Justified True Belief Triggers False Recall of "Knowing"

Derek Powell (derekpowell@ucla.edu)
Department of Psychology, University of California, Los Angeles
Los Angeles, CA 90095 USA

Zachary Horne (horne2@illinois.edu)
Department of Philosophy, University of Illinois, Urbana-Champaign
Urbana, IL 61802 USA

Angel Pinillos (pinillos@asu.edu)
Department of Philosophy, Arizona State University
Tempe, AZ 85281 USA

Keith J. Holyoak (holyoak@lifesci.ucla.edu)
Department of Psychology, University of California, Los Angeles
Los Angeles, CA 90034 USA

Abstract

Many philosophers once held that knowledge represents a justified true belief. Gettier (1963) challenged this view with thought experiments in which someone has a justified and true belief, but an element of luck is involved that disqualifies the belief from counting as knowledge. We tested lay beliefs using a semantic integration paradigm modeled after that of Gentner (1981). People read stories in which a character 'thought' something was true. On a subsequent recall task, readers sometimes falsely recalled the verb 'thought' as 'knew,' implicitly indicating that the reader had attributed knowledge to the character. False recall of 'knew' was more frequent when the story implied that the belief was true and justified than if it did not. Justified true belief triggered these recall errors even when an element of luck was involved. The present findings suggest that semantic integration provides an empirical paradigm suitable for investigating lay notions about knowledge.

Keywords: belief; knowledge; semantic integration; false memory; experimental philosophy

People's beliefs are the primary drivers of their actions, yet these beliefs are often uncertain—the products of limited information about the world and interconnections between other (often uncertain) beliefs. For this reason, a capacity for evaluating the status of different beliefs is important for individuals in directing their own rational behavior, and for predicting the behavior of others. Understanding these processes requires an analysis of the concept of *knowledge*: the distinction between what is known versus what is merely believed, imagined, hoped for, or assumed. Research on metacognition has examined how people assess their confidence in their own beliefs (e.g., Klayman et al., 1999; Tsai, Klayman & Hastie, 2008), and research on theory of mind (Premack & Woodruff, 1978) has examined how beliefs are attributed to others (e.g., Birch, 2005) but there has been very little psychological research examining the self-assessment and attribution of knowledge.

Making a decision about when to attribute knowledge, either to oneself or to another, hinges on one's conception of knowledge: it is a decision about whether or not the concept applies in a particular instance. Although other factors may play into this decision process, understanding the exact nature of the concept itself is essential to understanding the overall process of knowledge attribution. Philosophers have long contemplated the nature of knowledge, and have also developed a variety of methods for studying concepts. A common method involves thought experiments in which a specific scenario is imagined and a response intuited.

One of the most influential of these thought experiments was proposed by Edmund Gettier (1963). Named for their progenitor, "Gettier cases" challenge the traditional conception of knowledge. Prior to the 1960s, most philosophers thought that knowledge was to be analyzed as justified true belief. Today, many philosophers see Gettier cases as counterexamples to that analysis. Gettier cases are situations in which an agent has a true belief that is justified, but an element of luck is involved that disqualifies their cognitive state from being considered knowledge. To illustrate such a case, suppose that at 3:34pm an agent comes to believe it is 3:34pm by looking at her normally reliable watch. Suppose also that unbeknownst to the agent, her watch had been stopped for exactly 24 hours—she just happened to glance at her watch at the correct time. The agent's belief is not only true, but is also justified (since looking at one's normally reliable watch is a good way to form veridical beliefs about time of day). However, most philosophers judge that this agent does not *know* that it is 3:34pm, because her belief is true only by luck. If this judgment is correct, then this case is a counterexample to the traditional thesis that knowledge is justified true belief.

It remains an open question whether philosophers' conceptions of knowledge are shared by laypeople. Recently, experimental philosophers have sought to empirically investigate laypeople's intuitive judgments about philosophical thought experiments (for a review, see Knobe et al., 2012). In particular, researchers have turned their attention to examining whether philosophers' Gettier judgments are shared by laypeople (Weinberg, Nichols & Stich, 2001; Turri, forthcoming). Starmans and Friedman (2012) investigated laypeople's evaluations of Gettier cases by presenting participants with short vignettes that described agents forming beliefs under different circumstances. Three different versions of each scenario were created: the agents in the vignettes either formed a false belief, formed a justified true belief, or were "Gettiered"—the belief they formed was both justified and true, but was true only by luck. Starmans and Friedman then asked participants to judge whether the agents "knew" or "only believed" the proposition in question, and to rate how confident they were in their judgment. Participants attributed knowledge to agents in Gettier cases almost as readily as they did in cases of non-Gettiered justified true belief, suggesting that laypeople's concepts of knowledge may differ from those held by many philosophers. Their findings led Starmans and Friedman to conclude that laypeople view knowledge as justified true belief, in accord with the more traditional philosophical view.

However, the survey-based methodology used by Starmans and Friedman (2012) has limitations. Answers on such surveys may be influenced by demand characteristics (Orne, 1962). For example, if participants form some interpretation of the experimenter's hypothesis, they may attempt to confirm or disconfirm this hypothesis. In addition, participants commonly display apprehension about being evaluated (Weber & Cook, 1972). This may lead them to give responses they perceive as either socially desirable, or likely to be considered "correct," irrespective of their actual attitudes or answers. For research on folk concepts, this type of evaluative apprehension might be manifested as a sort of amateur philosophizing, or attempts to avoid being "tricked" by the experimenter.

A further concern with survey-based methods is that asking participants to make knowledge attributions fails to isolate their knowledge concepts from other decision processes that could influence their judgments downstream. This is a general problem for survey-based methods that applies not only to investigations of concepts for knowledge, but to investigations of lay concepts in general. Recently, social psychologists and experimental philosophers have investigated such concepts as intentional action (e.g., Knobe et al., 2012), causation (e.g., Livengood & Machery, 2007), and explanation (Braverman et al., 2012), but thus far their methods have been primarily survey-based. In fact, almost all work in experimental philosophy utilizes these methods.

Here we propose and test a new method for examining people's concepts, based on psychological research related to *semantic integration*. Research has shown that new memories are not formed in isolation, but are connected with other memories of past experiences as well as semantic knowledge. Even in the simplest of experimental contexts, such as learning lists of words, experiences are processed and given meaningful representations. For example, Roediger and McDermott (1995; also Deese, 1959) asked participants to memorize lists of words. The lists included many different words that were all semantically related to a single target word, the "critical lure." When participants were later asked to recall the words they had been presented with, they were about as likely to falsely recall the critical lure, which had never been presented, as any of the other words that actually appeared in the list.

Cognitive psychologists and psycholinguists have also applied false memory paradigms to research on language processing (e.g., Bransford & Franks, 1971; Flagg, 1976; Gentner, 1981). Researchers have been interested in how various pieces of semantic information are combined to form meaningful structured representations, or "discourse meanings," during language processing. Early research by Sachs (1967) revealed that memory for the meaning of sentences is more robust than memory for their specific wordings. During language processing, the original form of presented material is only stored temporarily, just long enough to be comprehended, whereas the material's meaning is encoded into long-term memory.

Bransford and Franks (1971) reasoned that if semantic information is integrated during language processing, and it is the meaning of a passage that is actually encoded into memory, then human memory ought to exhibit productivity. That is, it should be possible for exposure to several basic, interrelated propositions to produce false verbatim memory for more complex propositions that express their combination, even when these propositions were never themselves experienced. These sorts of productive memory errors have been taken as evidence for semantic integration across a variety of language comprehension contexts (Flagg, 1976; Owen, Bower & Black, 1979; Gentner, 1981; Sulin & Dooling, 1974; Thorndyke, 1976). To explain these findings, Gentner (1981) examined a model of language processing in which linguistic propositions are considered both individually and in the broader context of the story in which they appear. Her model assumes that when a sentence is read within the context of a larger passage, the discourse meaning that a reader forms may incorporate information not contained in the original sentence.

In evaluating this model, Gentner (1981) focused her investigations on a relatively well-analyzed area of linguistics, the meanings of verbs. She was able to make specific predictions about how manipulations of the contextual information given in a passage of text would affect later recall for verbs within that passage. To illustrate, consider the relationship between the general verb 'give' and the more specific verb 'pay'. An informal analysis suggests that 'to give' some item is to take some action that transfers ownership of that item to a recipient. 'Paying' is a more specific form of giving, in which the giver *owes* the

recipient. In her experiments, Gentner (1981) asked her participants to read paragraph-long stories that each featured a critical sentence containing some key verb of interest. For instance, one of these stories contained the critical sentence, "Max finally gave Sam the money." Two versions of this story were created, one that contained additional context explaining that Max owed Sam money, and a control story that lacked this information. After reading one version of the story, Gentner's participants performed a recall task, in which they were shown the critical sentence with the word 'gave' removed, and asked to fill in the word that had appeared in the story. In support of her predictions, Gentner found that participants who had been provided with the additional contextual information were more likely to falsely recall the more specific verb 'paid' as having appeared in the critical sentence than participants who had read the control story.

We aimed to turn this methodology on its head: Whereas Gentner (1981) used a false recall paradigm to examine how known semantic structures are integrated during language processing, we used a similar paradigm to examine the disputed semantic structure of the concept 'knowledge'. Following Gentner (1981), we constructed stories containing the generic cognition verb 'thought', and used false recall of the more specific verb 'knew' as a measure of the extent to which different contexts instantiate the semantic structure of 'knowledge'. By incorporating the relevant contextual information into our stories, we can examine whether the different scenarios imagined in various thought experiments differentially activate people's concept of knowledge.

Gentner's (1981) paradigm has several qualities that are desirable for our present purposes, relative to other semantic integration tasks. First, the use of free recall makes its results more compelling than tasks that rely on recognition judgments. Participants' responses to recognition tasks can be influenced by both true recollection as well as feelings of familiarity (Tulving, 1985). In contrast, explicit recall of the word 'knew' provides stronger evidence for the semantic activation of the concept. Second, this paradigm focuses responses onto a single specific word of interest, whereas other semantic integration paradigms often ask participants to evaluate larger semantic units, such as phrases or sentences (Waskan et al., under review). This specificity may help reduce ambiguity in investigations of individual concepts.

Experiment 1 serves as a proof of concept, demonstrating that semantic integration can be used to investigate laypeople's concept of knowledge. In Experiment 2 we examined the more controversial issue of how Gettier cases activate people's concept of knowledge.

Experiment 1

Method

In Experiment 1 we constructed two similar stories about a detective investigating a crime. In the first story, the detective forms the justified true belief (JTB condition) that his suspect is guilty: the omniscient narrator reveals that the detective's suspect committed the crime, and the detective uncovers evidence that his suspect is guilty. In the second story, the detective forms the unjustified belief (UB) that his suspect is guilty: he cannot find any useful evidence linking his suspect to the crime, and the narrator does not reveal whether the suspect is guilty. Following Gentner (1981), each story included a critical sentence (shown below) that itself contained a critical word (bolded below). Here, the critical word in both stories was 'thought', a generic cognition verb that could plausibly be recalled as 'knew', should the right conditions be met. This critical sentence was later presented with a blank in place of the critical word, and participants were asked to recall the word that appeared in the story.

Critical sentence: "Whatever the ultimate verdict would be, Dempsey **thought** Will was guilty."

Following prior research on semantic integration and false recall paradigms (Gentner, 1981) we predicted that participants who read the story in which Dempsey's belief is justified and true would be more likely to falsely recall the word "knew" than participants who read the story in which Dempsey's belief is unjustified. Of course, it will not come as a surprise if a justified true belief more closely resembles people's concept of knowledge than an unjustified belief. Experiment 1 examined this simple case in order to demonstrate the potential of the semantic integration paradigm.

Participants This experiment was conducted online, with 147 participants (91 female) recruited from Amazon's Mechanical Turk (mTurk). The mean age of the participants was 34 years. They were all paid $0.50 for their participation.

Procedure Participants were randomly assigned to the JTB and UB conditions, and were asked to read the corresponding story about the detective. Then, participants completed a distraction task, reading an approximately 1000-word selection from a fictional article on gamma ray bursts (taken from Waskan et al., under review). Timing controls ensured that participants spent an adequate amount of time attending to each section of the experiment.

In the recall task, participants were shown five sentences from the detective story, each missing one word that was replaced with an underscored blank space. They were instructed to type in the word that originally appeared in the story. The critical sentence was always presented first.

After the recall task, participants were asked a direct question to assess their understanding of whether Dempsey had knowledge or not. Following Starmans and Friedman (2012), they were asked: "Would you say that Dempsey knew Will was guilty, or only thought Will was guilty?" They indicated their choices as 'knew' or 'thought,' and then rated their confidence on a 1-5 Likert scale.

Results and Discussion

Recall task Participants' responses during the recall task were classified as either 'thought'-type responses or 'knew' responses. Words and phrases synonymous with 'thought' but neither stating nor implying knowledge were grouped together as 'thought'-type responses. In an effort to remain conservative, *only* the word 'knew' was counted toward the tally of 'knew' responses. Responses that were nonsense (i.e., were not verbs, were random letters typed in the blank, etc.) were excluded from analysis. After these exclusions, 64 participants remained in the JTB condition and 65 participants remained in the UB condition.

As predicted, participants recalled the word 'knew' significantly more often when they were assigned to the JTB condition than when they were assigned to the UB condition (39% vs. 18%; $X^2 = 5.72$, $p = .016$). This finding demonstrates that participants semantically integrated contextual information, specifying that Dempsey's belief was both true and justified, with the generic cognition verb 'thought,' leading to false recall of the more specific verb 'knew.'

Knowledge survey question As in Starmans and Friedman (2012), participants' 'knew' and 'thought' responses were assigned scores of 1 and -1 respectively, and these values were multiplied by the confidence ratings participants reported to produce a knowledge rating score. Knowledge ratings from participants in the JTB condition (-1.40) were significantly higher than in the UB condition (-3.66), $t(127) = 6.94$, $p < .001$.

Experiment 2

Method

In Experiment 2 we used the semantic integration paradigm to examine a more substantive question about knowledge. Specifically, we investigated the extent to which Gettier cases activate people's concept of knowledge. Experiment 2 used three stories, adapted from the detective stories of Experiment 1. In the first story, one character "Will" is guilty of a crime and "Dempsey," the detective in the story, finds authentic evidence of his guilt, forming the justified true belief that he is guilty (JTB condition). Meanwhile, another character "Beth", who is Will's girlfriend, observes the sequence of events that unfold and result in Dempsey thinking that Will is guilty. In the second story, Will is innocent of the crime, but is framed by his girlfriend Beth because she suspects that he is cheating on her. Dempsey finds evidence planted by Beth, and as a result forms the *false* belief that Will is guilty of the crime (FB condition). Finally, in the third story, Will is guilty of the crime, but he has eliminated all the authentic evidence of his crime. Beth, as part of a ploy to seek reprisals against Will, plants evidence that implicates him in the crime. Dempsey finds this evidence and forms the belief that Will is guilty. In this case, Dempsey's belief is both justified and true, but is only true by chance (Gettier condition).

Participants Experiment 2 was also conducted online, with 304 participants (164 female) recruited from mTurk. The mean age of participants was 31 years. All participants were paid $0.50 for their participation.

Procedure Participants were randomly assigned to one of the three conditions, and were asked to read the corresponding story about the detective. The distractor and recall tasks for Experiment 2 were the same as those in Experiment 1. After the recall task, participants were asked the same question about Dempsey's knowledge as in Experiment 1, and were also asked, "Should Dempsey have arrested Will?" Participants rated their confidence for both responses. After these questions, participants answered a pair of comprehension questions to ensure they had attended to central details of the story.

Results and Discussion

Recall task Some participants were excluded from analysis after failing the reading comprehension check. Others gave ambiguous free recall responses that did not fit into either the 'knew' or 'thought' response categories. After these exclusions, 259 participants remained in the final analysis. Participants' recall responses were classified according to the same criteria as in Experiment 1. Figure 1 shows the proportion of 'knew' responses in the three conditions.
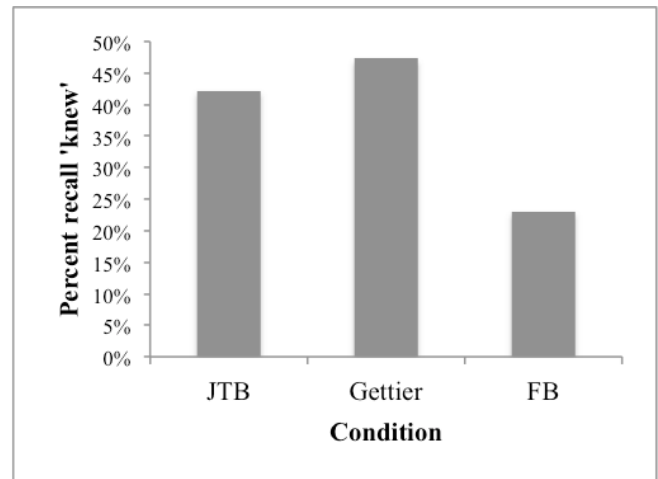

Figure 1: Percent false recall of 'knew' in critical sentence across conditions

False recall of 'knew' was observed significantly more often in the JTB and Gettier conditions as compared with the FB condition ($X^2 = 5.94$, $p = .015$, and $X^2 = 9.63$, $p < .01$, respectively). Interestingly, no significant difference was found between the frequencies of 'knew' recall in the JTB and Gettiered conditions ($X^2 = .20$, $p = .66$). Thus participants seemed to believe that agents in Gettier cases possess knowledge, and apparently drew no distinction between Gettier cases and non-Gettier cases of justified true

belief. This finding accords with those reported by Starmans and Freidman (2012), and stands in contrast to how epistemologists have understood the implications of Gettier cases.

Survey questions Participants' knowledge rating scores were calculated as in Experiment 1. This same procedure was also employed with the "arrest" question to calculate an action rating score, where a positive score indicated that participants endorsed Dempsey's arresting Will.

Participants' knowledge ratings differed significantly between conditions, $F(2,256) = 18.12$, $p < .001$. Post hoc tests using Bonferroni corrections indicated that participants knowledge ratings in both the JTB (0.62) and Gettier conditions (-0.11) were higher than those in the FB condition (-2.56; $p < .001$), but that knowledge ratings did not differ significantly between the JTB and Gettier conditions ($p = .65$). This finding is consistent with the results of the recall task, and replicates the pattern of results reported by Starmans and Friedman (2012). It is worth noting, however, that participants' actual knowledge rating scores are considerably lower in the present experiment than in that reported by Starmans and Friedman. This is likely due to the different experimental materials used, and in particular the quality of evidence depicted in the different stories: Starmans and Friedman's vignettes feature direct perceptual evidence, whereas our materials describe weaker physical and testimonial evidence.

Significant differences between conditions were also found for action ratings, $F(2, 256) = 52.81$, $p < .001$, where planned comparisons revealed the same pattern as for knowledge ratings: participants action ratings differed between the JTB (3.26) and FB (-0.12) conditions ($p < .001$), and between the Gettier (3.80) and FB conditions ($p < .001$), but not between the JTB and Gettier conditions ($p = .69$). Interestingly, participants endorsed Dempsey's action (arresting Will) even when Dempsey had been Gettiered. A number of philosophers hold that knowledge is intimately connected to action (Hawthorne & Stanley, 2008). In particular, they hold that it is intuitive that if a person knows some proposition, then it is acceptable for them to use that proposition in reasoning and in action. Our result is consistent with this thesis: in conditions where participants say that Dempsey knows Will is guilty, they also tend to say that he should have arrested Will.

General Discussion

Experiment 1 provided a proof-of-concept for the use of semantic integration tasks in examining people's concept of knowledge. As expected, a story in which a character forms a justified true belief activated people's concept of knowledge more strongly than a story in which this same character forms an unjustified belief, as evidenced by an increase in false recall of the verb 'knew' in place of 'thought.'

The results of Experiment 2 corroborated Starmans and Friedman's (2012) findings on laypeople's reactions to Gettier cases, while avoiding alternative interpretations that might be raised with survey-based methods. As our participants believed they were completing a memory task, it is very unlikely that their responses were affected by unwanted demand characteristics, or are indicative of some sort of performance error (Kauppinen, 2007; Cullen, 2010). Rather, we interpret their responses as reflecting their concept of knowledge, which apparently differs from the concept developed by philosophers who have considered Gettier cases.

Together, these results have important implications for modern philosophical research. After the work of Gettier (1963), not only do epistemologists almost universally assume that knowledge is *not* equivalent to justified true belief, but many of them also assume that they are investigating a conception of knowledge that is shared by laypersons (or at least that lay judgments inform philosophical theories of knowledge). In light of our findings, as well as those of Starmans and Friedman (2012), it appears that one of the above assumptions must be rejected.

The present findings also have implications for research on knowledge assessment and attribution. Broadly speaking, empirical findings indicate that research on knowledge assessment and attribution cannot rely on the conceptions of knowledge developed by philosophers as accurate guides to lay conceptions of knowledge. Rather, these findings underscore the need for empirical examinations of laypeople's concept of knowledge.

The present study demonstrates that semantic integration tasks provide a promising methodology for empirically investigating lay concepts, avoiding many of the pitfalls associated with survey-based methods. Semantic integration tasks minimize the likelihood that participants' responses are affected by demand characteristics, and make it possible to isolate the activation of concepts from downstream decision processes.

Directions for Future Research

The fact that 'knowledge' has a verb form, as well as approximate near-synonyms such as 'thought' and 'believed', allowed us to model our investigations of knowledge directly on Gentner's (1981) research on the semantic integration of verb meanings. Of course, not all concepts of interest to psychologists and philosophers will necessarily exhibit these desirable traits. Where this is not the case, other semantic integration tasks may be more appropriate. For example, these constraints would not apply to semantic integration tasks measuring recognition for sentences or phrases (e.g., Bransford & Franks, 1971; Owens, Bower & Black, 1979). As described earlier, this type of task presents some disadvantages (e.g., increased ambiguity from the assessment of recognition memory over larger semantic units). However, these disadvantages are not insurmountable. In particular, employing a remember-know procedure (Tulving, 1985) could help distinguish between genuine recollection and familiarity. With sufficient care, it

should be possible to craft phrases or sentences that unambiguously express whatever concept may be of interest to researchers (e.g., Waskan et al., under review).

The present investigation demonstrates the need for empirical research on knowledge attribution and knowledge concepts, and also illustrates a powerful method that may be applicable in future investigations. Further research is needed to explore both the implications of these early findings on Gettier cases, as well as other factors that are potentially relevant to knowledge, such as the salience of error (e.g., Schaffer & Knobe, 2012) and agents' practical interests (e.g., Pinillos, 2012). Moreover, semantic integration tasks may prove useful for the study of other philosophical concepts, including intention (Knobe et al., 2012) and causation (Livengood & Machery, 2007).

Acknowledgments

References

Birch, S. A. J. (2005). When knowledge is a curse. *Current directions in psychological science*, *14*(1), 25–29.

Bransford, J. D., & Franks, J. J. (1971). The abstraction of linguistic ideas. *Cognitive Psychology*, *2*, 331–350.

Braverman, M., Clevenger, J. Harmon, I., Higgins, A., Horne, Z., Spino, J., & Waskan, J. (2012). Intelligibility is necessary for explanation but accuracy may not be. *Proceedings of the Thirty-Fourth Annual Conference of the Cognitive Science Society. Sapporo, Japan.*

Cullen, S. (2010). Survey-driven romanticism. *Review of Philosophy and Psychology*, *1*, 275–296.

Deese, J. (1959). On the prediction of occurrence of particular verbal intrusions in immediate recall. *Journal of Experimental Psychology*, *58*, 17-22.

Flagg, P. (1976). Semantic integration in sentence memory? *Journal of Verbal Learning and Verbal Behavior*, *15*, 491–504.

Gentner, D. (1981). Integrating verb meanings into context. *Discourse Processes*, *4*, 349–375.

Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis*, 23, 121-123.

Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *Journal of Philosophy*, *105*, 571.

Kauppinen, A. (2007). The rise and fall of experimental philosophy. *Philosophical Explorations*, *10*, 95-118.

Klayman, J., Soll, J., González-Vallejo, C., & Barlas, S. (1999). Overconfidence: It depends on how, what, and whom you ask. *Organizational Behavior and Human Decision Processes*, *79*, 216–247.

Knobe, J., Buckwalter, W., Nichols, S., Robbins, P., Sarkissian, H., & Sommers, T. (2012). Experimental philosophy. *Annual Review of Psychology*, *63*, 81–99.

Livengood, J., & Machery, E. (2007). The folk probably don't think what you think they think: Experiments on causation by absence. *Midwest Studies in Philosophy*, *XXXI*, 107 – 127.

Orne, M. (1962). On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. *American Psychologist*, *17*(11), 776–783.

Owens, J., Bower, G., & Black, J. (1979). The "soap opera" effect in story recall. *Memory & Cognition*, *7*, 185–191.

Pinillos, A. (2012). Knowledge, experiments, and practical interests. In J. Brown & M. Gerkin (Eds.), *Knowledge Ascriptions* (pp. 192-219). New York: Oxford University Press.

Premack, D., & Woodruff, G. (1978). Does the chimpanzee have a theory of mind?. *Behavioral and brain sciences*, *1*, 515-526.

Roediger, H. L., & McDermott, K. B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 803-814.

Sachs, J. (1967). Recognition memory for syntactic and semantic aspects of connected discourse. *Perception & Psychophysics*, *2*, 437–442.

Schaffer, J., & Knobe, J. (2012). Contrastive knowledge surveyed. *Noûs*, *46*(4), 675-708.

Starmans, C., & Friedman, O. (2012). The folk conception of knowledge. *Cognition*, *124*, 272–283.

Sulin, R. A., & Dooling, D. J. (1974). Intrusion of a thematic idea in retention of prose. *Journal of Experimental Psychology*, *103*, 255–262.

Thorndyke, P. (1976). The role of inferences in discourse comprehension. *Journal of Verbal Learning and Verbal Behavior*, *15*, 437–446.

Tsai, C., Klayman, J., & Hastie, R. (2008). Effects of amount of information on judgment accuracy and confidence. *Organizational Behavior and Human Decision Processes*, *107*, 97–105.

Tulving, E. (1985). Memory and consciousness. *Canadian Psychology/Psychologie Canadienne*, *26*, 1-12.

Turri, J. (forthcoming). A conspicuous art: Putting Gettier to the test. *Philosopher's Imprint*.

Waskan, J., Clevenger, J., Harmon, I., Horne, Z., & Spino, J. (under review). Explanatory anti-psychologism overturned by lay and scientific case classifications.

Weber, S., & Cook, T. (1972). Subject effects in laboratory research: An examination of subject roles, demand characteristics, and valid inference. *Psychological Bulletin*, *77*, 273–295.

Weinberg, J. M., Nichols, S., & Stich, S. (2008). Normativity and epistemic intuitions. *Experimental Philosophy*, *29*, 17-45.