Knowledge Central:

A Central Role for Knowledge Attributions in Social Evaluations

John Turri

Ori Friedman

Ashley Keefner

Word count: 6170

Keywords: knowledge; theory of mind; social evaluations; decision-making; outside/inside evidence

Abstract

Five experiments (N = 1710) demonstrate the central role of knowledge attributions in social evaluations. In Experiments 1-3, we manipulated whether an agent believes, is certain of, or knows a true proposition and asked people to rate whether the agent should perform a variety of actions. We found that knowledge, more so than belief or certainty, leads people to judge that the agent should act. In Experiments 4-5, we investigated whether attributions of knowledge or certainty can explain an important finding on how people act based on statistical evidence, known as "the Wells effect." We found that knowledge attributions, but not certainty attributions, mediate this effect on decision-making.

Knowledge Central:

A Central Role for Knowledge Attributions in Social Evaluations

As social beings, people are motivated to gather and retain information about others, including information about their mental states. People give particular attention to others' knowledge states. "Know" is the most frequently used mental state verb in the English language, more common than even "think" and "want," in both adults and children (http://www.oxforddictionaries.com/words/the-oec-facts-about-the-language; Shatz, Wellman & Silber, 1983). People have two broad methods for judging what others know. First, they base these judgments on detailed information about an agent's circumstances, such as information about which kinds of evidence the agent has accessed (e.g., Buckwalter, 2014; Pratt & Bryant, 1990; Sodian, 1988; Sodian & Wimmer, 1987; Starmans & Friedman, 2012; Turri, Buckwalter & Blouw, 2014; Turri & Friedman, 2014; Turri 2014; Wimmer, Hogrefe & Perner, 1988; Woolley & Wellman, 1993). Second, people also base these judgments on general expectations about what others are likely to know (e.g., Birch & Bloom, 2003; Brennan & Williams, 1995; Fussell & Krauss, 1991, 1992; Lau, Chiu & Hong, 2001; Nickerson, Baddeley & Freeman, 1987; Thomas & Jacoby, 2013). For instance, children expect generic information to be more broadly known than specific information (Cimpian & Scott, 2012).

What are the consequences of attributing knowledge to others? Previous research suggests that knowledge attributions affect several kinds of social evaluation, including moral judgments, decisions about whom to trust, and decisions about how to interpret others' behavior (e.g.,

Schroeder & Linder, 1976; Sobel, 2009; Young & Saxe, 2011; Yuill & Perner, 1988).[1] For instance, one series of studies compared the effects of belief, knowledge, and certainty on decisions about whom to trust (Furrow & Moore, 1990; Moore, Pure & Furrow, 1990). Participants were asked to find an object while one agent said, "I know it is in the red box," and another said, "I think/am sure it is in the blue box." Participants were more likely to look in the red box, suggesting that attributions of knowledge imply greater certainty than certain other kinds of self-attributions. In this paper, we examine another possible consequence of attributing knowledge. We examine whether these attributions play an important role in judgments about whether agents are licensed to act.

There are often limits on whether people should act on their beliefs. For instance, suppose that some friends are at the beach. The health department recently tested the beach water and declared it dangerously unsafe. Even if the friends think the water is safe, this might not license them to take down a sign stating that the water is hazardous, or to encourage other people to swim in it. Limits on whether actions are licensed are particularly important in legal and medical decision-making — they determine whether a juror may vote to convict a defendant he thinks is guilty, and whether a doctor should undertake a risky medical intervention that she thinks will help a patient.

One factor that might license action is certainty, or one's confidence that a particular

[1] Much of this research has compared knowledge to ignorance rather than to other representational mental states, such as belief or certainty. One concern with this approach is that the difference between knowledge and ignorance is stark and does not isolate knowledge per se. Someone with a true belief about a topic is not ignorant, yet not all true beliefs are viewed as knowledge (Starmans & Friedman, 2012). A further concern is that people make inferences about belief and certainty from knowledge judgments. People infer that someone who knows a proposition also believes it (Buckwalter, Rose & Turri, 2013), and people expect that someone who knows is extremely confident (Wesson & Pulford, 2009). If people are likely to infer belief or certainty from knowledge, then to isolate what knowledge attributions contribute to social evaluations, we should compare the effects of knowledge to those of certainty and belief.

conclusion is true. If the friends are extremely confident that the water is safe, or the juror is especially certain of the plaintiff's guilt, perhaps this licenses them to act. But confidence in a conclusion may not always suffice. For instance, Wells (1992) demonstrated that mock jurors make different liability judgments across conditions where they assigned equal probability to a defendant's guilt (also see Niedermeier, Kerr & Messé, 1999 and Wright, MacEachern, Stoffer & MacDonald, 1996). Instead, participants' willingness to convict the defendant depended on the nature of the evidence suggesting that the defendant was guilty.

We propose that knowledge attributions might have a privileged role in licensing actions. When an agent is viewed as *knowing* some fact, as opposed to believing or being certain of it, people might be more likely to view the agent as licensed to act. One reason to expect this is that knowledge requires an agent's conclusion to be based on evidence connecting the agent's representation to the world in some appropriate way (Friedman & Turri, 2015; Starmans & Friedman, 2012). This connection to evidence might differentiate knowledge from belief and certainty, because these latter states might occur with weaker or less direct evidence. Broadly consistent with this, people often deny that agents are knowledgeable about some conclusion, even when they are extremely confident that it is true. For instance, in one study, participants denied that a woman knew that her lottery ticket was a loser, even though she concluded this after learning that it only had a 1-in-10,000,000 odds of being a winner (Friedman & Turri, 2015).

We tested whether knowledge plays this role in two quite different ways. In Experiments 1-3, participants read vignettes in which an agent believes, is certain of, or knows a relevant true proposition. Then participants rated whether the agent should perform a variety of actions. We found that knowledge, more so than belief or certainty, leads people to judge that the agent should

act. In Experiments 4-5, we provided a more specific test of whether knowledge is privileged in licensing action. We investigated whether attributions of knowledge or certainty can explain the "Wells effect" (Niedermeier, Kerr & Messé, 1999; Wells 1992; Wright, MacEachern, Stoffer & MacDonald, 1996), in which different forms of statistical evidence produce qualitatively different decisions about how one should act on evidence (e.g., whether a defendant should be convicted), despite inducing equivalent subjective probability estimates of the relevant underlying facts. We found that certainty attributions do not mediate this effect, whereas knowledge attributions do. This suggests that the different forms of statistical evidence produce different judgments about whether one is licensed to act because they produce different knowledge judgments.

Experiment 1

Method

Participants. Five hundred ninety-eight adult U.S. residents were tested (aged 18-68 years, mean age = 31; 96% reporting English as a native language; 236 female). Participants were recruited and tested online using Amazon Mechanical Turk and Qualtrics and compensated $0.40 for approximately 2-3 minutes of their time. Repeat participation was prevented by screening Amazon Mechanical Turk Worker IDs, which are uniquely identifying in the population sampled. The same recruitment procedures were used in all the experiments reported here. Eighty participants failed a comprehension question and were excluded from the analysis, but including them does not change the pattern of results reported below.

Materials and Procedure. Participants were assigned to one of twelve conditions in a 4 (Cover Story: Beach, Bagel, Restaurant, Carnival) × 3 (Mental State: Think, Know, Certain)

between-subjects design. Each participant read a single story, responded to a set of test and comprehension questions, then filled out a brief demographic survey.

We tested four cover stories to ensure that any findings were not due to incidental features of any one scenario or subject matter. Otherwise, we had no expectations for the Cover Story factor. The Beach story concerned a woman who visits a beach, learns that the water was recently declared unsafe, examines the water herself, and concludes that it is actually safe. The Bagel story concerned a woman who shops for food, reads a warning that the bagels may contain nuts, examines the bagels, and concludes that they do not contain nuts. The Restaurant story concerned a woman who enters a restaurant, reads that the bread is trans-fat-free, but concludes that the bread contains trans-fat. Finally, the Carnival story concerns a man who visits a carnival, gets in line for the Ferris wheel, then concludes that the Ferris wheel is unsafe.

The Mental State factor varied how the agent's conclusion was described. In the Think conditions, the agent "thinks" that the conclusion is true; in the Know conditions, the agent "knows" that it is true; and in the Certain conditions, the agent "is certain" that it is true. In all stories, the agent's conclusion is in fact true. Also, in all stories, the agent came to the mental state after briefly examining the object or setting that the mental state concerned. It might seem odd that brief examination could allow an agent to know or be certain of facts such as whether water is safe to swim in, or whether bagels contain nuts. However, there is an advantage to specifying this: it reduces the possibility of participants making assumptions about the manner in which the agent examined the object (or assumptions about the information gained from this examination) based on the agent's resulting mental state. For instance, it reduces the possibility of participants imagining that the examination was longer or more detailed when the agent came to know,

compared with when the agent instead came to believe.

A supplemental file includes the complete text of all stimuli, but to give readers a sense of the materials, here is the Beach story (Mental State manipulations bracketed and separated by a slash):

> The water at Metro Beach was recently tested and declared unsafe for swimming. However, although it is difficult to tell, the health department botched the test and, as a matter of fact, the water is perfectly safe for swimming. It's a hot summer day and Alicia decides to go to Metro beach. She briefly examines the water, and now she [thinks/knows/is certain] that the water is safe for swimming.

After reading the story, participants rated their agreement with several test statements about whether the agent should perform certain actions, given that the agent had drawn the relevant conclusion. For example, after reading the Beach story, participants were instructed, "Given that Alicia [thinks/knows/is certain] that the water is safe for swimming, please rate your agreement with each of the following statements:"

1. Alicia should go swimming.
2. Alicia should allow her children to go swimming.
3. Alicia should encourage other people at the beach to go swimming.
4. Alicia should tell other people at the beach that the water is safe for swimming.
5. Alicia should remove the sign at the beach that says the water is unsafe for swimming.
6. Alicia should tell the health department that their test results were wrong.

The order of statements was randomized. Responses were collected on a standard 7-point Likert

scale, 1 strongly disagree – 7 strongly agree, left-to-right on the participant's screen.

After responding to the test statements, participants went to a new screen where they answered a comprehension question (response options rotated randomly) and rated how sure the agent was of the conclusion (e.g., "How sure is Alicia that the beach water is safe?") on a 7-point Likert scale, 1 not at all sure – 7 completely sure. Participants could not return to a previous screen to change their answers.

Results

Preliminary analysis revealed that neither participant age nor participant gender affected response to the test statements. The same is true for all the experiments reported here. We will not discuss these demographic variables further.

We will focus on *mean evaluation score*s. For each participant, we calculated a mean evaluation score by summing the value of their response to each test statements and then dividing by the total number of statements. (See the supplemental file for a table with statistics for each dependent measure.) Higher mean evaluation scores indicate stronger agreement that the agent should engage in the relevant activities.

An analysis of variance revealed that mean evaluation score was affected by Cover Story, $F(3, 506) = 96.26$, $p < .001$, $\eta_p^2 = .363$; Mental State, $F(2, 506) = 28.45$, $p < .001$, $\eta_p^2 = .101$; and their interaction, $F(6, 506) = 2.34$, $p = .031$, $\eta_p^2 = .027$; all tests two-tailed. (See Figure 1.) Mean evaluation score was higher for Know than for Think for each cover story. (See Table 1.) Mean evaluation score was higher for Know than for Certain in the Beach story and showed a strong trend in this direction in the Bagel story (i.e., $p = .051$), but it did not differ in the Carnival and Restaurant stories. For the Beach story, confidence ratings were no different for Know ($M = 5.64$,

SD = 1.43) and Certain (*M* = 5.75, *SD* = 1.62), independent samples t-tests, *t*(87) = -0.33, *p* = .745. Similarly, for the Bagel story, confidence ratings were no different for Know (*M* = 5.65, *SD* = 1.48) and Certain (*M* = 5.75, *SD* = 1.87), *t*(74) = -0.26, *p* = .796. Thus, the higher mean evaluation scores for knowledge cannot be due to a difference in perceived confidence.

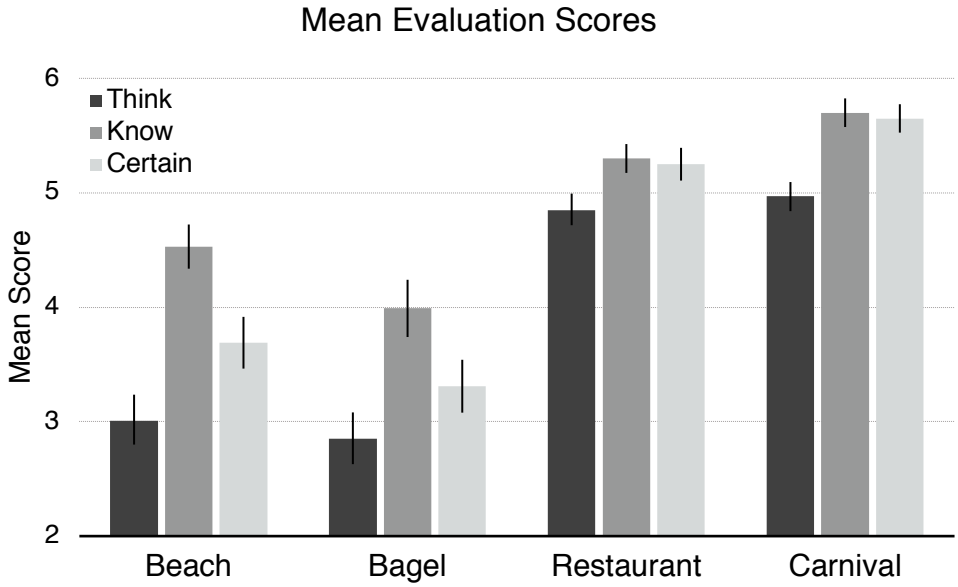

Figure 1. Experiment 1. Mean evaluation scores. Scales ran 1(SD) – 7 (SA). Error bars represent +/- one standard error of the mean.

Table 1. Experiment 1. Mean evaluation scores (MES). All tests 2-tailed.

Measure	Mean (SD)			Know vs. Think t-test			Know vs Certain t-test		
	Think	Know	Certain	t	df	p	t	df	p
Beach	n = 43	n = 45	n = 44						
MES	3.02 (1.38)	4.52 (1.30)	3.69 (1.50)	5.30	86	<.001	2.81	87	.006
Bagel	n = 28	n = 40	n = 36						
MES	2.85 (1.17)	3.99 (1.59)	3.31 (1.39)	3.21	66	.002	1.97	74	.051
Restaurant	n = 44	n = 48	n = 48						
MES	4.86 (0.88)	5.30 (0.87)	5.25 (0.99)	2.38	90	.019	0.22	94	.827
Carnival	n = 45	n = 49	n = 48						

MES	4.96 (0.87)	5.70 (0.87)	5.65 (0.86)	4.06	92	<.001	0.26	95	.799

Discussion

The results support the hypothesis that knowledge attributions play an important role in social evaluations. People consistently judged that knowledge licenses action more readily than true belief. The comparison between knowledge and certainty was more mixed. For two of the scenarios we tested, Beach and Bagel, people judged that knowledge licensed action more readily than certainty. But for the other two scenarios, Restaurant and Carnival, people judged knowledge and certainty to be equally good for licensing action. One explanation for this difference lies in the type of risk faced by the agent.

In the Beach and Bagel stories, greater risk resides in acting on the conclusion. For example, when Alicia convinces others that the beach water is safe, she faces the risk that she is wrong, and that heeding her advice will lead others to be sick; but if she keeps quiet about her conclusion, the worst possible outcome is that people will lose the opportunity to go in the water. By contrast, in the Restaurant and Carnival stories, greater risk resides in not acting on the conclusion. For instance, if the protagonist in Carnival does not act on his concern about the safety of the Ferris wheel, others could be seriously injured or killed; but if he acts and warns others, he only risks inconveniencing people while they wait for a mechanic to double-check the ride. One possibility, then, is that knowledge and certainty are equally good for licensing action when the risk resides in not acting while being right, whereas knowledge is superior when the risk resides in acting and being wrong. The next experiment tests this possibility.

Experiment 2

Method

Participants. Four hundred seven new participants were tested (aged 18-74 years, mean age = 32 years; 95% reporting English as a native language; 157 female). Fifty-seven participants failed a comprehension question and were excluded from the analysis, but including them does not change the pattern of results reported below.

Materials and Procedure. Participants were assigned to one of eight conditions in a 2 (Cover Story: Beach, Bagel) × 2 (Mental State: Know, Certain) × 2 (Risk: Not Acting, Acting) between-subjects design. The testing procedures were the same as in Experiment 1. The Mental State factor varied knowledge and certainty the same way as in Experiment 1. The basic Beach story was the same as in Experiment 1; the Bagel story was slightly modified to make the Risk manipulation more natural. The Risk factor varied whether the serious risk lay in not acting while being right, or acting and being wrong. For instance, in the Beach story for the Not Acting condition, the beach water was officially declared safe but the agent concludes that it is unsafe. By contrast, in the story for the Acting condition, the beach water is officially declared unsafe but the agent concludes that it is safe. In the Bagel story for the Not Acting condition, the bagels are labeled as nut-free but the agent concludes that they contain nuts. By contrast, in the story for the Acting condition, the bagels are labeled as containing nuts but the agent concludes that they are nut-free.

Results

Table 2 summarizes the findings. In the main text, we will again focus on mean evaluation scores

(see the Results section of Experiment 1 for an explanation of this). (See the supplemental file for a table with statistics for each dependent measure.) A two-way analysis of variance revealed that mean evaluation score was affected by Mental State $F(1, 344) = 8.04$, $p = .005$, $\eta_p^2 = .023$, Risk, $F(1, 344) = 264.99$, $p < .001$, $\eta_p^2 = .435$, and an interaction between Mental State and Risk, $F(1, 344) = 13.44$, $p < .001$, $\eta_p^2 = .038$. For each cover story, when the risk lay in not acting, mean evaluation scores for knowledge and certainty did not differ, but when the risk lay in acting, mean evaluation scores were higher for knowledge than for certainty. (See Figure 2, Table 2.) As in Experiment 1, higher mean evaluation scores for knowledge cannot be attributed to greater perceived confidence. For the Beach story, when the risk lay in acting, confidence ratings were no different for Know ($M = 5.67$, $SD = 1.27$) and Certain ($M = 5.73$, $SD = 1.66$), independent samples t-test, $t(84) = -0.16$, $p = .872$. Similarly, for the Bagel story, when the risk lay in not acting, confidence ratings were no different for Know ($M = 6.06$, $SD = 1.00$) and Certain ($M = 6.21$, $SD = 1.30$), $t(70) = -0.52$, $p = .604$.

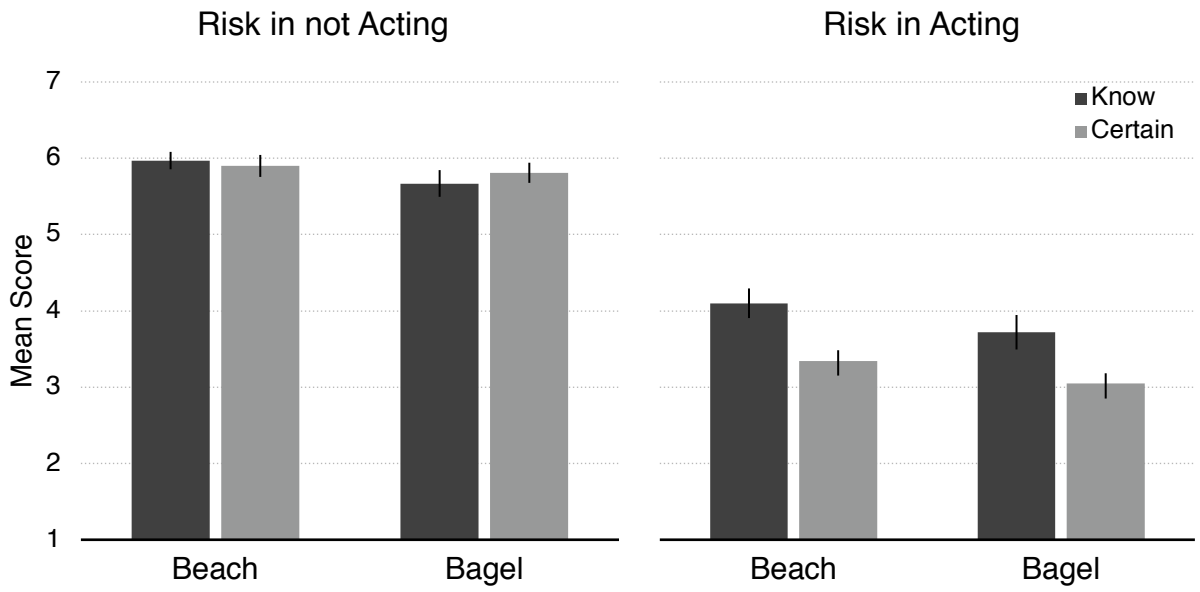

Figure 2. Experiment 2. Mean evaluation scores. Error bars represent +/- one standard error of the mean.

Table 2. Experiment 2. Mean evaluation scores (MES). All tests 2-tailed.

	Not Acting					Acting				
	Mean (SD)					Mean (SD)				
Measure	Know	Certain	T	df	p	Know	Certain	t	df	p
Beach	n = 49	n = 48				n = 46	n = 40			
MES	5.95 (0.81)	6.02 (0.94)	-0.37	95	.715	4.16 (1.36)	3.52 (1.30)	2.22	84	.029
Bagel	n = 45	n = 52				n = 33	n = 39			
MES	5.67 (1.13)	5.81 (0.91)	-0.69	95	.489	4.29 (1.51)	3.30 (1.42)	2.88	70	.005

Discussion

In light of the results from Experiment 1, we examined whether the difference between knowledge and certainty for licensing action occurs when the risk resides in acting but being wrong, but not when risk resides in not acting while being right. The results from the present experiment support

this hypothesis. Evaluations were affected by a significant interaction between the agent's mental state and the type of risk faced by the agent. More specifically, evaluation scores were equally high for knowledge and certainty when the risk lay in not acting while being right, but they were higher for knowledge when the risk lay in acting but being wrong.

One possible explanation of these findings is that they might reflect a ceiling effect. As can be seen in Figure 2 and Table 2, when the risk lay in not acting (while being right), mean evaluation scores were very high and close to the maximum possible score of 7. Hence, any differences between knowledge and certainty might have been masked by participants giving the highest ratings they were willing to give for the particular evaluations requested. Perhaps differences would be revealed for these items if we used a more sensitive scale.

One possible concern with some of the materials tested up until now is that it might seem odd to attribute knowledge or certainty to an agent based on, say, a brief examination of the water. One might wonder, then, whether we would observe a difference between knowledge and certainty when both mental states are based on good evidence, such as testing the water with a portable kit.[2] The next experiment addresses this concern.

Experiment 3

Method

Participants. Ninety-nine new participants were tested (aged 18-69 years, mean age = 34 years; 97% reporting English as a native language; 49 female). Data was not collected from one recruit

[2] Thanks to an anonymous referee for suggesting this variation of the case.

who did not sign the consent form.

Materials and Procedure. Participants were assigned to one of two conditions (Know, Certain) in a between-subjects design. The materials and procedures were similar to the beach stories from Experiment 2 when the risk resides in acting but being wrong. The main difference was that instead of ending with the sentence, "She briefly examines the water, and now she [knows/is certain] that the water is safe for swimming," the text ended with, "She tests the water with a portable kit. Based on the results, she [knows/is certain] that the water is safe."

Results

Independent samples t-tests showed that mean evaluation score was significantly higher in the Know condition ($M = 4.47$, $SD = 1.13$) than in the Certain condition ($M = 3.74$, $SD = 1.25$). Moreover, for each individual evaluative item, mean agreement was either significantly higher or trending higher in the Know condition. (See the supplemental file for a table with statistics for each dependent measure.)

Discussion

This experiment tested whether knowledge would license action more than certainty when the agent's mental state was based on stronger evidence. We found that knowledge did license action more than certainty when this occurred. Indeed, although we did not directly test this possibility, the results suggest that making the agent's evidence stronger might actually lead to an ever greater advantage for knowledge compared to certainty.

Experiment 4

The findings from Experiments 1-3 suggest that knowledge licenses action more than true belief and certainty. In this experiment, we took a different approach to investigating the role of knowledge in licensing action. We examined whether attributions of knowledge explain the Wells effect (Niedermeier, Kerr & Messé, 1999; Wells 1992; Wright, MacEachern, Stoffer & MacDonald, 1996), in which people's judgments about how to act on evidence (e.g., whether to convict a defendant) depend on the nature of the evidence. Specifically, we examined whether attributions of knowledge mediate this effect.

Wells (1992) demonstrated that mock jurors make different liability judgments across conditions where they assign equal probability to the defendant's guilt. Suppose that Smith's dog was hit by a bus but no one witnessed the accident. In one version of the case, jurors learn that 80% of the buses operating in the town belonged to the Blue Bus Company, and 20% belonged to the Grey Bus Company. People estimate that it is 80% likely that a Blue bus killed the dog, but they tend to *disagree* that the jury should find the Blue Bus Company liable. In another version of the case, jurors learn that shortly before the accident, a weigh station attendant made a log entry on the bus that eventually killed the dog. The log says, "Blue bus." It is known that 80% of "Blue bus" entries in the log correctly identify a Blue bus, and 20% incorrectly identify a Grey bus. Again people estimate that it is 80% likely that a Blue bus killed the dog, and they tend to *agree* that the jury should find the Blue Bus Company liable.

Related to this finding is the "gatecrasher paradox" (Cohen, 1981): if 95% of attendees snuck into the rodeo without paying, then why cannot the rodeo organizers sue a random attendee for non-payment? More generally related is the preference for "clinical" to "statistical" decision

procedures. In all of these examples, people are more willing to base decisions and actions on "observations" specific to the case at hand rather than empirically established statistics, even when decades of scientific evidence show that relying on the statistics produces better results (Dawes, Faust & Meehl, 1989, p. 1673; see also Dawes, 1996; Meehl, 1954).

One way to unify these findings draws on the distinction between "inside" and "outside" probability judgments (Lagnado & Sloman, 2004). "Inside" probabilistic information is specific to the case at hand (Kahneman & Tversky, 1982). For example, the weigh station attendant entered information about the very bus that hit the dog. "Outside" probabilistic information is generic and concerns distributions of properties or patterns. For example, a certain percentage of buses operating in town that day belonged to the Blue Bus Company. The attendant's log book and the distribution of buses each makes it likely that a Blue bus hit the dog, but only the former does so from the "inside." People are less likely to judge that the company is liable based on "outside" evidence.

Interestingly, the outside/inside difference also affects knowledge judgments (Friedman & Turri, 2015). People are more likely to attribute knowledge when an agent has inside evidence (i.e., evidence specific to the case at hand), compared with when the agent has outside evidence (i.e., evidence concerning a larger set of cases). This suggests an intriguing possibility consistent with the findings from Experiments 1-3: the outside/inside difference might affect judgments about how people are licensed to act *because* it affects judgments about what they know. That is, knowledge judgments might mediate the outside/inside effect on judgments about how people should behave. The present experiment examines this question.

Method

Participants. Three hundred two new participants were tested (aged 18-69 years, mean age = 32 years; 93% reporting English as a native language; 114 female).

Materials and Procedure. Participants were assigned to one of six conditions in a 3 (Story: Cab, Tumor, Rodeo) × 2 (Probability: Outside, Inside) between-subjects design. We had no expectations for the Story factor and included it to cover multiple scenarios familiar from the literature on outside/inside judgments. The Cab story was a simplified adaptation of Wells's original example about a lawsuit involving a motor vehicle destroying someone's belongings. The Tumor story is an example of an important medical decision made in a clinical context. The Rodeo story is a simplified version of the "gatecrasher paradox" story. The Probability factor varied whether the probability pertained to the case at hand (Inside) or generic distributional information (Outside).

After reading the story, participants rated their agreement with a knowledge attribution and a statement about what the agent should decide to do. The order of the two statements was rotated randomly and responses were collected on a standard 7-point Likert scale, 1 strongly disagree – 7 strongly agree, left-to-right on the participant's screen. Participants then went to a new screen and rated the probability of the relevant proposition, expressed as a percentage between one and one-hundred percent. The supplemental file contains all the stories and questions, but to give readers a sense of the materials, here are the story and questions for the Cab condition (Probability manipulation bracketed and separated by a slash):

> (Cab) Gary is suing the Blue Cab Company. Gary's prize-winning rose garden was destroyed by a taxi cab that drove on to his front lawn. During the trial, jurors learned

that only two cab companies operate in the town: the Blue Cab Company and the Green Cab Company. According to a computerized analysis of the video footage from when Gary's garden was destroyed, [80% of cabs on the road were Blue Cabs/the cab was 80% likely to be a Blue Cab].

1. The jurors should rule against the Blue Cab Company and make them pay damages to Gary.
2. The jurors know that a Blue Cab destroyed Gary's garden.
3. How probable is it that a Blue Cab destroyed Gary's garden?

Results

Table 3 summarizes the findings. For each cover story, knowledge judgments and "should" judgments were higher for inside probability than for outside probability, but estimations of the relevant proposition's probability did not differ. (See Table 3.) To test whether knowledge judgments mediated the effect of Probability on "should" judgments, for each Story we conducted a bootstrap mediation analysis (Hayes, 2013) with Probability as the independent variable, response to the "should" statement as the outcome variable, and knowledge judgments as potential mediator. In each case, knowledge judgments completely mediated the effect of Probability on the outcome. (See Figure 3.) For Cab, the 95% confidence interval for the direct effect was -0.26 to 0.78, and for the indirect effect it was 0.24 to 1.21. For Tumor, the 95% confidence interval for the direct effect was -0.22 to 0.72, and for the indirect effect it was 0.08 to 0.69. For Rodeo, the 95% confidence interval for the direct effect was -0.15 to 1.13, and for the indirect effect it was 0.09 to 0.98.

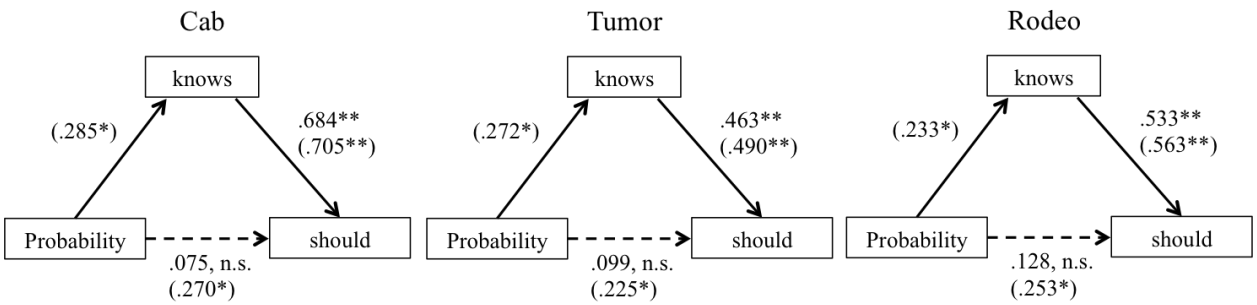

Figure 3. Experiment 4. Mediation results for the three stories (Cab, Tumor and Rodeo). Parenthetical values represent the strength of a simple regression between the two variables; values outside the parentheses represent the strength of the relationships in a model used to test for mediation. *p < .025, **p < .001.

Table 3. Experiment 4. Summary of findings. All tests 2-tailed.

	Mean (SD)				
Measure	Outside	Inside	t	df	p
Cab	n = 51	n = 49			
know	2.96 (1.91)	4.02 (1.68)	-2.95	98	.004
should	3.61 (1.73)	4.55 (1.66)	-2.78	98	.007
probable	78.02 (11.14)	76.84 (13.76)	0.47	98	.638
Tumor	n = 50	n = 50			
know	3.20 (1.95)	4.18 (1.54)	-2.79	92.78	.006
should	5.10 (1.39)	5.67 (1.11)	-2.27	97	.025
probable	78.70 (9.84)	75.71 (15.00)	1.17	97	.243
Rodeo	n = 51	n = 52			
know	4.29 (1.93)	5.13 (1.93)	-2.41	101	.018
should	3.78 (1.95)	4.75 (1.78)	-2.62	101	.010
probable	88.02 (17.93)	89.06 (17.35)	-0.30	101	.766

Discussion

Prior work has shown that the difference between inside and outside probability affects how people evaluate decisions and actions. This experiment tested the hypothesis that the outside/inside difference affects decision evaluations because it affects knowledge judgments. The results support the hypothesis. Across three very different scenarios, knowledge judgments completely mediated the outside/inside effect on evaluations of decisions and actions. The next experiment tests whether certainty attributions also mediate this effect.

Before proceeding, we will address a question about how to interpret the mediation result from the present experiment. Our analyses assumed a "knowledge model" whereby assignment to condition affects knowledge attributions, which in turn affect evaluative judgments. But there is an alternative "evaluative model" whereby assignment to condition affects evaluative judgments, which in turn affect knowledge judgments, perhaps as a post-hoc rationalization for the evaluative judgment.[3] It turns out that both mediation models are statistically significant, but overall the ratio of indirect to total effect is larger for the knowledge model (.526 [.289, .939], κ^2 = .128) than for the evaluative model (.438 [.244, .762], κ^2 = .117). (Preacher & Kelly, 2011 introduce the κ^2 ["kappa-squared"] effect size for mediation models.) This suggests that the knowledge model better fits the current data. However, the ratios are still sufficiently similar that we are not willing to

[3] We thank Associate Editor James Hampton for proposing this alternative.

confidently reject the evaluative model. Instead, we note that if the evaluative model turns out to be correct, it would still support the conclusion that knowledge attributions play an important role in social evaluations. For if the evaluative model is correct, then a natural explanation for this is that people are implicitly sensitive to the role that knowledge plays in licensing action, which motivates them to alter their knowledge attributions as a post-hoc justification.

Experiment 5

Method

Participants. Three hundred four new participants were tested (aged 18-69 years, mean age = 33 years; 93% reporting English as a native language; 129 female).

Materials and Procedure. The materials and procedures were the same as in Experiment 4, except that instead of judging whether the agent "knows" the relevant proposition, participants judged whether the agent "is certain."

Results

Table 4 summarizes the findings. For each cover story, "should" judgments were higher for inside probability than for outside probability; for Car and Rodeo, this difference passed the conventional threshold for statistical significance ($p < .05$), whereas for Tumor the difference was trending in the predicted direction ($p = .059$).[4] By contrast, neither certainty attributions nor estimations of the relevant proposition's probability differed significantly between inside and outside probability. To

[4] In keeping with all the other statistical reporting in this paper, this is the result of a two-tailed test, even though, in light of prior findings, we are entitled to a directional prediction. For a one-tailed test, the difference passes the conventional threshold for statistical significance.

test whether certainty judgments mediated the effect of Probability on "should" judgments, for each Story we conducted a bootstrap mediation analysis (Hayes, 2013) with Probability as the independent variable, response to the "should" statement as the outcome variable, and certainty judgment as potential mediator. In each case, certainty judgments failed to mediate the effect of Probability on the outcome. For Cab, the 95% confidence interval for the direct effect was 0.08 to 0.88, and for the indirect effect it was -0.27 to 0.78. For Tumor, the 95% confidence interval for the direct effect was -0.04 to 0.82, and for the indirect effect it was -0.07 to 0.20. For Rodeo, the 95% confidence interval for the direct effect was 0.43 to 1.53, and for the indirect effect it was -0.14 to 0.69.

Table 4. Experiment 5. Summary of findings. All tests 2-tailed.

	Mean (SD)				
Measure	Outside	Inside	t	df	p
Cab	n = 51	n = 50			
certain	3.37 (1.92)	3.72 (1.68)	-0.97	99	.335
should	3.82 (1.72)	4.56 (1.64)	-2.20	99	.030
probable	74.90 (14.23)	75.76 (14.51)	-0.30	99	.765
Tumor	n = 50	n = 52			
certain	3.58 (1.79)	3.81 (1.73)	-0.66	100	.514
should	5.34 (1.04)	5.77 (1.22)	-1.91	100	.059
probable	77.66 (9.77)	80.19 (3.70)	-1.71	62.31	.091
Rodeo	n = 51	n = 50			
certain	4.20 (1.76)	4.60 (1.68)	-1.18	99	.240
should	3.69 (1.91)	4.90 (1.45)	-3.60	93.03	<.001
probable	84.61 (21.97)	88.12 (17.96)	-0.88	99	.382

Discussion

Whereas the results from Experiment 4 provided evidence that knowledge judgments mediate the outside/inside effect on evaluations of decisions and actions, the results from this experiment failed to find evidence that certainty judgments mediate this effect. Taken together, the results from these two experiments support the conclusion that this effect occurs because of knowledge judgments but not because of certainty judgments.

General Discussion

The findings of five experiments show that knowledge attributions play an important role in judging that agents are licensed to act. In Experiments 1-3, we found that knowledge licenses action more than certainty or belief. Experiment 2 also found that differences between knowledge and certainty can depend on the type of risk faced by the agent. In situations where action carried greater risks than inaction, participants were more likely to judge that knowledgeable agents should act compared with agents who were certain; however, when the risks of inaction were greater than the risks of action, judgments for knowledge and certainty did not differ.

In Experiments 4-5, we found that knowledge attributions help to explain judgments about whether people should make decisions based on probabilistic evidence. Previous work showed that people are reluctant to endorse legal and medical decisions based on probabilistic evidence when this requires taking an "outside view" (i.e., judging a particular instance from information about a distribution) rather than an "inside view" (i.e., considering probabilistic evidence specifically concerning the case at hand). For instance, based on distributional information, people

are unwilling to say that a cab company is liable for an accident. Similarly, clinicians are often unwilling to make decisions or recommendations based on distributional information. In Experiment 4, we found that people's knowledge attributions entirely mediated the outside/inside effect on their evaluation of legal and medical decisions, whereas in Experiment 5, we found that people's certainty attributions did not mediate this effect. Together these findings demonstrate a central role for knowledge attributions in social evaluations.

These findings are interesting given that concepts like belief, certainty, and knowledge are closely connected, especially when they concern the same true proposition. People typically judge that knowledge requires true belief (Buckwalter, Rose & Turri, 2013; Starmans & Friedman, 2012), and previous findings and our present results suggest that people are often willing to infer that someone is certain of a proposition when they know it (Wesson & Pulford, 2009; Experiment 1 above). Close relations between these concepts may explain why research on "theory of mind" has been unclear about differences between how people consider knowledge and true belief in predicting the behavior of agents. For instance, attributing either mental state can be used to explain an agent's correct search for an object (i.e., she knew where it was, or she had a true belief about this). The present studies therefore advance understanding by showing that people do view these mental states as having different consequences for behavior. Whereas the present experiments demonstrated this by asking about which actions agents should perform based on their mental states, future research could examine whether these differences extend to predictions and explanations of actions.

The findings also suggest that knowledge attributions can play a pivotal role in legal and medical decision-making. For instance, in considering whether someone should be held liable for

harm, jurors might first ask themselves whether they know that the person is guilty. Before endorsing the decision made by another practitioner, a doctor might ask whether the practitioner knows that surgery is required. And before granting parole, a parole board might ask whether the clinician knows that the convict is rehabilitated. We emphasize that these are descriptive hypotheses about the role of knowledge attributions in legal and medical decision-making, not necessarily normative claims about how such decisions should be made.

Finally, we end with a hypothesis about why knowledge might be sensitive to the outside/inside difference. According to some recent philosophical theories, the ordinary concept of knowledge can be roughly defined as a justified true belief (e.g. Chisholm, 1989). However, this "justified true belief" approach has been heavily criticized and there is now a broad consensus that it fails to capture the ordinary concept (for reviews, see Ichikawa & Steup, 2012; Turri, 2012). In response to these developments, some theorists have proposed a radically different approach to the ordinary concept of knowledge, according to which knowledge is a representation of a fact caused or explained by the fact itself (e.g. Goldman, 1967; Turri, 2015). On this approach, knowledge involves, first and foremost, the detection of truths, either directly through perception or indirectly through competent inference from other known facts. That is, knowledge essentially involves the operation of cognitive mechanisms to detect and retain information. It need not always be based on or constituted by belief or justification, as those categories are ordinarily understood.

The results from the present studies suggest a possible point of convergence with these developments in the theory of knowledge. People's thinking about causation is sensitive to the outside/inside difference (Sloman & Lagnado, 2015). When making causal judgments, people do

not merely generate representations of variables linked by conditional probabilities. Instead, they also generate representations of the structure of relevant events and the mechanisms relating them to other events. This broadly corresponds to the difference between outside and inside information, as discussed above. If the concept of knowledge does essentially involve causation or mechanism — as the theories described in the previous paragraph suggest — then the outside/inside effect on knowledge judgments might be a special instance of the outside/inside effect on causal judgments more generally. That is, when we make knowledge attributions, we do not merely represent the person as having certain beliefs or evidence. Instead, we represent the person as being in an informational state with a certain causal or explanatory etiology.

Acknowledgments — For helpful conversation and feedback, we thank Wesley Buckwalter, Joshua Knobe, David Lagnado, and Angelo Turri. This research was supported by the Social Sciences and Humanities Research Council of Canada, the Natural Sciences and Engineering Research Council of Canada, and the Ontario Ministry of Economic Development and Innovation.

References

Birch, S. A., & Bloom, P. (2003). Children are cursed: An asymmetric bias in mental-state attribution. *Psychological Science*, *14*(3), 283-286. doi: 10.1111/1467-9280.03436

Brennan, S. E., & Williams, M. (1995). The feeling of another′ s knowing: Prosody and filled pauses as cues to listeners about the metacognitive states of speakers. *Journal of Memory and Language*, *34*(3), 383-398. doi: 10.1006/jmla.1995.1017

Buckwalter, W. (2014). Factive verbs and protagonist projection. *Episteme*, *11*(4), 391-409. doi: 10.1017/epi.2014.22

Buckwalter, W., Rose, D., & Turri, J. (2013). Belief through thick and thin. *Nous*, 1–28. doi:10.1111/nous.12048

Chisholm, R. (1989). *Theory of knowledge* (3rd ed.). Englewood Cliffs, NJ: Prentice Hall.

Cimpian, A., & Scott, R. M. (2012). Children expect generic knowledge to be widely shared. *Cognition*, 123(3), 419-433. doi: 0.1016/j.cognition.2012.02.003

Danovitch, J. H., & Keil, F. C. (2004). Should you ask a fisherman or a biologist? Developmental shifts in ways of clustering knowledge. *Child Development*, *75*(3), 918–931. doi:10.1111/j.1467-8624.2004.00714.x

David M. Sobel. (2009). Enabling conditions and children's understanding of pretense. *Cognition*, *113*(2), 177–188. doi:10.1016/j.cognition.2009.08.002

Dawes, R. M. (1996). *House of cards: Psychology and psychotherapy built on myth*. New York: Free Press.

Dawes, R., Faust, D., & Meehl, P. (1989). Clinical versus actuarial judgment. *Science*, *243*(4899), 1668–1674. doi:10.1126/science.2648573

Friedman, O., & Turri, J. (2015). Is probabilistic evidence a source of knowledge? *Cognitive Science*, *39*, 1062-1080. doi: 10.1111/cogs.12182

Fussell, S. R., & Krauss, R. M. (1991). Accuracy and bias in estimates of others' knowledge. *European Journal of Social Psychology*, *21*(5), 445-454. doi: 10.1002/ejsp.2420210507

Fussell, S. R., & Krauss, R. M. (1992). Coordination of knowledge in communication: Effects of speakers' assumptions about what others know. *Journal of personality and Social Psychology*, *62*(3), 378-391. doi: 10.1037/0022-3514.62.3.378

Goldman, A. I. (1967). A causal theory of knowing. *Journal of Philosophy*, *64*(12), 357–372. doi: 10.2307/2024268

Guglielmo, S., Monroe, A. E., & Malle, B. F. (2009). At the heart of morality lies folk psychology. *Inquiry*, *52*(5), 449–466. doi:10.1080/00201740903302600

Hayes, A. F. (2013). *Introduction to mediation, moderation, and conditional process analysis: A regression-based approach*. New York: Guilford Press.

Ichikawa, J. J., & Steup, M. (2012). The analysis of knowledge. *Stanford Encyclopedia of Philosophy* (Ed. E.N. Zalta). Retrieved August 1, 2015, from http://plato.stanford.edu/entries/knowledge-analysis/index.html

Kahneman, D., & Tversky, A. (1972). Subjective probability: a judgment of representativeness. *Cognitive Psychology*, *3*(3), 430–454. doi: 10.1016/0010-0285(72)90016-3

Kahneman, D., & Tversky, A. (1982). Variants of uncertainty. *Cognition*, *11*(2), 143–157. doi:10.1016/0010-0277(82)90023-3

Lagnado, D. A., & Channon, S. (2008). Judgments of cause and blame: The effects of

intentionality and foreseeability. *Cognition*, *108*(3), 754–770. doi:10.1016/j.cognition.2008.06.009

Lagnado, D., & Sloman, S. A. (2004). Inside and outside probability judgments. In D. J. Koehler & N. Harvey (Eds.), *Blackwell handbook of judgment & decision making* (pp. 157–176). Malden, Mass: Blackwell.

Lau, I. Y. M., Chiu, C. Y., & Hong, Y. Y. (2001). I know what you know: Assumptions about others' knowledge and their effects on message construction. *Social Cognition*, *19*(6), 587-600. doi: 10.1521/soco.19.6.587.20888

Lieberman, M. D. (2013). *Social: why our brains are wired to connect*. New York: Crown Publishers.

Lindskold, S., & Han, G. (1986). Intent and the judgment of lies. *The Journal of Social Psychology*, *126*(1), 129–130. doi:10.1080/00224545.1986.9713581

Malle, B. F., & Knobe, J. (1997). The folk concept of intentionality. *Journal of Experimental Social Psychology*, *33(3)*, 101–121. doi: 10.1006/jesp.1996.1314

Mars, R. B., Neubert, F.-X., Noonan, M. P., Sallet, J., Toni, I., & Rushworth, M. F. S. (2012). On the relationship between the “default mode network” and the ‘social brain’. *Frontiers in Human Neuroscience*, *6*. doi:10.3389/fnhum.2012.00189

Meehl, P. E. (1954). *Clinical vs. statistical prediction: a theoretical analysis and a review of the evidence*. Minneapolis, Minn.: University of Minnesota Press.

Moore, C., & Davidge, J. (1989). The development of mental terms: pragmatics or semantics? *Journal of Child Language*, *16*(3), 633–641. doi:10.1017/S030500090001076X

Moore, C., Pure, K., & Furrow, D. (1990). Children's understanding of the modal expression of

speaker certainty and uncertainty and its relation to the development of a representational theory of mind. *Child Development*, *61*(3), 722–730. doi:10.1111/j.1467-8624.1990.tb02815.x

Nickerson, R. S., Baddeley, A., & Freeman, B. (1987). Are people's estimates of what other people know influenced by what they themselves know? *Acta Psychologica*, *64*(3), 245-259. doi: 10.1016/0001-6918(87)90010-2

Niedermeier, K. E., Kerr, N. L., & Messé, L. A. (1999). Jurors' use of naked statistical evidence: Exploring bases and implications of the Wells effect. *Journal of Personality and Social Psychology*, *76*(4), 533. doi:10.1037/0022-3514.76.4.533

Pratt, C., & Bryant, P. (1990). Young children understand that looking leads to knowing (so long as they are looking into a single barrel). *Child Development*, *61*(4), 973-982. doi: 10.1111/j.1467-8624.1990.tb02835.x

Preacher, K. J., & Kelley, K. (2011). Effect size measures for mediation models: Quantitative strategies for communicating indirect effects. *Psychological Methods*, *16*(2), 93–115. http://doi.org/10.1037/a0022658

Rilling, J. K., Barks, S. K., Parr, L. A., Preuss, T. M., Faber, T. L., Pagnoni, G., et al. (2007). A comparison of resting-state brain activity in humans and chimpanzees. *Proceedings of the National Academic of Sciences*, *104(43)*, 17146-17151. doi: 10.1073/pnas.0705132104

Schilbach, L., Eickhoff, S. B., Rotarska-Jagiela, A., Fink, G. R., & Vogeley, K. (2008). Minds at rest? Social cognition as the default mode of cognizing and its putative relationship to the "default system" of the brain. *Consciousness and Cognition*, *17*(2), 457–467. doi:10.1016/j.concog.2008.03.013

Schroeder, D. A., & Linder, D. E. (1976). Effects of actor's causal role, outcome severity, and knowledge of prior accidents upon attributions of responsibility. *Journal of Experimental Social Psychology*, *12*(4), 340–356. doi: 10.1016/S0022-1031(76)80003-0

Shatz, M., Wellman, H. M., & Silber, S. (1983). The acquisition of mental verbs: A systematic investigation of the first reference to mental state. *Cognition*, *14*(3), 301–321. doi:10.1016/0010-0277(83)90008-2

Shaver, K. G. (1985). *That attribution of blame: causality, responsibility, and blameworthiness*. New York: Springer-Verlag.

Sloman, S. A., & Lagnado, D. (2014). Causality in thought. *Annual Review of Psychology*. doi:10.1146/annurev-psych-010814-015135

Sodian, B. (1988). Children's attributions of knowledge to the listener in a referential communication task. *Child Development, 59*(2)*,* 378-385. doi: 10.2307/1130317

Sodian, B., & Wimmer, H. (1987). Children's understanding of inference as a source of knowledge. *Child Development*, *58*(2), 424–433. doi: 10.2307/1130519

Starmans, C., & Friedman, O. (2012). The folk conception of knowledge. *Cognition*, *124*(3), 272–283. doi:10.1016/j.cognition.2012.05.017

Taylor, M., Lussier, G. L., & Maring, B. L. (2003). The Distinction Between Lying and Pretending. *Journal of Cognition and Development*, *4*(3), 299–323. doi:10.1207/S15327647JCD0403_04

Thomas, R. C., & Jacoby, L. L. (2013). Diminishing adult egocentrism when estimating what others know. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*(2), 473-486. doi: 10.1037/a0028883

Turri, J. (2012). In Gettier's wake. In S. Hetherington (Ed.), *Epistemology: The key thinkers* (pp. 214–229). London: Continuum.

Turri, J. (2013). The test of truth: An experimental investigation of the norm of assertion. *Cognition*, *129*(2), 279–291. doi:10.1016/j.cognition.2013.06.012

Turri, J. (2014). Skeptical appeal: The source-content bias. *Cognitive Science*. doi:10.1111/cogs.12153.

Turri, J. (2015). From virtue epistemology to abilism: theoretical and empirical developments. In C. B. Miller, M. R. Furr, A. Knobel, & W. Fleeson (Eds.), *Character: new directions from philosophy, psychology, and theology* (pp. 315-330). Oxford University Press.

Turri, J., Buckwalter, W., & Blouw, P. (2014). Knowledge and luck. *Psychonomic Bulletin & Review*. doi:10.3758/s13423-014-0683-5

Vincent, J. L., Patel, G. H., Fox, M. D., Snyder, A. Z., Baker, J. T., Van Essen, D. C., et al. (2007). Intrinsic functional architecture in the anaesthetized monkey brain. *Nature*, *447*(7140), 83–86. doi:10.1038/nature05758

Wells, G. L. (1992). Naked statistical evidence of liability: Is subjective probability enough? *Journal of Personality and Social Psychology*, *62*(5), 739. doi:10.1037/0022-3514.62.5.739

Wesson, C. J., & Pulford, B. D. (2009). Verbal expressions of confidence and doubt. *Psychological Reports*, *105*(1), 151-160. doi: 10.2466/PR0.105.1.151-160

Wimmer, H., Hogrefe, G.J., & Perner, J. (1988). Children's understanding of informational access as source of knowledge. *Child Development*, *59*(2), 386–396. doi: 10.2307/1130318

Woolley, J. D., & Wellman, H. M. (1993). Origin and truth: Young children's understanding of imaginary mental representations. *Child Development*, *64*(1), 1-17. doi: 10.1111/j.1467-8624.1993.tb02892.x

Wright, E. F., Maceachern, L., Stoffer, E., & Macdonald, N. (1996). Factors affecting the use of naked statistical evidence of liability. *The Journal of Social Psychology*, *136*(6), 677–688. doi:10.1080/00224545.1996.9712244

Young, L., & Saxe, R. (2011). When ignorance is no excuse: Different roles for intent across moral domains. *Cognition*, *120*(2), 202–214. doi:10.1016/j.cognition.2011.04.005

Yuill, N., & Perner, J. (1988). Intentionality and knowledge in children's judgments of actor‘s responsibility and recipient’s emotional reaction. *Developmental Psychology*, *24*(3), 358. doi:10.1037/0012-1649.24.3.358