simple theory suggested by Plato (1990, 339) and explicitly endorsed by many philosophers and psychologists since, knowledge is equivalent to justified true belief (e.g. Ayer, 1956; Bartsch & Wellman, 1995; Chisholm, 1957; Miller, Hardin, & Montgomery, 2003). The justified true belief or JTB theory maintains that if a person's belief is not only true, but held on some basis such as clear perceptual evidence or valid inference, then it constitutes knowledge.

The JTB theory has fallen sharply into disfavor among philosophers in recent decades (Dretske, 1983; Pollock & Cruz, 1999; Shope, 1983; Sosa, 1991; Williamson, 2000), primarily as the result of a class of intuitive counterexamples noticed by Edmund Gettier half a century ago (Gettier, 1963). Gettier's counterexamples involved beliefs that were justified and true, but did not register intuitively as knowledge. Philosophers have since devised a great variety of "Gettier cases." For example, consider the following (adapted from Lehrer, 1965):

Mike works in an office with two co-workers, Havit and Nogot. Nogot, who is generally a trustworthy person, has told Mike he owns a Ford, shown him his Ford key ring, and pointed out a Ford as his own in the company parking lot. So Mike believes that Nogot owns a Ford ... but in fact Nogot is faking it, all very skillfully, and doesn't own a car. Meanwhile Havit always rides his bike to work and pretends not to own a car, but secretly does have a Ford at home in his garage.

Does Mike know that someone in his office owns a Ford? This proposition is true, in virtue of the hidden facts about Havit's car ownership. Furthermore, Mike's belief in this proposition is

justified, given Nogot's skillful campaign. However, philosophers have found it intuitive that Mike does not *know* that someone in his office owns a Ford (Lehrer, 1965; Shope, 1983). Still, it is a live question whether philosophers are exceptional in finding such cases to be unsatisfying examples of knowledge, or whether adults generally distinguish between knowledge and justified true belief. Several studies have found evidence that laypeople may judge the agents of Gettier cases to lack knowledge (Cullen, 2010; Starman & Friedman, 2012; Weinberg, Nichols, & Stich, 2001; Wright, 2010). However, no study to date has investigated explicit attributions of justification for these cases, so the relationship between perceptions of knowledge and justified true belief remains unclear. The current study was therefore designed as a systematic exploration of attributions of both knowledge and justification for the kind of cases that philosophers have taken to separate knowledge from justified true belief.

1.2 Gettier Cases and the Causal Condition

Justified true belief can fall short of knowledge in many ways. Indeed, the variety of Gettier cases is broad enough that it is now widely agreed in philosophy that there is no simple supplement that can be added to the JTB conditions to yield a satisfactory analysis of knowledge, although numerous attempts have been made (Shope, 1983; Williamson, 2000). For example, one might think that the problem in the example given above is that Mike infers his justified and true conclusion from a belief that happens to be false. Such an observation prompted early theorists to add a "no false lemma" rule to the JTB analysis, stipulating that knowledge cannot arise from reasoning that essentially relies on a false belief (Clark, 1963). However, this suggestion is no longer seen as viable, both because there are cases in which the agent complies with that further condition but still seems not to know (Feldman, 1974; Saunders & Champawat, 1964), and because there are cases in which the agent violates the condition but does seem to know (Fitelson, 2010; Warfield, 2005).

Another proposed modification to the traditional JTB theory was the addition of a requirement that the agent's belief should be appropriately caused by the fact believed (A. Goldman, 1967). Some psychologists have endorsed this idea of adding a causal condition to the JTB analysis (e.g. Dienes & Perner, 1999, 739; Perner, 1991, 304). However the idea of analyzing knowledge by adding a causal condition to the JTB conditions is now widely rejected in philosophy, even by the originator of the causal theory, Alvin Goldman. Goldman came to reject his earlier theory in part because of a new type of Gettier case he discovered, in which a belief might be caused in the very same way paradigmatic instances of knowledge are caused—say, through ordinary perceptual observation—and yet still seem to fall short of knowledge (A. Goldman, 1976). If a hunter is lucky to be looking at the one real duck on a lake covered with realistic decoys, his belief that it is a duck may be true and justified, but does not seem to count as knowledge (assuming he cannot distinguish the duck from the decoys). Many philosophers have taken this type of example to show that justified true belief can fall short of knowledge even when there is nothing unusual about the causal connection between the agent's belief and the fact believed (Pollock & Cruz, 1999; Shope, 1983; Steup, 2008).

1.3 Gettier cases and Skeptical Pressure cases

Gettier cases involve a justified true belief that intuitively registers as falling short of knowledge; they share this characteristic with a distinct set of cases known in epistemology as “Skeptical Pressure” cases. In typical Gettier cases, the agent is really at risk of being mistaken: Mike is being lied to and the hunter is surrounded by things that look like ducks but are not. However, philosophers have noticed that we may also get an intuitive sense that someone lacks knowledge when risks of error are just mentioned as possible, even if these risks are not actually present (Hawthorne, 2004; Vogel, 1990). For example, one classic Gettier case involves a person who forms a belief about the time by glancing at a broken clock that happens to be stopped at the right time:

her belief is true, and justified, but it seems that the agent who looks at the broken clock does not really gain knowledge of the correct time (Russell, 1948; Scheffler, 1965). In the corresponding Skeptical Pressure version of the case, it is stipulated that the clock is actually working, but the hypothetical possibility of risk is stressed: individuals are told that clocks are sometimes broken and that the agent has not double-checked whether the clock is working. When possibilities of error are highlighted in this manner it can seem intuitive that the agent who glances at the clock does not really *know* what time it is, and is just lucky that the clock is working.

1.4 Evaluations of justification

In assessing the mental state of an agent, it is possible to evaluate not only the outcome—whether the agent gained a true or false belief, for example—but also the quality of the process by means of which the agent arrived at this outcome. Evaluations of justification are traditionally thought to apply strictly to the quality of the agent’s thought process, independent of its outcome (BonJour & Sosa, 2003; A. I. Goldman, 1979; Plantinga, 1993). True beliefs can be unjustified, as in lucky guessing; false beliefs can be justified, as in cases where a responsible agent has relied on misleading evidence. In theory, the agent who glances at a broken clock and arrives at a false belief about the time should be seen as having exactly as much justification as the parallel agent who glances at a working clock and arrives at a true belief. In practice, however, evaluations are often biased by our knowledge of outcomes (e.g. S. A. J. Birch, 2005; Fischhoff, 1975). Under the influence of outcome bias, it would seem that an agent glancing at a broken clock has reasoned less well when her resulting mental state is a false belief, even if this agent’s cognitive conduct is just the same as that of a parallel agent who glances at a working clock.

Lower levels of apparent justification would also be expected for the Gettier case agent who glances at a clock that simply happens to be stopped at the right time. Even if this agent could not have anticipated that she had chosen a risky strategy for telling the time, our privileged information

about the agent's setting will make the risks of her course of action salient to us. Even mentioning some possibility of error on the part of the agent could lower attributions of justification. By making the reader focus on the possibility of the risk of error, it may be possible to generate an impression that the agent in the case should be taking care to control for such risk, for example by double-checking whether the clock is working. One of the central biases of adult mental state ascription is epistemic egocentrism, or the tendency to evaluate others as though they shared our privileged information and concerns (Birch, 2004; SA Birch & Bloom, 2007; Camerer, Loewenstein, & Weber, 1989; Nickerson, 1999). Someone who is actively wondering whether a clock might be broken cannot just glance at that very clock to tell the time, but would rationally gather further evidence before making up her mind. If the reader of a scenario is encouraged to think about the possibility that the clock might be broken, the bias of epistemic egocentrism should incline this reader to evaluate the agent described in the scenario as though she shared the reader's own active worry (Nagel, 2010).

1.5 Gettier cases and mental state ascription

Gettier cases present especially intriguing challenges in perspective-taking. In evaluating a Gettier case, we recognize that the proposition believed by the agent in the case is true and that the agent has reasoned in a manner that would seem appropriate from his perspective. At the same time we are given some privileged information about the situation that makes another way of reasoning seem more suitable. If Mike had examined Nogot's situation more thoroughly, he would have revised his opinion about whether someone in his office owned a Ford; if he had also investigated Havit's garage, he would have returned to believing the original proposition, but on a quite different basis. Meanwhile, if the hunter had known about the decoys, he might have taken a closer look before making his judgment about the duck. Because cases like this involve various dimensions of epistemic appraisal, their evaluation should interest theorists of mental state attribution.

1.6 Objectives

1.6.1. Prevalence and universality of epistemic intuitions regarding Gettier cases

Gettier case intuitions have been examined in a few studies concerned with the nature of philosophical intuitions more broadly. One of these studies found that the Gettier case it tested (reproduced in Appendix B) triggered different responses from members of different ethnic groups, calling into question the universality of responses to these cases (Weinberg, et al., 2001). Another study found that responses to its Gettier case were somewhat susceptible to order effects: the majority of participants gave the expected response when the case followed a clear case of knowledge or ignorance, but not after a more ambiguous case (Wright, 2010). This variability in responding similarly raised concerns about the generalizability of philosophers' responses to Gettier cases. A further study confirmed that the majority of participants agree that the agent in a Gettier case lacks knowledge, but drew attention to the sensitivity of such responses to the precise form of the question being asked: participants were more likely to deny knowledge when their survey options were "really knows" / "only believes" than when the options were simply "knows" / "does not know" (Cullen, 2010).

A more recent study, by Starmans and Friedman (2012), examined responses to a broader range of cases: seven different stories were tested, only two of which elicited the standard response that the agent lacked knowledge. The two Gettier cases that worked were both what philosophers call "false lemma" cases; they involved a true conclusion being inferred from a false but justified belief, like our Havit/Nogot case above. Starmans and Friedman introduce their own terminology, labeling these as "apparent evidence" cases, in contrast to what they call "authentic evidence" cases. Their authentic evidence category covers those cases that do not involve any reasoning on the basis of a false belief, like our case of the hunter and the duck. Mainstream philosophers see both authentic and apparent evidence Gettier cases as involving failures of knowledge. Starmans and

Friedman suggest that laypeople react differently to apparent and authentic evidence cases, denying knowledge only for apparent evidence cases; they propose that laypeople follow the rule that knowledge is justified true belief originally arising from authentic evidence. However, this suggestion is not consistent with the three earlier studies, all of which found the majority of participants denying knowledge for Gettier cases employing authentic evidence.

Starman and Friedman (2012) also raise the question of whether participants might be willing to attribute knowledge in any case of true belief, and presented their participants with one pair of cases tailored to contrast true beliefs that they took to exhibit high and low levels of justification. In line with their expectations, knowledge was attributed in the high justification but not the low justification condition. They did not, however, question their participants directly on the level of perceived justification for these cases.

No study to date has examined the relationship between explicit attributions of justification and knowledge. This study was designed to bridge this gap, and also to explore responses to a variety of different Gettier cases, including cases with and without “apparent evidence” or reliance on a false lemma.

Skeptical Pressure cases provide another situation in which justified true beliefs might intuitively appear not to constitute knowledge, and these have yet to be subjected to empirical research. We aim to fill this gap by checking whether participants without philosophical training share a similar pattern of intuitive responses to these cases. Importantly, by including both Gettier cases and Skeptical Pressure cases, we can for the first time directly compare attributions of knowledge made for both types of situations. If knowledge is ordinarily perceived as nothing other than justified true belief, Gettier cases and Skeptical Pressure cases should ordinarily be perceived as cases of knowledge. However, if the standard philosophical intuitions about knowledge are also shared by non-philosophers, then laypeople should see the agents of these cases as lacking

knowledge. In our study, we also compare responses to Gettier Cases and Skeptical Pressure cases to those for instances of ordinary true and false belief, thus allowing us to further evaluate how these unique cases of justified true belief are separate from more straightforward cases.

One final issue is also addressed for the first time in the present study. Understood literally, the verb “to know” presupposes the truth of its complement clause, so that the sentence “John knows that the door is locked, but the door isn’t locked” is semantically unacceptable (Kiparsky & Kiparsky, 1970). Known as “factivity,” this property of the verb “know” is not shared by verbs such as “think,” which are classified as nonfactive: it is fully acceptable to say “John thinks that the door is locked, but the door isn’t locked.” However, we can also use the verb “to know” in a projected or non-literal sense, as in the sentence “John *knew* that his team was going to win, and then they didn’t” (Holton, 1997). Someone who asserts this sentence is best understood as expressing the thought that it seemed very much to John that he knew. This is similar to the emphasis accompanying irony, where the emphasis on the projected term is a sign of non-literal usage. When knowledge attributions are being assessed through a survey method, it is particularly important to gauge whether participants mean what they say literally. For this reason, when participants are asked whether an agent knows, it is useful to follow up affirmative responses with a further question probing whether a literal knowledge attribution is intended.

1.6.2 Assessing individual differences in epistemic intuition

Earlier studies of Gettier cases seemed to show variability in responding to a single Gettier case based on individual differences (i.e., ethnicity) and contextual influences (i.e., problem order). In order to explore this issue more completely, we examine whether variation in these intuitions correlates with a number of different individual difference factors, including ethnicity. We further explored this issue by examining other basic demographic variables, such as gender and age.

There are also some individual differences more directly linked to the problem of knowledge ascription during Gettier cases. Because the ascription of knowledge fundamentally deals with understanding other people and their minds, it might be the case that those naturally more inclined to empathize with others could ascribe knowledge more readily during these unique cases. Trait empathy, typically seen as an emotional form of perspective-taking or “feeling for” others (Preston & de Waal, 2002), was therefore also examined in this study. Another individual difference that might influence responses to Gettier cases is philosophical training; consequently, we examined whether those with philosophical education respond similarly to those without.

2. Methods and Materials

2.1 Participants

Canadian undergraduate students participated in this study for academic credit ($N = 238$). Individuals who did not complete the survey were removed ($N = 16$), along with those who completed the survey so quickly they were unlikely to have read and considered all items seriously (i.e., less than 13 minutes for approximately 70 questions; $N = 14$). The remaining sample included 208 individuals (37 males), averaging in age 21.8 years ($SD = 4.6$, 1 not reporting). Participants self-identified as being from a variety of ethnic backgrounds: 70 as White (33.7%), 53 as South Asian (25.5%), 24 as East Asian (Chinese, Korean, Vietnamese, etc.; 11.5%), 15 as Latin American (7.2%), 14 as Black (6.7%), and 32 as Other (including Arab, West Asian, First Nation, and unspecified; 15.4%). Around a quarter ($N = 48$; 23.1%) reported having taken at least one philosophy course at university, with most having taken only one course ($N = 33$, $M = 1.87$, $SD = 1.68$).

2.2 Materials and procedure

All questionnaires were administered using an online survey client (qualtrics.com). Participants first completed the belief and knowledge survey, followed by the empathy and demographic questionnaires. At the end of the survey participants were invited to give feedback on

their experiences during the study, including the factors they considered when evaluating how justified characters were in their beliefs.

2.2.1 Belief and Knowledge Survey

Participants read and responded to questions about 16 vignettes, randomly ordered, half of which were experimental with the remainder acting as filler. Filler vignettes were the same for all participants, and included a variety of scenarios about justified and unjustified belief formation. The eight experimental vignettes were made up of four story types: (1) Gettier cases, (2) Skeptical Pressure cases, (3) Standard True Belief, and (4) Justified False Belief. Participants responded to two versions of each story type. For Gettier cases, the two versions corresponded to two slightly different types of cases, those based on either authentic evidence (where the agent does not form any false belief, as in our example 1A in Appendix A) or apparent evidence (where the agent does initially form a false belief, as in our example 2A).

Participants were randomly assigned to one of four survey versions, to counterbalance the type of story associated with each individual vignette (i.e., characters and situation). This ensured that within each survey, each vignette was tested in association with a different story type, while keeping the word count constant across vignettes (see Appendix A). In all survey versions, a 9th Gettier vignette appeared at the end (not counterbalanced). This vignette was the one employed by Weinberg and colleagues (2001), in an attempt to replicate their effects of ethnicity on knowledge ascription.

Because Gettier cases are seen as illustrating the possibility of justified true belief without knowledge, it was important to assess whether participants viewed the agent as possessing a belief that was both justified and true. Participants were therefore asked to evaluate the truth of the key proposition in the story, with the story present on-screen alongside the question (e.g., “According to the story, what kind of stone is in the necklace that Emma tries on?”). If a participant failed to

answer this most basic question incorrectly, they were redirected to the next case without answering any further questions.

Participants who passed the initial comprehension screen answered a series of questions pertaining to the agent's mental state. First they evaluated whether the agent believed the key proposition (belief ascription; e.g., "Does Emma think that the stone in the necklace is a diamond?"). Participants who answered "No" were not asked further questions about the agent's belief and were instead redirected to the final question of the trial (i.e., a memory question). Participants who answered "Yes" to the belief ascription question were then presented with two sets of questions regarding (a) how justified the agent was in forming his/her belief, and (b) whether to ascribe knowledge to the agent. The order for belief justification and knowledge ascription was counterbalanced. For belief justification, participants were asked to rate how justified the agent was in his/her belief (e.g., "How justified is Emma in thinking that the stone is a diamond?") using a seven-point Likert scale ranging from 1 (*completely unjustified*) to 7 (*completely justified*). For knowledge ascription, participants were asked whether or not the agent was knowledgeable of the key proposition (e.g., "Does Emma know whether or not the stone is a diamond?"). Three response options were provided: (1) Yes, she knows, (2) No, she doesn't know, and (3) Unclear – not enough information provided in the story. For participants who responded "No" (2) or "Unclear" (3), no further questions about the agent's knowledge were asked. For participants who responded "Yes" (1), a follow-up question was presented to further gauge the robustness of their knowledge attribution: "In your view, which of the following sentences better describes Emma's situation? (a) Emma knows that the stone is a diamond (b) Emma feels like she knows that the stone is a diamond, but she doesn't actually know that it is." Responses to both sets of knowledge ascription questions were then used to generate ordinal knowledge ascription scores. Participants who answered "No" to the first question were classified as attributing 'immediate' denial of knowledge

and received a score of 0. Participants who answered “Yes” to the first question but then chose option (b) for the second question were classified as attributing ‘delayed’ denial of knowledge and received a score of 1. Finally, participants who answered “Yes” to the first question and chose option (a) for the second question were classified as giving the strongest attribution of knowledge and received a score of 2.

The final question of each trial was a memory question that asked participants to recall a specific detail from the story (e.g., “In the story you just read, what did Emma say as she tried on the necklace?”). Participants had the option to fill in a blank or to select ‘do not recall.’ Participants who volunteered an answer were then asked to rate their confidence in that answer using a 5-point Likert scale (1 = not at all, 5 = absolutely certain). We hypothesized that participants who were more engaged with the story might be more likely to give the philosophically standard response.

2.2.2 Empathy

Participants completed a self-report measure of empathy in the form of the Toronto Empathy Questionnaire (TEQ) (Spreng, McKinnon, Mar, & Levine, 2009). The TEQ consists of 16 items and is an empirical single-factor derivation of self-reported empathy, characterizing empathy as a primarily emotional process. It has been demonstrated to predict task-based measures of mental inference, has good internal consistency, and good test-retest reliability (Spreng et al., 2009).

2.2.3 Demographics

All participants also completed a demographic questionnaire that included questions regarding age, gender, ethnicity, and prior training in philosophy.

3. Results

Separate analyses were conducted for participants’ responses to questions of (a) belief ascription, (b) knowledge ascription, and (c) belief justification.

3.1 Belief Ascription Scores

To examine if responses to belief ascription questions (e.g., “Does Emma think that the stone in the necklace is a diamond?”) varied between story types, a Generalized Estimating Equation (GEE) was conducted where the dependent variable was belief denial (coded as 0) or ascription (coded as 1). A significant effect of story type was observed ($\chi^2 = 32.36$, $df = 3$, $N = 208$, $p < .001$). Follow-up comparisons between story types found a greater likelihood of ascribing belief in cases of: (1) Standard True Belief versus Skeptical Pressure ($\chi^2 = 25.66$, $df = 1$, $N = 208$, $p < .001$, $B = 1.11$, $Exp(B) = 3.04$), (2) Standard True Belief versus Gettier Cases ($\chi^2 = 23.31$, $df = 1$, $N = 208$, $p < .001$, $B = 1.17$, $Exp(B) = 3.21$), and (3) Standard True Belief versus False Belief ($\chi^2 = 27.42$, $df = 1$, $N = 207$, $p < .001$, $B = 1.32$, $Exp(B) = 3.74$) ($\alpha = .008$, to control for multiple comparisons). A summary of response frequencies is presented in Table 1.

Table 1. *Frequency of Belief Ascription Responses by Story Version and Story Type.*

	Frequency of Responses	
	Belief Denial	Belief Ascription
Standard TB	24 (6.4 %)	351 (93.6 %)
Skeptical Pressure	58 (17.2 %)	279 (82.8 %)
Gettier	60 (18.0 %)	273 (82.0 %)
False Belief	66 (20.4%)	258 (79.6%)

3.2 Knowledge Ascription Scores

To examine the effects of story type on knowledge ascription (e.g., “Does Emma know whether or not the stone is a diamond?”) including our follow-up questions assessing strength of ascription (i.e., “In your view, which of the following sentences better describes Emma’s situation...”) a GEE analysis was again conducted. Knowledge ascription, the dependent variable,

was operationalized as an ordinal score, as described above (0 = immediate knowledge denial, 1 = delayed knowledge denial, 2 = unwavering knowledge ascription).

A significant effect of story type was observed ($\chi^2 = 247.97$, $df = 3$, $N = 207$, $p < .001$), prompting follow-up comparisons between story types ($\alpha = 0.008$). There was a greater likelihood of ascribing knowledge in cases of: (1) Standard True Belief versus Skeptical Pressure ($\chi^2 = 81.52$, $df = 1$, $N = 204$, $p < .001$, $B = 1.35$, $Exp(B) = 3.87$), (2) Standard True Belief versus Gettier Cases ($\chi^2 = 107.81$, $df = 1$, $N = 202$, $p < .001$, $B = 1.62$, $Exp(B) = 5.07$), (3) Standard True Belief versus False Belief ($\chi^2 = 231.60$, $df = 1$, $N = 204$, $p < .001$, $B = 2.71$, $Exp(B) = 14.99$), (4) Skeptical Pressure versus False Belief ($\chi^2 = 59.48$, $df = 1$, $N = 197$, $p < .001$, $B = 1.45$, $Exp(B) = 4.27$), and (5) Gettier versus False Belief ($\chi^2 = 49.08$, $df = 1$, $N = 196$, $p < .001$, $B = 1.13$, $Exp(B) = 3.09$). A summary of response frequencies for each story type is presented in Table 2.

Table 2. *Frequency of Knowledge Ascription Responses Across both Time Periods*

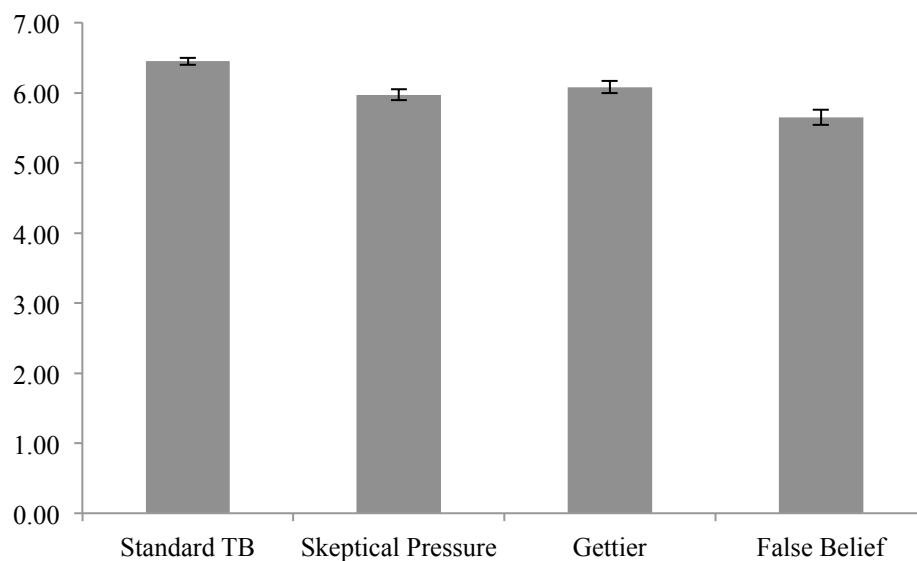
	Frequency of Responses		
	Immediate Knowledge Denial	Delayed Knowledge Denial	Unwavering Knowledge Ascription
Standard TB	49 (14.4 %)	35 (10.3 %)	256 (75.3%)
Skeptical Pressure	89 (33.0 %)	69 (25.6 %)	112 (41.5%)
Gettier	108 (41.1%)	62 (23.6%)	93 (35.4 %)
False Belief	166 (67.8 %)	44 (18.0 %)	35 (14.3 %)

3.2 Belief Justification Scores

To examine the effects of story type on participants' belief justification ratings a one-way repeated measures ANOVA was conducted. A significant effect of story type was observed (Wilks's $\Lambda = .66$, $F(3, 135) = 22.99$, $p < .001$, $\eta^2 = .34$). Follow up comparisons using t -tests ($\alpha = 0.008$)

revealed that the beliefs in Standard True Belief cases ($M = 6.45, SD = .60$) were evaluated as more justified than those in False Belief Cases ($M = 5.65, SD = 1.27; t(160) = 7.07, p < .001, d = .56$), Skeptical Pressure cases ($M = 5.97, SD = .92; t(179) = 6.17, p < .001, d = .46$), and Gettier cases ($M = 6.08, SD = 1.01; t(180) = 5.36, p < .001, d = .40$) (Figure 2). Justification ratings for Gettier and Skeptical Pressure cases fell between those for ordinary True Belief cases and False Belief cases, but did not significantly differ from each other ($p > .05$) (Figure 2). These similar justification ratings are consistent with the hypothesis that a similar mechanism, such as the bias of epistemic egocentrism, is at work in both these types of cases.

Figure 2: Average Justification Rating by Case



3.3 Authentic versus Apparent Evidence in Gettier Cases

To examine whether participant responses varied between Gettier cases that included authentic or apparent evidence, a separate set of analyses was conducted. Specifically, the effect of evidence type was examined for the main dependent measures (i.e., belief ascription, knowledge ascription, and belief justification).

For belief ascription, a GEE analysis indicated that participants were 4.44 times more likely to ascribe belief in cases with authentic evidence compared to cases with apparent evidence ($\chi^2 = 20.28$, $df = 1$, $N = 201$, $p < .001$, $B = 1.49$). Similarly, participants were 2.28 times more likely to ascribe knowledge immediately in cases with authentic evidence compared to those with apparent evidence ($\chi^2 = 14.62$, $df = 1$, $N = 179$, $p < .001$, $B = 1.10$). However, there were no significant differences in participants' belief justification ratings between the two Gettier case types based on a paired samples *t*-test.

3.4 Examining individual differences

In light of the fact that Gettier and Skeptical Pressure cases resulted in reduced knowledge ascription compared to standard True Belief cases, we explored how individual differences might relate to variability in responding. We first examined whether those with philosophical training responded differently compared to those without. Students who had taken one or more philosophy courses did not respond differently than those who had not, with respect to knowledge ascription and belief justification, across all case types (all $ps > .10$). The difference that most closely approached statistical significance was knowledge ascription for Gettier cases, with those who had taken a philosophy course reporting less knowledge in these cases compared to those who had not, $t(177) = 1.49$, $p = .14$, $d = .26$.

Next, we examined the possible influence of age, gender, and ethnicity on responses to these types of cases. Age was not related to responses based on Spearman's rho (to account for non-normality of the data; all $ps > .25$), and men and women did not respond differently for most cases, although there was a trend for women to attribute more knowledge in Skeptical Pressure cases, $t(178) = -1.87$, $p = .06$, $d = .35$. To examine whether ethnicity had an influence on participants' knowledge ascription and belief justifications, separate repeated measures ANOVAs were conducted. For both knowledge ascription and belief justification, no significant main effect was

found for ethnicity and there was no interaction between ethnicity and story type. This suggests that for the stories used in the current study, there were few differences in response pattern as a function of age, gender, and ethnic background.

In order to attempt a replication of the previous study examining ethnicity and differences in knowledge ascription, we looked at responses to the Gettier case used by the past researchers (Weinberg et al., 2001; see Appendix B). We presented their case verbatim, although they asked participants to select between “Really Knows” and “Only Believes,” in contrast to our response options. Note that in this Gettier case the agent originally has authentic evidence and does not reason on the basis of a false belief. In our study, as in the original study, a majority of participants (59%) denied that the agent had knowledge, in line with the standard philosophical treatment of the case. In the original study, also conducted with undergraduate participants at a North American University, a minority of East Asians (43%) and South Asians (39%) indicated that the agent had mere belief rather than knowledge, contrasted with 74% of ‘Western’ participants. In our study, however, a similar discrepancy did not emerge (see Table 3). Two separate one-way ANOVAs confirmed that for this particular Gettier case, there were no significant ethnic group differences for knowledge ascription or belief justification (all p s > .05). Findings of cross-cultural similarity should be read with caution, however, given that our minority participants probably included at least some second- or third-generation immigrants who would be expected to show greater cultural similarity to the majority.

Lastly, we investigated whether variability in empathy was related to knowledge attribution and ratings of justification. Correlations did not reveal any association between empathy and knowledge ascription, for any of the case types (Spearman’s Rho, all ρ s < .11, all p s > .15). With respect to belief justification, however, greater empathy was related to higher ratings of justification for Gettier cases ($\rho = .16, p = .028$) and Standard True Belief cases ($\rho = .14, p = .044$).

Table 3: *Knowledge Ascription for “Buick Story” by Ethnicity*

Ethnic Group	Immediate Knowledge Denial	Delayed Knowledge Denial	Unwavering Knowledge Ascription	Total
White	27 (42.9%)	16 (25.4%)	20 (31.7%)	63
East Asian	10 (47.6%)	5 (23.8%)	6 (28.6%)	21
South Asian	17 (40.5%)	7 (16.7%)	18 (42.9%)	42
Latin	3 (25.0%)	2 (16.7%)	7 (58.3%)	12
Black	3 (27.3%)	5 (45.5%)	3 (27.3%)	11
Other	7 (25.0%)	3 (10.7%)	18 (64.3%)	28
Total	67	38	72	177

Note. Percentage of total respondents for each ethnic group is shown in brackets.

3. Discussion

According to the classical analysis, knowledge is nothing more than justified true belief. However, for almost fifty years now, the dominant view in philosophy has been that there are intuitive differences between knowledge and justified true belief. The results of our study indicate that for people in general, the cases philosophers have taken to reveal those intuitive differences—Gettier cases—are indeed widely judged to be instances of justified true belief without knowledge. This was true for Gettier cases that involve “apparent evidence,” or the initial formation of a false belief, but also for cases where no false belief is formed. Moreover, this pattern appeared to be quite robust. Gettier case intuitions did not differ greatly as a function of age, gender, and ethnicity. Scenarios of a second philosophically interesting category, Skeptical Pressure cases, were also largely

seen as exemplifying justified true belief without knowledge, again in keeping with most philosophical theorizing about these intuitions.

Although the broad pattern of responses was consistent with the standard philosophical handling of the cases we examined, participants did not exhibit perfect consensus on these cases. Although Standard True Belief cases attracted the highest rates of knowledge attribution, these cases were not always rated as exemplifying knowledge. Some variation here is not particularly surprising, however. In the context of a study in which many vignettes concern deceptive appearances and problematic testimony, participants may have been hesitant to take stories of standard true belief formation at face value (Aichhorn, et al., 2009). Similarly, the Justified False Belief stories were not always judged to involve a failure of knowledge. Although a large majority of participants did deny knowledge for these stories, it is potentially troubling that the rate of knowledge ascription was not lower still. If truth is a necessary condition for knowledge, as most philosophers and psychologists take it to be, then it is unclear why participants would be willing to ascribe knowledge in these cases. Performance factors may have played a role here. Some of the stories were long and relatively confusing, and the multiple-choice comprehension screen questions may sometimes have been passed by chance rather than on the basis of attention to the story. It is also possible that responses reflect an acquiescence bias (Schuman & Presser, 1981). Further work is needed to determine whether these factors play a role.

Although we have demonstrated that the rates of knowledge ascription for Gettier cases and Skeptical Pressure cases are less than for cases of Standard True Belief, the reasons behind this remain unclear. Gettier cases typically involve circumstances that are trickier than they seem to the agent of the case. Unbeknownst to the agent, the clock is broken, the suspect has an identical twin, there are two people with the same name, and so forth (see Appendix A). We do not normally guard against such circumstances, in that ordinary belief formation does not demand cautious double-

checking (e.g. Evans, 2006; Kahneman & Frederick, 2005). In Gettier cases, although the agent's more casual way of thinking results in the same true belief as a more cautious and thorough strategy, the reason this judgment falls short of knowledge might be because we feel we would have adopted a more demanding cognitive strategy. This might be especially true in light of our privileged knowledge about the agent's environment. Ordinarily, we trust a label in a jewelry store saying that a stone is a diamond; if we knew that a dishonest employee had been substituting many diamonds in the store for lesser stones, we might double-check.

Since Skeptical Pressure cases produce similarly negative intuitions about knowledge, this supports the idea that just thinking about a more elaborate cognitive strategy can be enough to set it as a benchmark against which judgment is measured. Once we reflect on the sheer possibility of being tricked by a fake stone, we enter a frame of mind in which it seems appropriate to double-check before making a judgment. As a result, we might judge the agent negatively, based on this standard.

For both Gettier and Skeptical Pressure cases, participants tended to see the agent of the case as justified in believing the key proposition. When prompted to explain what factors played a role in their assessments of justification, many participants described their judgment as a perspective-taking exercise, or "trying to put myself in their shoes." They reported imagining themselves in the place of the agent: "I put myself in their position and see what would I have done in that particular situation"; "I considered how I would think if I were in the same situation." Our capacity for perspective-taking processes has limitations, however, as evidenced by the fact that Justified False Belief stories attracted significantly lower justification ratings than Standard True Belief stories. When we know that the clock is broken and showing the wrong time, it is hard for us to set aside this privileged knowledge in evaluating the justification of the agent who glances at it and forms a mistaken belief (Susan Birch & Bloom, 2004; SA Birch & Bloom, 2007; Nickerson, 1999).

Although there was an overall tendency not to attribute knowledge for Gettier and Skeptical pressure cases, some participants did attribute knowledge in these cases. The reasons for this are unclear, although it is noteworthy that variability in responding was not systematically linked to age, ethnicity, gender, or training in philosophy. Simple performance error might again be an explanation for this variability. These 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 (cf. Keysar, Lin, & Barr, 2003). 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 (e.g., Bob does initially know that Jill owns an American car, prior to the theft), 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 in the story early rather than following it through every twist. Further work is needed to examine and potentially control for these factors. Lastly, we found that individuals higher in empathy were more slightly more likely to be generous in their attributions of justification for Standard True Belief and Gettier cases. To the extent that attributions of justification require endorsing the perspective of another, it makes sense that more empathetic individuals would attribute higher levels of justification in at least some circumstances.

3.1 Conclusions

In evaluating another person's judgment, we can examine both the content of this judgment and the path taken to reach it. The person whose true belief was acquired through a bad path (wishful thinking, a lucky guess) is not ordinarily seen as knowing. Gettier cases introduce another level of complexity: a path that would ordinarily be good enough for knowledge (say, trusting the word of a

generally reliable co-worker) may not be good in these circumstances, even if the belief that it produces turns out to be true. Because these cases separate dimensions of epistemic appraisal that often run together, they have the potential to show us something new about the nature of these appraisals. One might have thought that intuitive attributions of knowledge would be composed in a straightforward fashion from evaluations of the agent's inner perspective ("Is his belief justified?") and his outer environment ("Is his belief true?"). Intuitions that knowledge is absent in Gettier cases challenge this simple compositional approach. If the real relationship between intuitive evaluations of knowledge and belief is more complex than the classical JTB theory had maintained, its complexity might be illuminated by further study of these intriguing intuitive responses.

4. Appendices

4.1 Appendix A: Sample stimuli (presented to participants without the descriptive labels):

1. Q: Does Emma know whether or not the stone is a diamond?

- A. (Gettier case – Authentic Evidence): Emma is shopping for jewelry. She goes into a nice-looking store, and selects a diamond necklace from a tray marked “Diamond Earrings and Pendants”. “What a lovely diamond!” she says as she tries it on. Emma could not tell the difference between a real diamond and a cubic zirconium fake just by looking or touching. In fact, this particular store has a very dishonest employee who has been stealing real diamonds and replacing them with fakes; in the tray Emma chose almost all of the pendants had cubic zirconium stones rather than diamonds (but the one she chose happened to be real).
- B. (Skeptical Pressure variant): Emma is shopping for jewelry. She goes into a nice-looking store, and spends some time looking at several different displays. She tells the salesperson that she is looking for a diamond necklace with a classic design. She always likes to try things on before she makes up her mind about them, and the salesperson shows her several options. Emma selects a diamond pendant from a tray marked “Diamond Earrings and Pendants”. “What a lovely diamond!” she says as she tries it on. Emma could not tell the difference between a real diamond and a cubic zirconium fake just by looking or touching.
- C. (Standard True Belief): Emma is shopping for jewelry. She goes into a nice-looking store, and spends some time looking at various different displays. She tells the salesperson that she is looking for a simple diamond necklace with a classic design. She always likes to try things on before she makes up her mind about them, and she asks the salesperson to show her quite a wide variety of different items, which he brings out for her one tray at a time. Emma selects a diamond necklace from a tray marked “Diamond Earrings and Pendants”. “What a lovely diamond!” she says as she tries it on.
- D. (Justified false belief): Emma is shopping for jewelry. She goes into a nice-looking store. She looks at several displays, then selects a necklace from a tray marked “Diamond Earrings and Pendants”. “What a lovely diamond!” she says as she tries it on. Emma could not tell the difference between a real diamond and a cubic zirconium fake just by looking or touching. In fact, this particular store has a dishonest employee who has been stealing real diamonds and replacing them with fakes; in the tray Emma chose from, all of the necklaces – including the one she tried on – had cubic zirconium stones rather than diamonds.

2. Q: Does Luke know whether or not one of his co-workers recently vacationed in Las Vegas?

- A. (Gettier Case – Apparent Evidence): Luke works in an office with two other people, Victor and Monica. All winter Victor has been describing his plans to go to Las Vegas on his vacation, even showing Luke the website of the hotel where he has reservations. When Victor is away on vacation, Luke sees Victor’s Facebook photos of himself with Vegas landmarks in the background, together with status updates about how much he is enjoying his trip. When he gets back to work, Victor talks a lot to Luke about how much fun he had vacationing in Las Vegas. However, Victor didn’t really go on the trip; he has just been

pretending. His tickets and reservations were cancelled because his credit card was maxed out, and he secretly stayed home in Markham, very skillfully faking the Facebook pictures using Photoshop. Meanwhile, Monica just spent a weekend vacationing in Las Vegas, but kept this a secret from all her co-workers.

- B. (Skeptical Pressure variant): Luke works in an office with two other people, Victor and Monica. They all get along fairly well, chatting at the water cooler when work is slow. All winter Victor has been telling Luke about his plans to go to Las Vegas on his vacation, even showing Luke the website of the hotel where he has reservations. When Victor is away on vacation, Luke sees Victor's Facebook photos of himself with Vegas landmarks in the background, together with status updates about how much he is enjoying his trip. Victor really did have a good time in Vegas, and was sad to go back home to Markham. When he gets back to work, Victor talks a lot to Luke about how much fun he had on vacation. The resolution on the Facebook photos was low, however; Luke could not tell the difference between real vacation pictures from Las Vegas and fakes created with Photoshop.
- C. (Standard True Belief): Luke works in an office with two other people, Victor and Monica. Their work gets boring sometimes – all three of them are in the accounting department of a large office supply chain – but the salary and benefits are good, and they get along fairly well, chatting at the water cooler when work is slow. All winter Victor has been telling Luke about his plans to go to Las Vegas on his vacation, even showing him the website of the hotel where he has reservations. When Victor is away on vacation, Luke sees Victor's Facebook photos of himself with Vegas landmarks in the background, together with status updates about how much he is enjoying his trip. Victor really did have a good time in Vegas, and was sad to go back home to Markham. When he gets back to work, Victor talks a lot to Luke about how much fun he had on vacation.
- D. (Justified false belief): Luke works in an office with two other people, Victor and Monica. They get along pretty well. All winter long Victor has been talking about his plans to go to Las Vegas on his vacation, even showing Luke the website of the hotel where he has reservations. When Victor is away on vacation, Luke sees Victor's Facebook photos of himself with Vegas landmarks in the background, together with status updates about how much he is enjoying his trip. When he gets back to work Victor talks a lot to Luke about how much fun he had vacationing in Vegas. However, Victor didn't really go on the trip; he has just been pretending. His tickets and reservations were cancelled because his credit card was maxed out, and he secretly stayed home in Markham, very skillfully faking the Facebook pictures using Photoshop. Incidentally, Monica never takes vacations, and all her co-workers know this about her.

4.2 Appendix B: Gettier Case from Weinberg, Nichols and Stich (2001):

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?

5. References

- Aichhorn, M., Perner, J., Weiss, B., Kronbichler, M., Staffen, W., & Ladurner, G. (2009). Temporoparietal junction activity in theory-of-mind tasks: Falseness, beliefs, or attention. *Journal of Cognitive Neuroscience*, 21(6), 1179-1192.
- Ayer, A. J. (1956). *The Problem of Knowledge*. Baltimore: Penguin Books.
- Bartsch, K., & Wellman, H. (1995). *Children talk about the mind*. New York: Oxford University Press.
- Birch, S., & Bloom, P. (2004). Understanding children's and adults' limitations in mental state reasoning. *Trends in cognitive sciences*, 8(6), 255-260.
- Birch, S., & Bloom, P. (2007). The curse of knowledge in reasoning about false beliefs. *Psychological Science*, 18(5), 382.
- Birch, S. A. J. (2005). When Knowledge Is a Curse Children's and Adults' Reasoning About Mental States. *Current Directions in Psychological Science*, 14(1), 25-29.
- BonJour, L., & Sosa, E. (2003). *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues*. Malden, MA: Blackwell.
- Chisholm, R. M. (1957). *Perceiving: A Philosophical Study*. Ithaca: Cornell University Press.
- Clark, M. (1963). Knowledge and Grounds: A Comment on Mr. Gettier's Paper. *Analysis*, 24(2), 46.
- Cullen, S. (2010). Survey-driven romanticism. *Review of Philosophy and Psychology*, 1, 275-296.
- Dienes, Z., & Perner, J. (1999). A theory of implicit and explicit knowledge. *Behavioral and Brain Sciences*, 22(05), 735-808.
- Dretske, F. (1983). *Knowledge and the Flow of Information*. Cambridge: MIT Press.
- Evans, J. S. B. T. (2006). The heuristic-analytic theory of reasoning: Extension and evaluation. *Psychonomic Bulletin & Review*, 13(3), 378 - 395.
- Feldman, R. (1974). An Alleged Defect in Gettier Counterexamples. *Australasian Journal of Philosophy*, 52(1), 68-69.
- Fischhoff, B. (1975). Hindsight is not equal to foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance*, 1(3), 288.
- Fitelson, B. (2010). Strengthening the case for knowledge from falsehood. *Analysis*, 70(4), 666-669.
- Goldman, A. (1967). A Causal Theory of Knowing. *The Journal of Philosophy*, 64, 357-372.
- Goldman, A. (1976). Discrimination and Perceptual Knowledge. *The Journal of Philosophy*, 73(20), 771-791.
- Goldman, A. I. (1979). What is Justified Belief? In G. S. Pappas (Ed.), *Justification and Knowledge* (pp. 1-23). Dordrecht: D. Riedel.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. New York: Oxford University Press.
- Holton, R. (1997). Some telling examples: A reply to Tsohatzidis. *Journal of pragmatics*, 28(5), 625-628.
- Kahneman, D., & Frederick, S. (2005). A model of heuristic judgment. *The Cambridge handbook of thinking and reasoning*, 267-293.
- Keysar, B., Lin, S., & Barr, D. (2003). Limits on theory of mind use in adults. *Cognition*, 89(1), 25-41.
- Kiparsky, P., & Kiparsky, C. (1970). Fact. In M. Bierwisch & K. Erich (Eds.), *Progress in Linguistics* (pp. 143-173). The Hague and Paris: Mouton.
- Lehrer, K. (1965). Knowledge, truth and evidence. *Analysis*, 25(5), 168.
- Miller, S., Hardin, C., & Montgomery, D. (2003). Young children's understanding of the conditions for knowledge acquisition. *Journal of Cognition and Development*, 4(3), 325-356.
- Nagel, J. (2010). Knowledge ascriptions and the psychological consequences of thinking about error. *Philosophical Quarterly*, 60(239), 286-306.
- Nickerson, R. S. (1999). How we know--and sometimes misjudge--what others know: Imputing one's own knowledge to others. *Psychological Bulletin*, 125(6), 737-759.

- Perner, J. (1991). *Understanding the Representational Mind*. MIT Press Cambridge, MA.
- Plantinga, A. (1993). *Warrant: The current debate*. New York: Oxford University Press.
- Plato (1990). *The Theaetetus of Plato* (M. J. Levett, Trans.). Indianapolis: Hackett.
- Pollock, J. L., & Cruz, J. (1999). *Contemporary theories of knowledge* (Vol. 35): Rowman & Littlefield Publishers, Inc.
- Russell, B. (1948). *Human Knowledge: its scope and its limits*. New York: Simon & Schuster.
- Saunders, J. T., & Champawat, N. (1964). Mr. Clark's Definition of 'Knowledge'. *Analysis*, 25(1), 8-9.
- Scheffler, I. (1965). *Conditions of Knowledge*. Chicago: Scott, Foresman.
- Schuman, H., & Presser, S. (1981). *Questions and answers in attitude surveys*: Academic Press New York.
- Shope, R. K. (1983). *The analysis of knowing: A decade of research*: Princeton University Press.
- Sosa, E. (1991). *Knowledge in perspective*. New York: Cambridge University Press.
- Spreng, R. N., McKinnon, M. C., Mar, R. A., & Levine, B. (2009). The Toronto Empathy Questionnaire: Scale development and initial validation of a factor-analytic solution to multiple empathy measures. *Journal of personality assessment*, 91(1), 62-71.
- Starmans, C., & Friedman, O. (2012). The folk conception of knowledge. *Cognition*, xxx-xxx.
- Steup, M. (2008). The analysis of knowledge. *Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/entries/knowledge-analysis/>.
- Vogel, J. (1990). Are there counterexamples to the closure principle? In M. Roth & G. Ross (Eds.), *Doubling: Contemporary Perspectives on Skepticism* (pp. 13-27). Dordrecht: Kluwer.
- Warfield, T. A. (2005). Knowledge from falsehood. *Philosophical Perspectives*, 19(1), 405-416.
- Weinberg, J. S., Nichols, S., & Stich, S. (2001). Normativity and Epistemic Intuitions. *Philosophical Topics*, 29(1), 429-460.
- Williamson, T. (2000). *Knowledge and its Limits*. New York: Oxford University Press.
- Wright, J. C. (2010). On intuitional stability: The clear, the strong, and the paradigmatic. *Cognition*, 115(3), 491-503.