UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title
Knowledge and Luck

Permalink

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 36(36)

ISSN
1069-7977

Authors
Turri, John
Blouw, Peter
Buckwalter, Wesley

Publication Date
2014

Peer reviewed

Knowledge and Luck

John Turri (john.turri@gmail.com)
Department of Philosophy, 200 University Avenue West
Waterloo, ON N2L 3G1 Canada

Peter Blouw (pblouw@uwaterloo.ca)
Department of Philosophy, 200 University Avenue West
Waterloo, ON N2L 3G1 Canada

Wesley Buckwalter (jwbuckwa@uwaterloo.ca)
Department of Philosophy, 200 University Avenue West
Waterloo, ON N2L 3G1 Canada

Abstract

Nearly all success is due to some mix of ability and luck. But some successes we attribute to the agent's ability, whereas others we attribute to luck. To better understand the criteria distinguishing credit from luck, we conducted a series of studies on knowledge attributions. Knowledge is an achievement that involves reaching the truth. But many factors affecting the truth are beyond our control and reaching the truth is often partly due to luck. Which sorts of luck are compatible with knowledge? We find that knowledge attributions are highly sensitive to lucky events that change the explanation for why a belief is true. By contrast, knowledge attributions are surprisingly insensitive to lucky events that threaten but ultimately fail to change the explanation for why a belief is true. These results shed light on our concept of knowledge, help explain apparent inconsistencies in prior work on knowledge attributions, and constitute progress toward a general understanding of the relation between success and luck.

Keywords: luck; success; evaluation; knowledge attribution

Introduction

Reasoning about what people know is central to our lives and is often essential for predicting and evaluating human action. Whether someone knows that a certain action is impermissible affects whether we excuse or punish them for it (Hart, 1959; Duff, 1990; Malle & Nelson, 2003). Moreover, if you know something, then you're entitled to base actions on it and communicate it to others. By contrast, if you don't know something, then it seems that you should be more cautious about basing actions on it or communicating it to others (Unger, 1975; Williamson, 2000; Hawthorne, 2004; Stanley, 2005; Fantl & McGrath, 2009; Turri, 2011a; Turri, 2013a; Buckwalter & Turri, under review). Indeed, researchers have recently argued that knowledge sets the standard for the two main forms of human pedagogy, assertion (telling someone a fact) and instructional demonstration (showing someone how something is done), which makes knowledge an essential ingredient of human culture and civilization (Buckwalter & Turri, in press).

Knowledge is an achievement that involves reaching the truth through cognitive ability (Aristotle, 350 BCE; Reid,

1785; Sosa, 2007; Greco, 2010). Nearly all human achievement, including knowledge, is due to some mix of ability and luck. Ludwig Wittgenstein had in mind this inevitable role of luck when he wrote, "It is always by favour of Nature that one knows something" (1975: §505). But not just any sort of luck is compatible with knowing (Plato, 380 BCE; Russell, 1948, p. 113). For example, blindly guessing the truth won't yield knowledge. But beyond this obvious starting point, which sorts of luck are viewed as compatible with knowledge?

Insightful work in philosophy and psychology has focused on one or another aspect of the relationship between knowledge and luck (Gettier, 1963; Unger, 1968; Harman, 1973; Goldman, 1976; Engel, 1992; Weinberg, Nichols, & Stich, 2001; Pritchard, 2005; Cullen, 2010; Wright, 2010; Beebe & Shea, 2013; Buckwalter, 2013; Starmans & Friedman, 2012; Nagel, San Juan, & Mar, 2013; Starmans & Friedman 2013), but the relationship has not been systematically investigated from a psychological perspective. Other psychological work has investigated how a neighboring concept related to luck, "deviant causation" (Searle, 1983), affects people's attributions of intentional action and judgments of moral responsibility (Malle & Knobe, 1997; Pizarro et al., 2003; Knobe, 2003; Malle, 2006; Mele & Cushman, 2007). Recent work has also identified at least one form of luck, broadly related to deviant causation, that doesn't seem to diminish knowledge attribution (Starmans & Friedman, 2012), although subsequent work has claimed to undermine these findings (Nagel, San Juan, & Mar, 2013).

Philosophers and psychologists have developed many thought experiments intended to pump intuitions about the relationship between knowledge and luck (e.g. Shope, 1983; Beebe & Buckwalter, 2010; Wright, 2010; Starmans & Friedman, 2012; Turri, in press). These thought experiments vary widely and focus on different features that different authors have considered relevant to luck. Unfortunately, contributors to the debate have failed to exercise appropriate control over the thought experiments. The variety of cases exhibits three important structural differences. First, many differ in whether the agent initially *perceives* a state of affairs that makes his or her belief true (a "truth-maker," for

short). In some examples, the agent perceives a truth-maker, but in others the agent perceives a convincing fake. Second, many examples differ in whether the agent's perceptual relation *remains intact* throughout. Sometimes the agent perceives a certain truth-maker and events threaten to *disrupt* that truth-maker, but the threat ultimately fails. Other times, the threat succeeds in disrupting the original truth-maker, which is then replaced by a "backup" truthmaker. Third, many examples differ in how closely the originally perceived truth-maker and backup truth-maker *resemble* one another. Sometimes they very closely resemble one another, while other times they differ greatly.

Theorists have argued at length over the intuitively correct verdict in each of these cases — does the agent *know* or *only believe* the truth? — and the best explanation for that verdict (for an overview, see Shope, 2002). We will not enter into the theoretical debate here, at least not directly (see the General Discussion for ways that our results could indirectly shed light on the theoretical debate). Instead, we are interested in whether people's ordinary judgments about knowledge are sensitive to the three luck-related factors identified in the previous paragraph: *threat*, *disruption*, and *replacement*.

More specifically, we will test the effect of the following things on ordinary knowledge attributions: (i) A threat to the detection of truth: When an agent perceptually detects a truth-maker, what is the effect of a salient threat to her ability to detect it? (ii) The disruption of a truth-maker: What is the effect of a successful threat that fundamentally changes the explanation for why her belief is true? (iii) The resemblance between the initial truth-maker and the backup truth-maker that replaces it: What is the effect of making them very similar or dissimilar?

The results of our investigations will highlight some general lessons about the ordinary concept of knowledge and provide a framework for integrating prior psychological findings on knowledge attributions and resolving some potentially troubling inconsistencies in the literature. Ultimately it turns out that there is no true and general lesson to be drawn about the relationship between knowledge and luck. Instead, there are several more specific lessons about the relationship between knowledge and different luck-related factors, each of which affect knowledge attributions in interesting and importantly different ways. Our results will also shed light on the potential theoretical usefulness of a peculiar class of cases, known as "Gettier cases," that have recently generated controversy in the psychological literature (Starmans & Friedman, 2012; Turri, 2013b; Nagel, San Juan, & Mar, 2013; Nagel, Mar, & San Juan 2013; Starmans & Friedman, 2013). The General Discussion examines in greater detail the relationship between our findings and prior work on these issues.

Experiment 1

Method

Participants. Eight hundred and thirteen new participants were tested (aged 18–75 years, mean age = 31.25, 326 female, 96% reporting English as a native language). We excluded data from 85 participants who failed comprehension questions. Including data from these participants did not affect the results reported below.

Materials and Procedure. Participants were randomly assigned to one of seven conditions and read a single story. (See Table 1.) Each condition featured a different story. The basic storyline featured an agent, Emma, admiring jewelry in a fancy department store (based on a story in Nagel, San Juan, & Mar, 2013). Emma purchases a stone from the diamond display, puts it in her pocket, browses for another minute, then leaves the store. The different versions of the story vary whether the stone is a real diamond or a fake, whether there is a threat to the stone remaining in Emma's pocket, whether the threat fails or succeeds, and whether any other stone also ends up in Emma's pocket. The different stories manipulated whether the Emma detects an initial truth-maker, whether Emma's truth-detection is saliently threatened, whether the threat succeeds in disrupting the initial truth-maker, and whether the backup truth-maker is highly similar or dissimilar to the initial.

In all the stories, Emma purchases a stone from a jewelry store, puts it in her pocket, and soon walks out of the store. In all the stories that involve detection, the stone she purchases is a diamond. In all the stories that do not involve detection, the stone is a fake. In all the stories that involve similar backup truth-makers, the backup truth-maker is that, one way or another, a diamond is put into Emma's pocket before she leaves the store. In all the stories that involve dissimilar backup truth-makers, the backup truth-maker is that a real diamond was secretly sewn into Emma's pocket by a previous owner long ago.

After reading the story, participants answered a series of comprehension questions about whether Emma has a diamond in her pocket when leaving the store, whether it's reasonable for Emma to think that there is a diamond in her pocket, and why Emma thinks that there is a diamond in her pocket. All participants then answered the same test question about whether, as Emma leaves the store, she knows that there is a diamond in her pocket. Questions were always asked in the same order and response options rotated randomly.

Results and Discussion

Gender did not affect attribution rates (Male, 54%; Female, 53%), Fisher's exact test, p = 0.886, n.s., so the analyses that follow collapse across gender.

Assignment to condition affected rates of knowledge attribution: 12–90%, χ^2 (df = 6, N = 813) = 242.21, $p <$ 0.001, Cramer's V = 0.546 (Figure 1). Rates of knowledge attribution did not differ between Knowledge Control (90%)

Table 1: Experiment 1: Description of stories used in the seven conditions.

| Condition	Description
1.Knowledge Control	The stone Emma purchases is a diamond. She walks out of the store and nothing else happens.
2.Failed Threat	The stone Emma purchases is a diamond. A skilled jewel thief tries to steal it from her pocket before she leaves the store, but he fails.
3.Detection Similar Replacement	The stone Emma purchases is a diamond. A skilled jewel thief tries to steal it from her pocket before she leaves the store, and he succeeds. Someone secretly slips a diamond into Emma's pocket before she leaves the store.
4.Detection Dissimilar Replacement	The stone Emma purchases is a diamond. A skilled jewel thief tries to steal it from her pocket before she leaves the store, and he succeeds. Long ago, Emma's grandmother secretly sewed a diamond into the pocket of Emma's coat.
5.No Detection Similar Replacement	The stone Emma purchases is a fake. A skilled jewel thief tries to steal it from her pocket before she leaves the store, and he succeeds. Someone secretly slips a diamond into Emma's pocket before she leaves the store.
6.No Detection Dissimilar Replacement	The stone Emma purchases is a fake. A skilled jewel thief tries to steal it from her pocket before she leaves the store, and he succeeds. Long ago, Emma's grandmother secretly sewed a diamond into the pocket of Emma's coat.
7.Ignorance Control	The stone Emma purchases is a fake. She walks out of the store and nothing else happens.
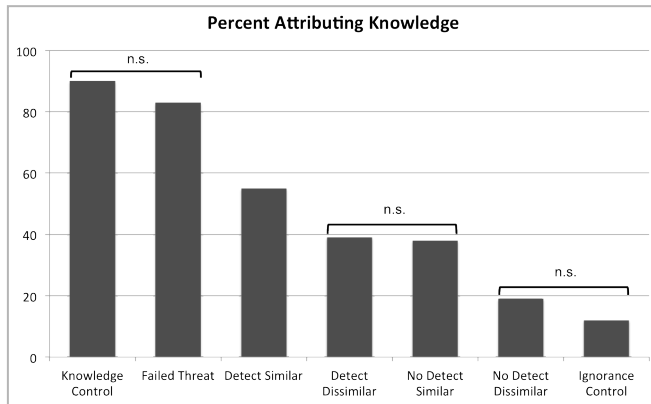
Figure 1: Experiment 1: Percent of participants attributing knowledge across conditions. Except where non-significance is indicated, significance for all comparisons at the p < 0.01 level.

and Failed Threat (83%), Fisher's, $p = .113$, n.s., and far exceeded chance rates in both conditions, binomial tests, ps < 0.001. A failed threat to truth-detection does not significantly depress knowledge attribution and is viewed as fully consistent with knowledge. Knowledge attribution in Detect Similar (55%) was significantly lower than in Knowledge Control, Fisher's, $p < .001$, Cramer's V = .367, non-significantly above chance, binomial, $p = .195$, and significantly higher than in Ignorance Control, Fisher's, $p < .001$, Cramer's V = .416. Rates of knowledge attribution differed significantly depending on whether the replacement truth-maker was similar or dissimilar to the initial truth-maker. This was true regardless of whether the agent had initially detected a truth-maker: rates of knowledge attribution were higher in Detect Similar than in Detect

Dissimilar (39%), Fisher's, $p = .016$, Cramer's V = .156, and they were higher in No Detect Similar (38%) than in No Detect Dissimilar (19%), Fisher's, $p = .006$, Cramer's V = .207.

Replication of Findings: Experiments 2-4

We replicated each of the results reported in Experiment 1 with separate additional studies. The results not only replicate earlier findings but also demonstrate that the findings generalize to other contexts and are not due to the use of a particular type of story. Experiment 1 featured stories set in indoor social contexts involving owned artifacts and other human agents with malicious intentions (e.g. whether a shopper knows she has a diamond). By contrast, Experiments 2-4 feature stories set in outdoor contexts involving natural kinds (e.g. whether an ecologist knows she saw a species of ground squirrel).

Experiment 2 (N=135, aged 18-59 years, mean age = 29.1 years, 41 female) replicates the effect on knowledge attributions of a salient but failed threat to the truth of a perceptual judgment (see Table 2). For example, and to give a sense of the materials in Experiments 2-4, the following vignette was used to further test the effect of failed threats on knowledge attribution:

Darrel is an ecologist collecting data on red speckled squirrels in Canyon Falls national park. The park is divided into ten zones and today Darrel is working in Zone 3. While scanning the river valley with his binoculars, Darrel sees a small, bushy-tailed creature with distinctive red markings on its chest and belly. The red speckled ground squirrel is the only native

species with such markings. Darrel records in his journal, "At least one red speckled ground squirrel in Zone 3 today." Ecologists are unaware that a non-native species of prairie dog recently began invading the park. These prairie dogs also have red markings on their chest and belly. When these prairie dogs tried to invade Zone 3, the red speckled ground squirrels were unable to completely drive them away. Still, the animal Darrel is looking at is indeed a red speckled ground squirrel.

We again found that when a threat *failed* to prevent an agent from detecting the truth, participants attributed knowledge at rates exceeding chance, $p < .001$, test proportion = .5, in this case that "Darrel knows that there is at least one red speckled ground squirrel in Zone 3 today." Indeed these rates were not significantly different from a case where no such threat was mentioned, Fisher's $p = .164$, n.s.. By contrast, participants overwhelmingly declined to attribute knowledge, $p < .001$, when a threat *succeeded* in preventing an agent from detecting the truth, for instance when the animal Darrel was looking at above was indeed one of the prairie dogs. Again the luck involved in such a threat failing is viewed as fully consistent with knowledge.

Table 2: Experiment 2: Percentage of participants attributing knowledge.

	No Threat	Threat	No Detection
Knows	81%	67%	16%

Experiment 3 (*N*=141, aged 18-66 years, mean age = 31.2 years, 57 female) replicates the effect on knowledge attributions of an unnoticed change in the explanation for why the agent's belief is true (See Table 3). The results support the claim that an unnoticed replacement-by-backup does affect whether it is viewed as knowledge, χ^2 (df = 2) = 40.16, $p < .001$, Cramer's V = .534. When replacement does not occur, knowledge attribution is at ceiling but when it does occur, knowledge attribution is significantly lower, Fisher's, $p = .017$, Cramer's V = .254. Nevertheless, such a replacement is widely viewed as consistent with knowing: most participants in Replacement attributed knowledge even though the belief was ultimately true because of the backup truth-maker, $p = .033$, test proportion = .5. However, participants' willingness to attribute knowledge when replacement occurred seemed to depend crucially on whether the initial truth-maker was detected, as demonstrated by the very low rates of knowledge attribution in No Detection, significantly below chance, $p < .001$.

Table 3: Experiment 3: Percentage of participants attributing knowledge.

	Normal Detection	Replacement	No Detection
Knows	88%	66%	23%

Experiment 4 (*N*=149, aged 19-63 years, mean age = 30 years, 42 female) replicates the effect of similarity or dissimilarity of the initial and replacement truth-makers on knowledge attributions. The results continue to support the view that when disruption and replacement occurs, knowledge attribution is affected by how similar or dissimilar the replacement truth-maker is to the original, χ^2 (df = 2) = 25.73, $p < .001$, Cramer's V = .416. People were more likely to attribute knowledge when the replacement was highly similar to the original than when it was dissimilar, Fisher's, $p = .014$, one-tailed, Cramer's V = .267, and when the replacement was dissimilar than when there was no detection at all, Fisher's, $p = .012$, Cramer's V = .260. As expected, No Detection was overwhelmingly viewed as a case of ignorance, with rates of knowledge attribution falling far below chance, binomial, $p < .001$, test proportion = .5. Knowledge attribution was also below chance in Dissimilar, binomial, $p = .009$. Knowledge attribution in Similar differed neither from chance, binomial, $p = .392$, nor from the rate of knowledge attribution in Replacement from Experiment 3, binomial, $p = .248$, test proportion = .66.

Table 4: Experiment 4: Percentage of participants attributing knowledge.

	Similar	Dissimilar	No Detection
Knows	57%	31%	10%

Though comparing results across these experiments is fraught, we still note the impressive consistency of knowledge attribution in structurally analogous conditions in Experiment 1 and replications of Experiments 2-4. All knowledge controls consistently resulted in approximately 80-90% knowledge attribution. Conditions involving failed threats were consistently treated similarly to clear cases of knowledge. Cases serving as ignorance controls consistently resulted in approximately 10-15% knowledge attribution. Cases of luck involving similar replacement consistently resulted in approximately 50-60% knowledge attribution. The similarities among these findings suggest that the ordinary concept of knowledge is deeply sensitive to the structural features of cases that we have identified.

General Discussion

Like any human achievement, knowledge is usually due to a mix of ability and luck. But the effect of luck on knowledge attributions is not well understood. We tested the effect of three luck-related factors on knowledge attribution — threat, disruption, and replacement — and found that the relation is complex and often surprising. We suggest that our investigation is profitably viewed as a case study in how

people judge the relationship between success and luck. We focused on a centrally important class of cognitive evaluations, knowledge judgments, which in everyday life are often implicated in further important questions about how people will behave, how they ought to behave, the extent to which we credit or blame them for outcomes, and whether we excuse or punish them for transgressions.

We made three main findings. First, we found that knowledge attributions are insensitive to the luck of a salient threat that fails to prevent someone from perceptually detecting the truth. Second, we found that knowledge attributions are sensitive to the luck involved with an unnoticed disruption and change in the explanation for why a belief is true. Third, we found that when the explanation changes for why a belief is true, knowledge attributions are sensitive to the way in which the truth is replaced or restored. People are more inclined to attribute knowledge when the "backup" or replacement truth-maker is similar to the original truth-maker than when it is dissimilar. Although we have treated the similarity/dissimilarity distinction as a dichotomy in these initial attempts to test its effect, we acknowledge that it is probably better thought of as a continuum. Further work is needed to understand how subtler changes in similarity affect knowledge attributions.

There has been recent controversy in psychology and philosophy over whether laypeople attribute lucky knowledge to agents in an intriguing range of cases known as "Gettier cases" (Gettier 1963; see Turri 2012 for a review). Studying Gettier cases is theoretically useful because they provide an excellent "opportunity to test which factors affect knowledge attributions" (Starmans & Friedman, 2013, p. 2). Some researchers have found that laypeople do tend to attribute knowledge in some Gettier cases (Starmans & Friedman, 2012), some researchers claim that laypeople do not (Nagel, San Juan, & Mar, 2013), and others observed mixed results depending on the method of questioning (Turri, 2013b). Our findings suggest an explanation for the seemingly inconsistent prior findings and theorizing on Gettier cases: knowledge attributions are sensitive to different forms and combinations of luck and prior research on Gettier cases has not controlled for all the sensitivities identified here. Indeed, by some criteria (e.g., Zagzebski, 1996; Lewis, 1996; Pritchard 2005), researchers would count the stories used for five separate conditions in Experiment 1 as Gettier cases: Failed Threat, Detect Similar, Detect Dissimilar, No Detect Similar, and No Detect Dissimilar. But if intuitions about Gettier cases vary as widely as our results indicate — from patterns that closely resemble responses to paradigmatic knowledge (Failed Threat) to patterns that closely resemble responses to paradigmatic ignorance (No Detect Dissimilar) — then "Gettier case" is a theoretically useless category. The fact that something is a Gettier case would be consistent with its being both overwhelmingly judged knowledge and overwhelmingly judged ignorance, thereby masking differences that radically affect the psychology of

knowledge attributions and depriving the category of any diagnostic or predictive value.

An important question for further research is why knowledge attributions are selectively sensitive to some but not other forms of luck. Answering this question might also help shed light on the social function of knowledge attributions as well as neighboring domains such as moral evaluation. Ethicists and moral psychologists have asked similar questions about the role of luck in evaluating an agent's moral status (Williams & Nagel 1976; Young et al. 2010). But whatever the ultimate explanation for this particular pattern of selective sensitivity, one thing is certain. We are fallible social beings deeply interested in keeping track of what people deserve and are responsible for. Whether it's a question of knowledge or morality or athletics, we seek criteria to distinguish genuine achievements from lucky outcomes. Our results have helped to expose some of those criteria.

References

Aristotle. (350 BCE/1941). Posterior analytics. In R. McKeon (Ed.), G. R. G. Mure (Trans.), *The basic works of Aristotle*. New York: Random House.

Beebe, J. R., & Buckwalter, W. (2010). The epistemic side-effect effect. *Mind & Language*, *25*(4), 1–25.

Beebe, J. R., & Shea, J. (2013). Gettierized Knobe effects. *Episteme*, *10*(03), 219–240.

Blouw, P., Buckwalter, W., & Turri, J. Knowledge, luck, and Gettier cases. In preparation.

Booth, J. R., Hall, W. S., Robison, G. C., & Kim, S. Y. (1997). Acquisition of the Mental State Verb Know by 2- to 5-Year-Old Children. *Journal of Psycholinguistic Research*, 26(6), 581–603.

Buckwalter, W. (2013). Gettier made ESEE. *Philosophical Psychology*, 1–16.

Buckwalter, W. (under review). Factive verbs and protagonist projection.

Buckwalter, W., & Turri, J. In press. Telling, showing and knowing: a unified account of pedagogical norms. *Analysis*.

Buckwalter, W., & Turri, J. Under review. Action, truth and knowledge.

Colaco, D., Buckwalter, W., Stich, S., & Machery, E. (under review). Epistemic Intuitions in Fake Barn Cases.

Cullen, S. (2010). Survey-Driven Romanticism. *Review of Philosophy and Psychology*, 1(2), 275–296. doi:10.1007/s13164-009-0016-1

Duff, R. A. (1990). *Intention, Agency and Criminal Liability*. Blackwell.

Engel, M., Jr. (1992). Is epistemic luck compatible with knoweldge? *The Southern Journal of Philosophy*, *30*(2), 59–75. doi:10.1111/j.2041-6962.1992.tb01715.x

Fantl, J., & McGrath, M. (2009). *Knowledge in an Uncertain World*. Oxford University Press.

Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis*, 23(6), 121–123.

Goldman, A. I. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, 73(20), 771–791.

Harman, G. (1973). *Thought*. Princeton, NJ: Princeton University Press.

Hart, H. (1959). Prolegomenon to the principles of punishment. *Proceedings of the Aristotelian society*, 60, 1–26.

Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.

Knobe, J. (2003). Intentional action in folk psychology: An experimental investigation. *Philosophical Psychology*, 16(2), 309–324. doi:10.1080/09515080307771

Malle, B. F., & Knobe, J. (1997). The folk concept of intentionality. *Journal of Experimental Social Psychology*, 101–121.

Malle, B. (2006). Intentionality, Morality, and Their Relationship in Human Judgment. *Journal of Cognition and Culture*, 6(1), 87–112.

Malle, B. F., & Nelson, S. E. (2003). Judging mens rea: the tension between folk concepts and legal concepts of intentionality. *Behavioral Sciences & the Law*, 21(5), 563–580.

Mele, A. R., & Cushman, F. (2007). Intentional Action, Folk Judgments, and Stories: Sorting Things Out. *Midwest Studies in Philosophy*, 31(1), 184–201.

Nagel, J., Juan, V. S., & Mar, R. A. (2013). Lay denial of knowledge for justified true beliefs. *Cognition*, 1–10. doi:10.1016/j.cognition.2013.02.008

Nagel, J., Mar, R., & Juan, V. S. (2013). Authentic Gettier cases: A reply to Starmans and Friedman. *Cognition*, 1–4. doi:10.1016/j.cognition.2013.08.016

Pizarro, D. A., Uhlmann, E., & Bloom, P. (2003). Causal deviance and the attribution of moral responsibility. *Journal of Experimental Social Psychology*, 39(6), 653–660.

Plato. (1997/380 BCE). The Republic. G. M. A. Grube & C. D. C. Reeve (Trans.). In J. M. Cooper (Ed.), *Plato: Complete Works*. Indianapolis: Hackett.

Pritchard, D. (2005). *Epistemic luck*. New York: Oxford University Press.

Reid, T. (1785/2002). *Essays on the intellectual powers of man*. (D. Brookes, Ed.) University Park: Pennsylvania State University Press.

Russell, B. (1948). *Human knowledge: Its scope and limits*. New York: Routledge.

Searle, J. (1983). *Intentionality*. Cambridge: Cambridge University Press.

Shope, R. K. (1983). *The Analysis of Knowing: A Decade of Research*. Princeton, New Jersey: Princeton University Press.

Shope, R. K. (2002). Conditions and analyses of knowing. In P. K. Moser (Ed.), *The Oxford handbook of epistemology*. Oxford: Oxford University Press.

Sosa, E. (2007). *A virtue epistemology: apt belief and reflective knowledge*, volume I. Oxford University Press.

Stanley, J. (2005). *Knowledge and practical interests*. Oxford University Press.

Starmans, C., & Friedman, O. (2012). The folk conception of knowledge. *Cognition*, 124(3), 272–83.

Starmans, C., & Friedman, O. (2013). Taking "know" for an answer: A reply to Nagel, San Juan, and Mar. *Cognition*, 1–4. doi:10.1016/j.cognition.2013.05.009

Turri, J. (2011a). The Express Knowledge Account of Assertion. *Australasian Journal of Philosophy*, 89(1).

Turri, J. (2011b). Manifest failure. *Philosophers' Imprint*, 11(8), 1–11.

Turri, J. (2012). In Gettier's wake. In S. Hetherington (Ed.), *Epistemology: The key thinkers*. Continuum.

Turri, J. (2013a). Knowledge and suberogatory assertion. *Philosophical Studies*. doi:10.1007/s11098-013-0112-z

Turri, J. (2013b). A conspicuous art: putting Gettier to the test. *Philosophers' imprint*.

Turri, J. (in press). *Epistemology: a guide*. Wiley-Blackwell.

Unger, P. (1968). An analysis of factual knowledge. *Journal of Philosophy*, 65(6), 157–170.

Unger, P. (1975). *Ignorance: a case for skepticism*. Oxford: Oxford University Press.

Weinberg, J. M., Nichols, S., & Stich, S. (2001). Normativity and epistemic intuitions. *Philosophical Topics*, 29(1-2), 429–460.

Williams, B., & Nagel, T. (1976). Moral luck. *Proceedings of the Aristotelian society*, 50, 115–151.

Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.

Wittgenstein, L. (1975). *On certainty*. (D. Paul & G. Anscombe, Trans., G. Anscombe & G. H. V. Wright, Eds.). Malden, MA: Blackwell.

Wright, J. C. (2010). On intuitional stability: the clear, the strong, and the paradigmatic. *Cognition*, 115(3), 491–503. doi:10.1016/j.cognition.2010.02.003

Young, L., Nichols, S., & Saxe, R. (2010). Investigating the Neural and Cognitive Basis of Moral Luck. *Review of Philosophy and Psychology*, 1, 333–349.

Zagzebski, L. T. (1996). *Virtues of the mind: an inquiry into the nature of virtue and the ethical foundations of knowledge*. Cambridge: Cambridge University Press.