

Knowledge Attributions and Lottery Cases: A Review and New Evidence*

John Turri

john.turri@gmail.com

Abstract: I review recent empirical findings on knowledge attributions in lottery cases and report a new experiment that advances our understanding of the topic. The main novel finding is that people deny knowledge in lottery cases because of an underlying qualitative difference in how they process probabilistic information. “Outside” information is generic and pertains to a base rate within a population. “Inside” information is specific and pertains to a particular item’s propensity. When an agent receives information that 99% of all lottery tickets lose (outside information), people judge that she does not know that her ticket will lose. By contrast, when an agent receives information that her specific ticket is 99% likely to lose (inside information), people judge that she knows that her ticket will lose. Despite this difference in knowledge judgments, people rate the likelihood of her ticket losing the exact same in both cases (i.e. 99%). The results shed light on other factors affecting knowledge judgments in lottery cases, including formulaic expression and participants’ own estimation of whether it is true that the ticket will lose. The results also undermine previous hypotheses offered for knowledge denial in lottery cases, including the hypotheses that people deny knowledge because they either deny justification or acknowledge a chance for error.

Keywords: social cognition; knowledge; belief; truth; evidence; inside/outside information; formulaic expression; philosophical method

Word count: 6780

* This is a draft (January 2016) of a paper forthcoming in *The Lottery Problem*, ed. I. Douven (Cambridge University press). Please check before citing or circulating.

Introduction

Appeals to ordinary thought and talk are frequent in philosophy. For example, Aristotle claimed that a theory of happiness should be evaluated in light of “what is commonly said about it” (Aristotle 350 BCE, 1098b, 9-11). He judged it improbable that a widely held view would be completely incorrect (1098b, 28-29). John Locke claimed that a theory of knowledge should be informed by how we ordinarily talk about knowledge (Locke 1690, book 4.11.3-8). More recently, J.L. Austin advised that “ordinary language” should get “the first word” in philosophical theorizing (Austin 1956, p. 11), and Wilfrid Sellars argued that identifying the defining features of ordinary thought — “the manifest image” — is “a task of the first importance” for philosophers (Sellars 1963, ch. 1). Relatedly, John Hawthorne claims that describing “the relatively inchoate thoughts and arguments that guide our intuitions is part of the job of a philosopher” (Hawthorne 2004, p. 16).

Many lines of research in contemporary epistemology are based on assumptions about ordinary thought and talk about knowledge — that is, about alleged central tendencies in “commonsense” or “folk” epistemology. This approach is evident, for instance, in the enormous recent literature on the semantics and pragmatics of knowledge attributions, where philosophers frequently appeal to how we would ordinarily think, talk, or act in certain situations (e.g. Goldman 1979; Dretske 1981; Stroud 1984; Craig 1990; Sosa 1991; DeRose 1992; Cohen 1999; Rysiew 2001; Stanley 2005; Brown 2005; Turri 2010). The tendency is also evident in work on knowledge judgments in lottery cases. For example, one epistemologist writes, “No matter how high the

odds that the ticket will not win, it strikes us that [Smith] doesn't know that [the] ticket will not win" (Vogel 1990, p. 292). Such claims are treated as "data" or "uncontroversial" starting points to which a theory of knowledge is accountable (Hawthorne 2004, p. 8).

Beyond making claims about our tendency to deny knowledge in lottery cases, epistemologists have made strong claims about the psychological processes informing our judgments. They propose to identify the factor "which prevents us from judging that subjects know" in lottery cases (DeRose 1996, p. 269), or "the relevant psychological forces driving" our knowledge judgments (Hawthorne 2004: 14). For example, one philosopher writes:

We can see ... that in the paradigm lottery situation, something like the following often goes on: The ascriber divides possibility space into a set of subcases, each of which, from the point of view of the subject, is overwhelmingly likely to not obtain, but which are such that the subject's grounds for thinking that any one of the subcases does not obtain is not appreciably different than his grounds for thinking that any other subcase does not obtain. [In other words], the relative strength of epistemic position with regard to each subcase is not appreciably different. In general, what is often at the root of the relevant lottery intuition is a division of epistemic space into a set of subcases with respect to which one's epistemic position seems roughly similar. Once such a division is effected, a parity of reasoning argument can kick in against the suggestion that one knows that a particular subcase does not obtain, namely: If one can know that *that* subcase does not obtain, one can know of each subcase that it does not

obtain. But it is absurd to suppose that one can know of each subcase that it does not obtain. (Hawthorne 2004, pp. 14-15)

The proposed process involves at least four distinct parts, instigated by the question whether an agent knows that a lottery ticket will lose. First, we partition possibility space into a set of roughly equiprobable elements, each of which corresponds to one ticket in the lottery winning, including one element corresponding to the ticket in question winning. Second, we judge that if the agent can know that the ticket in question will lose, then the agent can know that all of the tickets will lose. Third, we reject the conclusion that the agent can know that all of the tickets will lose. Fourth, and finally, on that basis we infer that the agent cannot know that the ticket in question will lose. The author offers no evidence that this description even approximately captures part of the psychological process involved. And, unfortunately, humans do not have the power to simply “see” that cognitive processes of such complexity occur. Anyone familiar with research on cognitive processing will recognize how challenging it would be to provide evidence for a model this complicated.

Philosophers have offered simpler psychological accounts of why we deny knowledge in lottery cases. For example, according to the *justification account*, we deny knowledge because we deny that the agent is justified or rational in believing that the ticket is a loser, and we think that knowledge requires justified belief (Nelkin 2000; Sutton 2007, pp. 48-53). According to the *chance account*, we deny knowledge because we think that there is a chance the agent is wrong about whether the ticket will lose, and we think that such a chance is incompatible with knowing (Cohen 1988, p. 196; Lewis 1996, p. 557).

To characterize ordinary thought and talk, philosophers often draw on their own experiences, social observation, and reflections about what we would say about certain situations (Locke 1690, bk. 4.11.3-8; MacIver 1938; Ducasse 1941, ch. 10; Fodor 1964; Wittgenstein 1975; Jackson 1998, ch. 2). Although this is a natural, and perhaps even unavoidable, starting point, this armchair methodology is severely limited. For example, one worry is that philosophers, being human, are susceptible to the false-consensus effect, whereby we tend to overestimate the extent to which others share our views (for a review, see Marks & Miller 1987). Another worry is that many cases philosophers use to illustrate what is “natural” or “intuitive” are polluted by thought-experimenter bias (Turri in 2016a). One result is that contemporary philosophers often badly mischaracterize commonsense and ordinary practice (for recent reviews, see Turri 2016a; Turri 2016c; Buckwalter & Turri 2018).

Nevertheless, we can overcome some limitations of armchair methodology and gain better evidence by using experimental and statistical tools from cognitive and social science. This is the approach adopted by the growing interdisciplinary field of experimental epistemology. Recent work in this field has begun investigating the knowledge judgments people actually make and why they make those judgments.

In one series of studies, approximately 70%-90% of participants denied knowledge in stereotypical lottery cases involving extremely long odds of winning (Turri & Friedman 2014). This provides evidence that people are inclined to deny knowledge in lottery cases when the evidence is statistical. However, this same series of studies also provided evidence against the justification account for why people deny knowledge in such cases. For example, in one study, only

9% of participants attributed knowledge, even though 80% attributed justification (i.e. judged that the agent was justified in thinking that the ticket was a loser) (Turri & Friedman 2014, Experiment 1). This result is unlikely if people deny knowledge because they deny justification.

In another study, results from a closely matched control condition provided evidence against the chance account for why people deny knowledge in lottery cases (Turri & Friedman 2014, Experiment 2). One group of participants read a stereotypical lottery case where there is only a one-in-ten-million chance that the agent's ticket wins; only 20% of these participants attributed knowledge and 88% judged that there was a chance the ticket was a winner. On its own, this is consistent with the chance account. However, another group of participants read a case where a newscaster announces the winning numbers and there is only a one-in-ten-million chance that he misreports the winning number; 66% of these participants attributed knowledge and 90% judged that there was a chance the agent's ticket was a winner. Comparing these two conditions, there is a large difference in knowledge attribution (20% versus 66%) despite a nearly identical rate attributing the chance of error (88% versus 90%). This is unlikely if people deny knowledge in lottery cases because they attribute a chance of error.

Researchers also provided evidence for a new hypothesis that might help explain why people deny knowledge in lottery cases: the *formulaic account*. According to this account, formulaic expression helps to explain why people tend to deny knowledge in lottery cases. Perhaps people are habituated to respond with the clichéd “You never know” in typical lottery cases. The formulaic account predicts that people are more likely to attribute knowledge in non-stereotypical lottery cases and in non-lottery cases where the odds are held constant. Results from two studies

supported these predictions.

In one study, participants read about an agent who receives a ten dollar bill with a ten digit serial number on it (Turri & Friedman 2014, Experiment 4). In one condition (“lotto”), the agent then learns that the serial number is over 99% unlikely to be this week’s winning lottery number, and she concludes that it is not the winning number. In another condition (“phone”), the agent instead learns that the serial number is over 99% unlikely to be Barack Obama’s mobile phone number. Participants in the lotto condition were asked whether the agent knows that the serial number is not the winning lottery number; participants in the phone condition were asked whether the agent knows that the serial number is not Obama’s phone number. In the lotto condition, 50% of participants attributed knowledge, which is significantly higher than the approximately 20% who attributed knowledge in earlier studies. In the phone condition, 63% of participants attributed knowledge, which was marginally higher than in the lotto condition.

In another study, participants also read about an agent who receives a ten dollar bill with a ten digit serial number on it (Turri & Friedman 2014, Experiment 5). The agent then learns from a statistician that the serial number “is just as likely to be Brad Pitt’s mobile phone number as it is to win this week’s lottery.” In one condition (“lotto”), the agent then concludes that the number will not win this week’s lottery. In the other condition (“phone”), the agent concludes that the number is not Brad Pitt’s phone number. The rate of knowledge attribution was higher in the phone condition.

Overall, then, the studies just reviewed suggest that philosophers were correct about central tendencies in knowledge judgments about some lottery cases. But their proposed explanations

are probably incorrect, and the central tendencies are importantly different in other lottery cases. The same line of research provided some initial support for the hypothesis that formulaic expression helps to explain knowledge denial in some lottery cases, but work is needed to further test this hypothesis and to identify other factors affecting people's knowledge judgments.

One potentially relevant factor is the difference between "outside" and "inside" probabilistic information. It is well documented that people often misunderstand probabilities (e.g. Kahneman & Tversky 1972). But even when they understand the probabilities, interesting differences arise depending on the statistical information's character. For example, mock jurors make different liability judgments across conditions where they assign equal probability to the defendant's guilt (Wells 1992). 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, and 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, but in this case they tend to *agree* that the jury should find the Blue Bus Company liable.

Outside information is generic and concerns distributions of properties or base rates. For example, in one version of the story about the bus, a certain percentage of buses operating in

town that day belonged to the Blue Bus Company. Inside information concerns a specific item in the case at hand. For example, in the other version of the story about the bus, the weigh station attendant entered information about the very bus that hit the dog. The attendant's observation 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" information.

Recent studies provide strong evidence that the outside/inside difference affects knowledge attributions. In one study, participants read a story about Gary, who is suing the Blue Cab Company (Turri, Friedman & Keefner 2017, Experiment 4). 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. At this point, the story ends in one of two ways. One group of people was told that 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" (outside information). The other group was told that according to the analysis, "the cab was 80% likely to be a Blue Cab" (inside information). People were then asked to rate the probability that a Blue cab destroyed the garden, whether the jurors know that a Blue cab destroyed the garden, and whether the jurors should rule against the Blue Cab Company. People in both groups responded to the probability question similarly — on average, they judged the probability to be approximately 80% — but they gave significantly different answers to the questions about what the jurors know and should decide. In other words, the outside/inside difference affected knowledge judgments and decision evaluations, but it did not affect probabili-

ty estimates. Outside information made people significantly less likely to attribute knowledge and to judge that the jurors should rule against the company. (For related results, see Friedman & Turri 2015, Experiments 3A and 3B.)

This line of research suggests a hypothesis worth testing, which might be called the *placement account*. According to this account, people deny knowledge in typical lottery cases because the statistical evidence seems to consist of outside information, and they do not think that outside information suffices for knowledge. This hypothesis, along with prior findings, predicts that if we could make the probabilistic evidence about the ticket seem like inside information, then people's willingness to attribute knowledge would increase significantly. It also predicts that an increase in knowledge attribution will occur even though people acknowledge a chance that the ticket will win.

In the next section, I report an experiment that tests the placement account by comparing knowledge judgments about lottery cases in conditions involving either outside or inside probabilistic information. The experiment also provides a further test of the formulaic account by comparing knowledge judgments about lottery cases to knowledge judgments about non-lottery cases. In this instance, the non-lottery cases involved an article of clothing.

The experiment differs methodologically from previous work on knowledge judgments about lottery cases, in two main ways. First, whereas previous work collected knowledge judgments using dichotomous response options, I used more sensitive scaled measures, which allow for more powerful statistical tests. Second, previous work adopted a "stipulate and question" approach to some crucial details of the cases tested. For instance, it is stipulated that the lottery

ticket is a loser and then participants who later fail a comprehension question about this detail are excluded from the analysis. But in ordinary life we typically do not have such privileged information prior to judging whether someone has knowledge. Instead we have to judge whether the ticket will win for ourselves, in the same context where we pass judgment on whether someone knows. In this respect the stipulate-and-question approach's ecological validity is diminished. Accordingly I adopted a slightly different approach for the experiment reported here. Instead of stipulating privileged details, I allowed participants to judge those details for themselves. In particular, I allowed them to judge a range of factors that philosophers typically assume are relevant to knowledge judgments, including whether the relevant proposition is true, whether the agent believes the proposition, the quality of the agent's evidence, and whether the agent should assert the proposition.

In addition to conferring greater ecological validity, the experimental procedures used below provide another advantage. They allow me to analyze responses and assess whether knowledge denials in lottery cases are due to a range of underlying judgments, including whether people themselves doubt that the ticket is a loser. Some philosophers have assumed that knowledge denials in lottery cases are not due to doubts about the target proposition's truth value (e.g. Hawthorne 2004, p. 8), but this has never been tested. In other words, the following experiment also allows me to test a *truth account* of knowledge denial in lottery cases, which says that people deny knowledge because they doubt the truth, which they view as essential to knowledge.

Experiment

Method

Two hundred and nine participants were tested (aged 18-99, mean age = 34 years; 97 female; 97% reporting English as a native language). Participants were U.S. residents, recruited and tested online using Amazon Mechanical Turk (AMT) and Qualtrics.

Participants were randomly assigned to one of four conditions in a 2 (Item: ticket, beret) × 2 (Information: outside, inside) between-subjects design. All participants read a brief scenario, responded to six test items, then completed a brief demographic questionnaire. The basic story was about Janice, who is asked a question about something of hers. The item is either a lottery ticket or a French beret. Janice then gathers information about the item. The information is either generic and pertains to a property's distribution within a population (outside), or specific and pertains to a particular item's propensity to have a property (inside). Here is the text of the stories, with the Information manipulation in brackets (outside/inside):

(Ticket) Janice is asked whether her lottery ticket is a loser. [She conducts research on lottery tickets, which shows that 99% of all lottery tickets lose. / She conducts a statistical analysis of her ticket's numbers, which shows that her specific numbers are 99% likely to lose.]

(Beret) Janice is asked whether her French beret contains lidinium fibers. [She conducts research on French berets, which shows that 99% of all French berets contain lidinium fibers. / She conducts a chemical analysis of her French beret,

which shows that her French beret is 99% likely to contain lidinium fibers.]

Beneath the story were five test statements in a question matrix. The statements included a belief attribution, an evaluation of the key proposition's truth-value, an evaluation of Janice's evidence, an evaluation of whether Janice should make an assertion, and a knowledge attribution.

For the ticket scenario, these were the statements:

Janice thinks that her lottery ticket is a loser. (think)

It's true that Janice's lottery ticket is a loser. (true)

Janice has good evidence that her lottery ticket is a loser. (evidence)

Janice should say that her lottery ticket is a loser. (assert)

Janice knows that her lottery ticket is a loser. (know)

For the beret scenario, these were the statements:

Janice thinks that her French beret contains lidinium. (think)

It's true that Janice's French beret contains lidinium. (true)

Janice has good evidence that her French beret contains lidinium. (evidence)

Janice should say that her French beret contains lidinium. (assert)

Janice knows that her French beret contains lidinium. (know)

Responses were collected on a standard 7-point Likert scale, 1 "strongly disagree" – 7 "strongly agree," left-to-right on the participant's screen. The order of test statements was randomized.

Participants then advanced to a new screen and rated the percentage chance that the key proposition was true. They were asked one question, depending on which item they read about:

On a scale of 0% to 100%, how likely is it that Janice's ticket is a loser?

On a scale of 0% to 100%, how likely is it that Janice’s beret contains lidinium?

Responses were collected in a text box below the question. Finally, participants advanced to a new screen and completed a demographic questionnaire.

Results

Preliminary regression analyses revealed no effect of participant age or sex on any of the dependent variables, so I omitted these demographic variables from all other analyses and will not discuss them any further. Table 1 contains descriptive statistics for all the dependent variables in the four conditions. (See the Appendix for all tables.) There were main effects of Item or Information on all the dependent variables except for the percentage task, which was affected by neither factor. There were no significant interactions. (See Table 2.)

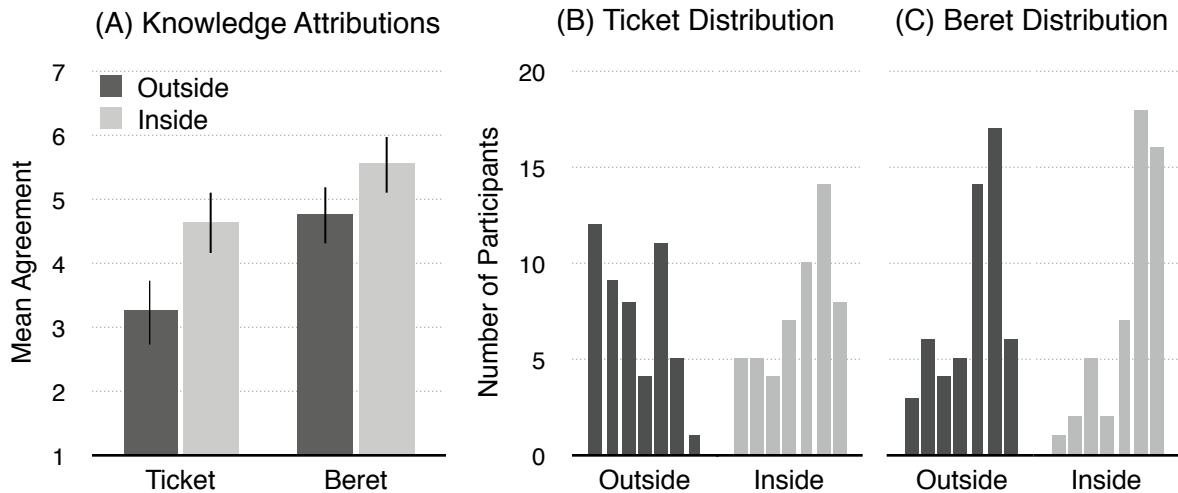


Figure 1. Mean knowledge attributions (Panel A) and the distribution of responses (Panels B and C) in the four conditions. Scales ran 1 (strongly disagree) – 7 (strongly agree). Error bars represent bootstrapped 95% confidence intervals.

Knowledge attributions were significantly higher when Janice had inside information than when she had outside information. This same pattern was observed whether the item was a ticket or a beret. (See Figure 1 and Table 3.) Knowledge attributions were significantly higher when Janice focused on her ticket than when she focused on her beret. This same pattern was observed whether she acquired outside information or inside information. (See Figure 1 and Table 4.) When Janice had outside information about her ticket, people tended to deny knowledge: the mean, median and mode were all below the scale's midpoint. (See Figure 1 and Tables 1 and 5.) In all three other conditions, people tended to attribute knowledge: the mean, median and mode were all above the midpoint. (See Figure 1 and Tables 1 and 5.)

Because there was no interaction effect on any dependent variable, I conducted one multiple linear regression analysis to investigate which factors significantly predicted participants' knowledge attributions. The predictors included both independent variables (Item and Evidence) and all five remaining dependent variables (belief attributions, truth evaluations, evidence evaluations, assertability attributions, and the percentage rating). The model was statistically significant and explained most (62%) of the variance in knowledge attributions. (See Table 6.) The only predictors that made a unique statistically significant contribution were truth evaluations and assertability attributions. That is, when controlling for the effect of the other predictors in the model, only participants' evaluations of the truth and assertability significantly predicted knowledge judgments. A more selective model including only truth evaluations and assertability attributions as predictors was also statistically significant and explained almost the exact same amount (61%) of the variance in knowledge attributions. (See Table 6.)

Finally, in the entire data set, 61% of participants (127 of 209) agreed with the knowledge attribution to one extent or another (i.e. selected “slightly agree,” “agree,” or “strongly agree” of the seven response options), which significantly exceeds chance rates (binomial test, $p < .001$, test proportion = .429). Of the participants who agreed to one extent or another, 96% (122 of 127) indicated that there was a chance that the known proposition was false (i.e. answered less than 100% on the percentage task). In the entire data set, 15% of participants (31 of 209) “strongly agreed” with the knowledge attribution. Of these participants, 94% (29 of 31) indicated that there was a chance that the known proposition was false.

Conclusion

The results from a new behavioral experiment advance our understanding of knowledge judgments in lottery cases. People were more likely to agree that an agent knows that a lottery ticket is a loser based on “inside” information (i.e. specific information about a particular item’s propensity) than on “outside” information (i.e. generic information about a property’s distribution in a population). Indeed, the central tendency was to *attribute* knowledge based on inside information and to *deny* knowledge based on outside information. This supports the placement account, which says that people deny knowledge in lottery cases because the probabilistic information seems generic rather than specific. The observed difference in knowledge attributions occurred despite the fact that people judged it equally likely that the ticket would lose based on outside and inside information. This contributes to a growing body of evidence in cognitive psy-

chology and experimental philosophy that there are qualitative differences in how people process probabilistic information from various sources, which has consequences for important social judgments, such as knowledge attributions and evaluations of how people should behave (see Wells 1992; Friedman & Turri 2015; Turri, Friedman & Keefner 2017; Turri 2018).

People were more likely to attribute knowledge that a beret contains certain fibers than to attribute knowledge that a lottery ticket is a loser, even though they judged those outcomes to be equally likely. This further supports the formulaic account, which says that formulaic expression depresses knowledge attribution in lottery cases. The “inside” beret case and the “inside” lottery case differed in more than just the item under consideration. In particular, the inside beret case involved a “chemical” test, whereas the inside lottery ticket case involved a “statistical” test. Thus this difference could be due to a greater trust in chemistry than in statistics. However, this alternative explanation seems unlikely because the “outside” versions of the beret and lottery cases produced the same asymmetry in knowledge attributions but contained no such confound. In each of the outside cases, the agent conducted “research” showing that 99% of the items had the relevant property. Of course, it is possible that participants tended to imagine different research being conducted for hats than for lottery tickets, but there is no reason to think that this actually happened and drove the finding.

People’s knowledge judgments were very well predicted by a combination of two other judgments: their evaluations of whether the target proposition was true and whether the agent should assert it. This provides some support for the truth account, which says that people deny knowledge in lottery cases because they doubt the truth of the proposition. By contrast, people’s

knowledge judgments were not significantly predicted by belief attributions or their evaluation of the agent's evidence for the target proposition. This further undermines the justification account, which says that people deny knowledge in lottery cases because they deny justification. It also adds to a string of recent findings which, taken together, constitute a compelling case that knowledge attributions are frequently not based on belief attributions (see Myers-Schulz & Schwitzgebel 2013; Murray, Sytsma & Livengood 2013; Turri & Buckwalter 2017; Turri 2016d; Turri, Buckwalter & Rose 2016). This calls into question the widespread assumption in philosophy that knowledge requires belief, as those categories are ordinarily understood.

Overall, a majority of participants in the study attributed knowledge and nearly all who attributed knowledge also judged that there was a chance the known proposition was false. This further undermines the chance account, which says that people deny knowledge in lottery cases because they acknowledge a chance of error. It also fits with other recent findings on knowledge attributions involving unreliable agents. For example, in one series of studies, the vast majority of participants attributed knowledge to an agent who is wrong 90% of the time and whom participants categorized as unreliable (Turri 2016b, Experiments 1-3). In this same series of studies, people's knowledge attributions were insensitive to whether the agent was highly reliable or unreliable (Experiments 4-5), and people attributed knowledge despite judging that the knower easily might have made a mistake (Experiment 9). At this point, it is clear that any form of infallibilism about knowledge is radically revisionary. This pertains to views requiring some form of unrestricted infallibility (e.g. Neta 2011; BonJour 2010; Unger 1975), as well as "contextualist" or "relevant alternative" approaches requiring infallibility relative to some restricted set of possibil-

ities (e.g. Dretske 1970; Stine 1976; Goldman 1976; Lewis 1996; DeRose 2009; Cohen 2013).

In closing, I would like to emphasize something noted in previous research on knowledge judgments about lottery cases (see the “General Discussion” in Friedman & Turri 2015 and in Turri & Friedman 2014). Current evidence shows that there will be no simple, general explanation for people’s judgments in lottery cases. What might at first appear to be consequential details end up not mattering, and what might appear to be inconsequential details end up mattering. Further work is required in order to understand the full range of factors informing judgments about these cases. In other words, it seems likely that research on this issue has reached the end of the beginning, rather than the beginning of the end.

Acknowledgments — For helpful feedback and discussion, I thank Ori Friedman, Ashley Keefner, Vanessa Lam, and Angelo Turri. This research was supported by the Social Sciences and Humanities Research Council of Canada, the Ontario Ministry of Economic Development and Innovation, and the Canada Research Chairs program.

Appendix

Table 1. Descriptive statistics for all dependent variables in the four conditions. N = number of participants. Mo = mode. Md = median. M = mean. SD = standard deviation.

Measure	Outside								Inside							
	Ticket (N = 50)				Beret (N = 55)				Ticket (N = 53)				Beret (N = 51)			
	Mo	Md	M	SD	Mo	Md	M	SD	Mo	Md	M	SD	Mo	Md	M	SD
think	6	6	5.30	1.57	7	6	5.98	1.03	6	6	5.64	1.02	7	6	6.22	1.06
true	4	4	3.58	1.77	6	5	4.96	1.61	6	4	4.51	1.89	6	6	5.88	1.29
evidence	6	6	5.52	1.53	7	6	6.15	1.18	6,7	6	5.81	1.29	7	7	6.27	1.12
assert	4	4	3.84	1.77	6	6	5.58	1.32	5	5	5.06	1.60	6	6	6.00	1.26
know	1	3	3.24	1.83	6	5	4.75	1.72	6	5	4.62	1.89	6	6	5.55	1.57
percent	99	99	98.2	5.52	99	99	95.7	11.93	99	99	98.0	4.78	99	99	97.5	5.73

Table 2. Univariate analyses of variance for all dependent variables. df = degrees of freedom. F = F-statistic. p = p-value. η_p^2 = partial eta squared.

Measure	Factor											
	Item				Information				Item*Information			
	df	F	p	η_p^2	df	F	p	η_p^2	df	F	p	η_p^2
think	1, 205	14.63	<.001	.067	1, 205	3.07	.081	.015	1, 205	0.11	.743	.001
true	1, 205	36.12	<.001	.150	1, 205	16.23	<.001	.073	1, 205	0.01	.981	.001
evidence	1, 205	9.39	.002	.044	1, 205	1.40	.238	.007	1, 205	0.21	.648	.001
assert	1, 205	42.01	<.001	.170	1, 205	15.57	<.001	.071	1, 205	3.72	.055	.018
know	1, 205	25.00	<.001	.109	1, 205	20.20	<.001	.090	1, 205	1.42	.235	.007
percent	1, 205	2.16	.144	.010	1, 205	0.56	.456	.003	1, 205	0.90	.343	.004

Table 3. Independent samples t-tests comparing mean knowledge attributions in outside and inside information conditions for each item (ticket, beret). N = number of participants. M = mean. SD = standard deviation. t = t-statistic. df = degrees of freedom. p = p-value. MD = mean difference. 95% CI = ninety-five percent confidence interval for MD. d = Cohen's d.

Item	Outside			Inside			t	df	p	MD	95% CI	d
	N	M	SD	N	M	SD						
Ticket	50	3.24	1.83	53	4.62	1.89	-3.77	101	<.001	-1.38	-2.11, -0.66	0.75
Beret	55	4.75	1.72	51	5.55	1.57	-2.51	104	.014	-0.80	-1.44, -0.17	0.50

Table 4. Independent samples t-tests comparing mean knowledge attributions in ticket and beret conditions for each information type (outside, inside).

Information	Ticket			Beret			t	df	p	MD	95% CI	d
	N	M	SD	N	M	SD						
Outside	50	3.24	1.83	55	4.75	1.72	-4.35	103	<.001	-1.51	-2.19, -0.82	0.86
Inside	53	4.62	1.89	51	5.55	1.57	-2.71	102	.008	-0.93	-1.60, -0.25	0.54

Table 5. One sample t-tests on mean knowledge attributions (test value = 4). t = t-statistic. df = degrees of freedom. MD = mean difference. 95% CI = ninety-five percent confidence interval for MD. d = Cohen's d.

Item	Outside						Inside					
	t	df	p	MD	95% CI	d	t	df	p	MD	95% CI	d
Ticket	-2.95	49	.005	-0.76	-1.28, -0.24	0.42	2.39	52	.020	0.62	0.10, 1.14	0.33
Beret	3.21	54	.002	0.75	0.28, 1.21	0.44	7.06	50	<.001	1.55	1.11, 1.99	0.99

Table 6. Multiple linear regressions predicting knowledge attributions. B = unstandardized beta coefficient. SE(B) = standard error of B. β = standardized beta coefficient. t = t-statistic. p = p-value.

Full model						Selective model					
Predictor	B	SE (B)	β	t	p	Predictor	B	SE (B)	β	t	p
Constant	.089	1.17		0.08	.940	Constant	-.105	0.28		-0.38	.705
Item	-.059	0.19	-.015	-0.31	.755	true	.426	0.06	.406	6.99	<.001
Information	.289	0.18	.075	1.64	.104	assert	.514	0.07	.450	7.75	<.001
think	.066	0.08	.042	0.80	.427						
true	.425	0.07	.405	6.55	<.001						
evidence	-.021	0.08	-.014	-0.26	.793						
assert	.487	0.07	.427	6.62	<.001						
percent	-.004	0.01	-.015	-0.33	.742						

Note. Full model: $F(7, 201) = 46.05$, $p < .001$, $R^2 = .616$. Reference class for Item: ticket. Reference class for Information: outside. Selective model: $F(2, 206) = 160.23$, $p < .001$, $R^2 = .609$.

References

- Aristotle. (1941). Posterior analytics. In R. McKeon (Ed.), G. R. G. Mure (Trans.), The basic works of Aristotle. New York: Random House.
- Austin, J. L. (1956). A plea for excuses. *Proceedings of the Aristotelian Society*, 57, 1–30.
- BonJour, L. (2010). The myth of knowledge. *Philosophical Perspectives*, 24(1), 57–83. <http://doi.org/10.1111/j.1520-8583.2010.00185.x>
- Brown, J. (2005). Comparing contextualism and invariantism on the correctness of contextualist intuitions. *Grazer Philosophische Studien*, 69(1), 71–100.
- Cohen, S. (1988). How to be a Fallibilist. *Philosophical Perspectives*, 2, 91–123.

- Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, 13, 57–89.
- Cohen, S. (2013). Contextualism defended. In M. Steup, J. Turri, & E. Sosa (Eds.), *Contemporary debates in epistemology* (2nd ed., pp. 69–75). Malden, Mass.: Wiley-Blackwell.
- Craig, E. (1990). *Knowledge and the state of nature: an essay in conceptual synthesis*. Oxford: Oxford University Press.
- DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52(4), 913–929.
- DeRose, K. (1996). Knowledge, assertion and lotteries. *Australasian Journal of Philosophy*, 74(4), 568–580.
- DeRose, K. (2009). *The case for contextualism*. Oxford: Oxford University Press.
- Dretske, F. I. (1970). Epistemic operators. *The Journal of Philosophy*, 67(24), 1007–1023.
- Dretske, F. I. (1981). *Knowledge and the flow of information*. Cambridge, Mass.: MIT Press.
- Ducasse, C. J. (1941). *Philosophy as a science: its matter and its method*. New York: Oskar Piest.
- Fodor, J. A. (1964). On knowing what we would say. *The Philosophical Review*, 73(2), 198–212.
- Friedman, O., & Turri, J. (2015). Is probabilistic evidence a source of knowledge? *Cognitive Science*, 39(5), 1062–1080.
- Goldman, A. I. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, 73(20), 771–791.
- Goldman, A. I. (1979). What is justified belief? In G. Pappas (Ed.), *Justification and knowledge*. Dordrecht: Reidel.

- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.
- Jackson, F. (1998). *From metaphysics to ethics*. Oxford: Oxford University Press.
- Kahneman, D., & Tversky, A. (1972). Subjective probability: a judgment of representativeness. *Cognitive Psychology*, 3(3), 430–454.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4), 549–567.
- Locke, J. (1975). *An essay concerning human understanding*. (P. H. Nidditch, Ed.). Oxford: Clarendon Press.
- MacIver, A. M. (1938). Some questions about 'know' and "think." *Analysis*, 5(3), 43–50.
- Marks, G., & Miller, N. (1987). Ten years of research on the false-consensus effect: An empirical and theoretical review. *Psychological Bulletin*, 102(1), 72. <http://doi.org/10.1037/0033-2909.102.1.72>
- Murray, D., Sytsma, J., & Livengood, J. (2013). God knows (but does God believe?). *Philosophical Studies*, 166(1), 83–107. <http://doi.org/10.1007/s11098-012-0022-5>
- Myers-Schulz, B., & Schwitzgebel, E. (2013). Knowing that P without believing that P. *Nous*, 47(2), 371–384.
- Nelkin, D. K. (2000). The lottery paradox, knowledge, and rationality. *The Philosophical Review*, 109(3), 373–409.
- Neta, R. (2011). A refutation of Cartesian fallibilism. *Nous*, 45(4), 658–695.
- Rysiew, P. (2001). The Context-Sensitivity of Knowledge Attributions. *Nous*, 35(4), 477–514. <http://doi.org/10.1111/0029-4624.00349>
- Sellars, W. (1963). *Science, perception and reality*. Atascadero, CA: Ridgeview Publishing Com-

pany.

- Sosa, E. (1991). *Knowledge in Perspective*. Cambridge: Cambridge University Press.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford University Press.
- Stine, G. C. (1976). Skepticism, relevant alternatives, and deductive closure. *Philosophical Studies*, 29, 249–261.
- Stroud, B. (1984). *The significance of philosophical skepticism*. Oxford: Clarendon Press.
- Sutton, J. (2007). *Without Justification*. MIT Press.
- Turri, J. (2010). Epistemic invariantism and speech act contextualism. *Philosophical Review*, 119(1), 77–95. <http://doi.org/10.1215/00318108-2009-026>
- Turri, J. (2016a). Knowledge judgments in “Gettier” cases. In J. Sytsma & W. Buckwalter (Eds.), *A companion to experimental philosophy* (pp. 337–348). Wiley-Blackwell.
- Turri, J. (2016b). A new paradigm for epistemology: from reliabilism to abilism. *Ergo*, 3(8), 189–231.
- Turri, J. (2016c). How to do better: toward normalizing experimentation in epistemology. In J. Nado (Ed.), *Advances in experimental philosophy and philosophical methodology* (pp. 35–51). London: Bloomsbury Academic.
- Turri, J. (2016d). Vision, knowledge, and assertion. *Consciousness and Cognition*, 41(C), 41–49.
- Turri, J. (2018). Epistemic modals and alternative possibilities. *Erkenntnis*, 83(5), 1063–1084.
- Turri, J., & Buckwalter, W. (2017). Descartes’s schism, Locke’s reunion: completing the pragmatic turn in epistemology. *American Philosophical Quarterly*, 54(1), 25–46.
- Buckwalter, W., & Turri, J. (2018). Moderate scientism in philosophy. In J. R. Ridder, R. Peels,

- & R. van Woudenberg (Eds.), *Scientism: prospects and problems*. Oxford University Press.
- Turri, J., & Friedman, O. (2014). Winners and losers in the folk epistemology of lotteries. In J. R. Beebe (Ed.), *Advances in experimental epistemology* (pp. 45–69). Bloomsbury.
- Turri, J., Buckwalter, W., & Rose, D. (2016). Actionability judgments cause knowledge judgments. *Thought*, 5(3), 212–222.
- Turri, J., Friedman, O., & Keefner, A. (2017). Knowledge central: a central role for knowledge attributions in social evaluations. *Quarterly Journal of Experimental Psychology*, 70(3), 504–515.
- Unger, P. (1975). *Ignorance: a case for skepticism*. Oxford: Oxford University Press.
- Vogel, J. (1990). Are there Counterexamples to the Closure Principle? In M. D. Roth & G. Ross (Eds.), *Philosophical Studies Series* (Vol. 48, pp. 13–27). Springer. http://doi.org/10.1007/978-94-009-1942-6_2
- Wells, G. L. (1992). Naked statistical evidence of liability: Is subjective probability enough? *Journal of Personality and Social Psychology*, 62(5), 739. <http://doi.org/10.1037/0022-3514.62.5.739>
- Wittgenstein, L. (1975). On certainty. (G. Anscombe & G. H. V. Wright, Eds., D. Paul & G. Anscombe, Trans.). Malden, MA: Blackwell.