



Cognitive Science (2014) 1–19

Copyright © 2014 Cognitive Science Society, Inc. All rights reserved.

ISSN: 0364-0213 print / 1551-6709 online

DOI: 10.1111/cogs.12182

# Is Probabilistic Evidence a Source of Knowledge?

Ori Friedman,<sup>a</sup> John Turri<sup>b</sup>

<sup>a</sup>*Department of Psychology, University of Waterloo*

<sup>b</sup>*Department of Philosophy, University of Waterloo*

Received 27 March 2013; received in revised form 7 February 2014; accepted 21 March 2014

---

## Abstract

We report a series of experiments examining whether people ascribe knowledge for true beliefs based on probabilistic evidence. Participants were less likely to ascribe knowledge for beliefs based on probabilistic evidence than for beliefs based on perceptual evidence (Experiments 1 and 2A) or testimony providing causal information (Experiment 2B). Denial of knowledge for beliefs based on probabilistic evidence did not arise because participants viewed such beliefs as unjustified, nor because such beliefs leave open the possibility of error. These findings rule out traditional philosophical accounts for why probabilistic evidence does not produce knowledge. The experiments instead suggest that people deny knowledge because they distrust drawing conclusions about an individual based on reasoning about the population to which it belongs, a tendency previously identified by “judgment and decision making” researchers. Consistent with this, participants were more willing to ascribe knowledge for beliefs based on probabilistic evidence that is specific to a particular case (Experiments 3A and 3B).

*Keywords:* Knowledge; Folk epistemology; Probabilistic evidence; Theory of mind

---

## 1. Introduction

Judging what others know is a socially important ability. Suppose you want to know whether a local zoo has zebras. You ask your friend Zoe, and she says that it does. Judging that Zoe knows this implies that her answer is trustworthy and that you can safely base your actions on it. But if it is something she only believes, you might be less sure. Likewise, her knowledge means that it would be redundant to tell Zoe that there are zebras at the zoo. Such examples show that knowledge is a socially important status relevant to trust, communication, and coordinating action. Knowledge attributions are not the

---

Correspondence should be sent to John Turri, Department of Philosophy, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada. E-mail: john.turri@gmail.com; or Ori Friedman, Department of Psychology, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada. E-mail: friedman@uwaterloo.ca

only possible way for a community of rational beings to flag good sources of information and make decisions about whom to trust or what to do. But knowledge attributions are in fact one main way that humans accomplish those things, which has motivated considerable research on knowledge attributions in both philosophy and cognitive science (Beebe & Buckwalter, 2010; Craig, 1990; Fantl & McGrath, 2009; Hawthorne, 2004; Locke, 1690; Schaffer & Knobe, 2012; Stanley, 2005; Turri, 2011a, 2013b; Williamson, 2000).

Philosophers have long theorized about the necessary conditions on knowledge (e.g., Chisholm, 1989; Descartes, 1641; Dretske, 1981; Gettier, 1963; Nozick, 1981; Plato, 370 BCE). More recently, psychologists and experimental philosophers have examined this question experimentally.<sup>1</sup> Their studies reveal that in judging whether an agent's belief counts as knowledge, people typically require that the belief is true and that it is justified—it must be based on evidence and not just on guesses or loose conjecture (Buckwalter, in press; Starmans & Friedman, 2012). However, not all forms of evidence lead to knowledge. For example, knowledge is not gained from misleading evidence (Starmans & Friedman, 2012; Turri, 2013a). Suppose Zoe came to believe zebras are at the zoo because she saw some mules convincingly disguised as zebras (Dretske, 1970). Although her belief will still be true if real zebras are somewhere else in the zoo, the truth of her belief would have nothing to do with the evidence supporting it. People deny knowledge for beliefs based on such “apparent evidence.”<sup>2</sup>

People may deny knowledge for beliefs based on other kinds of evidence. Many philosophers argue that knowledge is not gained from *probabilistic* evidence (Dretske, 1981; Lewis, 1996; Neta, 2011; for a contrary, minority view, see Turri, 2011b, in press a). For example, in considering “lottery cases,” philosophers deny that considering the poor odds of winning can lead one to know that a particular lottery ticket is a loser. Lay people might also deny that knowledge is gained from probabilistic evidence, because beliefs based on probabilistic evidence are *fallible* (i.e., they leave open the chance of error). In lottery cases, the odds do not conclusively rule out the (extremely unlikely) possibility that a certain ticket is a winner. People might also deny knowledge if they view beliefs based on probabilistic evidence as unjustified (Nelkin, 2000; Sutton, 2007). Probabilistic evidence might justify the belief that a certain ticket is very probably a loser, but without justifying the stronger belief that the ticket is a loser.

People might also deny that probabilistic evidence leads to knowledge because basing a belief on probabilistic evidence often involves drawing a conclusion about a *particular case* (e.g., this lottery ticket) by considering in relation to a larger *class* or *distribution* (e.g., all tickets in this lottery drawing). Such judgments are described as taking an “outside view” because they are based on facts external to the case at hand (Kahneman & Tversky, 1982; Lagnado & Sloman, 2004; also see Griffin, Koehler, & Brenner, 2007). People might be unwilling to attribute knowledge based on distributional information because they are reluctant to base conclusions on distributional information in *many* kinds of judgments. Relevant examples arise in legal decision-making. For example, discussion of the “gatecrasher paradox” (Cohen, 1981) centers on why a random individual attending a rodeo cannot be sued for non-payment, even though distributional evidence suggests that he probably did not pay (i.e., *most* audience members sneaked in without paying).

Wells (1992) notes that legal cases that rely on evidence not specifically tied to the particular case (termed “naked” statistical evidence) are typically thrown out of court. Moreover, Wells provides experimental evidence that laypeople and judges are typically unwilling to make liability judgments on the basis of such evidence (also see Arkes, Shoots-Reinhard, & Mayes, 2012).

Whereas people often neglect distributional information, they are typically willing to take an “inside view” and base judgments on information *specific* to a particular case, *even when this information is probabilistic*. This tendency to weigh specific information over distributional information arises in legal decision-making (e.g., Arkes et al., 2012; Koehler, 2001), and it is apparent in experiments on base-rate neglect more generally (see Barbey & Sloman, 2007, for an overview of base-rate neglect, including discussion of manipulations increasing use of base-rates). For example, when people judge the probability that a taxicab involved in an accident was green, they strongly weigh the reliability of the testimony of a witness who identified it as green (information specific to the case), while neglecting the distribution of green and blue taxis in the city (distributional information; e.g., Bar-Hillel, 1980; Tversky & Kahneman, 1982). If reliance on specific information extends to judgments about knowledge, people might ascribe knowledge for probabilistic evidence when it provides information *specific* to a particular case.

We conducted a series of experiments to examine whether people ascribe knowledge for true beliefs based on probabilistic evidence. Experiments 1, 2A, and 2B demonstrate that people often deny knowledge for such beliefs, especially in comparison to cases where beliefs are based on perceptual evidence or definitive testimony. These experiments also show that this reluctance does not arise either because beliefs based on probabilistic evidence are unjustified, or because such beliefs leave open the possibility of error. Experiments 3A and 3B also show this and provide evidence that people deny knowledge because they distrust basing conclusions about an individual on reasoning about the population to which it belong (“outside view”). However, people are more willing to ascribe knowledge for beliefs based on probabilistic evidence that is specific to a particular case (“inside view”).

## 2. Experiment 1

To first examine whether people ascribe knowledge for true beliefs based on probabilistic evidence, participants read brief stories based on classic thought problems from philosophy. In each story, an agent comes to have a true belief on the basis of evidence that is very strong, but which does not guarantee utter certainty. After reading the story, participants were asked questions assessing whether the agent is knowledgeable, whether her belief is justified, and whether there is a chance her belief could have been false. To assess knowledge ascriptions for beliefs based on probabilistic evidence, participants read a story (“Lotto”) in which an agent believes a lotto ticket is a loser; this story is based on classic “lottery cases” (e.g., Hawthorne, 2004; Vogel, 1990). For comparison, other participants read stories in which the agent’s belief is based on perceptual evidence, which could be fallible. In these stories, an agent sees an animal and correctly concludes it is a

zebra, even though there was a small chance that it could be a disguised mule. One version of this story is set in a zoo (“Zoo”) and is closely based on Dretske’s zebra/mule thought experiments (Dretske, 1970). We also included a second version of the story with the animal spotted in a farmer’s field (“Farm”).

## 2.1. Methods

### 2.1.1. Participants

One hundred thirty-eight participants (aged 18–63 years,  $M_{\text{age}} = 30$  years; 42 female) were tested. Data were excluded from two participants who failed comprehension questions designed to ensure that participants were reading scenarios carefully. In this study, and those subsequent, participants were recruited and tested using an online platform (Amazon Mechanical Turk and Qualtrics) and were paid \$0.25 for approximately 2 min of time. Participation was restricted to those residing in the United States. Repeat participation, within and across experiments, was prevented by allowing each AMT Worker to complete the task only once; as an added precaution, we screened data for responses from participants with identical Worker IDs (there were none).

### 2.1.2. Materials and procedure

Participants were randomly assigned to read one of three stories: Lotto, Zoo, or Farm. A between-subject design was used, with participants reading one story only, because previous work suggests that knowledge ascriptions are sometimes subject to order effects (Wright, 2010). In Lotto, Abigail comes to believe that a certain lotto ticket is a loser. It is left tacit that she concludes this from the strong odds that it is a loser, and that there is a small possibility that the ticket could be a winner. Here is the story:

Abigail is out shopping with her son. In a store, they see a man with a super lotto ticket. Abigail’s son says, “I bet that ticket’s not a loser. It might win the jackpot!” Abigail answers, “It is a losing ticket.” And Abigail is exactly right: The ticket is a loser.

In Zoo and Farm, the belief is based on perceptual evidence. In both stories, Abigail sees a black-and-white striped animal and correctly concludes that it is a zebra, and not a painted mule. Here is the text for these stories, with differences between the stories indicated in square brackets.

Abigail is [visiting the zoo/driving past a farm] with her son. In a [pen/field], they see a black-and-white striped animal. Abigail’s son says, “I bet that’s not a real zebra. It might be a painted mule!” Abigail answers, “It is a real zebra.” And Abigail is exactly right: The animal is a zebra.

After reading the story, participants answered a series of eight questions. These included three comprehension questions to make sure that participants read the story

carefully. Participants were also asked a knowledge question, a justification question, and a question about the chance of error. Following the knowledge and justification questions, participants were asked questions to rate how confident they were in their answers to these questions. Questions were as follows; response options are in brackets.

1. The animal/ticket is a \_\_\_\_\_. [zebra/mule]/[loser/winner]
2. Abigail and her son were \_\_\_\_\_. [visiting the zoo/driving past a farm/out shopping]
3. Abigail \_\_\_\_\_ that the animal/ticket is a zebra/loser. [knows/only believes]
4. How confident are you in the answer you just gave? [1: not at all confident—10: completely confident]
5. What did Abigail and her son see? [A black-and-white striped animal/A man with a super lotto ticket]
6. Abigail is \_\_\_\_\_ in thinking that the animal/ticket is a zebra/loser [justified/unjustified]
7. How confident are you in the answer you just gave? [1: not at all confident—10: completely confident]
8. Was there at least some chance, no matter how small, that the animal/ticket was a mule/winner?<sup>3</sup>

In this and all subsequent experiments, questions were always presented in the same order; response options were always rotated randomly, except for confidence Likert scales. The story remained at the top of the screen on all pages of the survey; participants could not go back to a previous page to alter an answer. Participants filled out a brief demographic survey after testing.

## 2.2. Results and discussion

Of chief interest was whether people would attribute knowledge at different rates across the three stories. If people judge that knowledge is less likely to result from probabilistic evidence than from perception, they should attribute knowledge less in Lotto than in Farm or Zoo. To evaluate this prediction, we examined responses to the knowledge question. Knowledge attributions varied across the Lotto (11% of responses), Zoo (78% of responses), and Farm (73% of responses) stories,  $\chi^2(2, N = 138) = 53.18, p < .001$ , Cramer's  $V = .621$ , all tests two-tailed. Participants attributed knowledge more in the Zoo than Lotto story,  $\chi^2(1, N = 93) = 43.13, p < .001$ , Cramer's  $V = .681$ , and more in the Farm than Lotto story,  $\chi^2(1, N = 92) = 37.27, p < .001$ , Cramer's  $V = .637$ ; knowledge attribution did not differ between the Zoo and Farm stories,  $\chi^2(1, N = 91) = .30, p = .583$ . Further tests examined whether, for each story, knowledge was attributed at rates exceeding chance. Participants mostly attributed knowledge in the Zoo story, Binomial sign test,  $N = 46, k = 10, p < .001$ , and in the Farm story,  $N = 45, k = 12, p = .002$ , but denied knowledge in the Lotto story,  $N = 47, k = 5, p < .001$ .

These findings suggest that people judge that knowledge is less likely to result from probabilistic evidence than from perception and, moreover, that people deny that knowl-

edge is gained from probabilistic evidence. However, the varying rates of knowledge attributions could perhaps have resulted because participants felt that the belief in the Lotto story was less justified than in the other stories, or more open to the chance of error. Broadly consistent with this, assessments of whether the belief was justified varied across the stories,  $\chi^2(2, N = 138) = 10.77, p = .005$ , Cramer's  $V = .279$ , as did assessments of whether there was a chance of error,  $\chi^2(2, N = 138) = 12.09, p = .002$ , Cramer's  $V = .296$ ; Table 1 shows the percentage of people affirming justification and chance of error in each condition in this experiment, as well as in those subsequent. To ensure that these differences in justification and chance-of-error judgments were not responsible for the differences in knowledge attributions across the stories, further analyses examined only those participants who ascribed justification and affirmed the chance of error ( $N = 30$  in Zoo, 32 in Farm, and 38 in Lotto). Limiting analysis to these participants, we continue to observe an effect of condition on knowledge ascription,  $\chi^2(2, N = 100) = 29.78, p < .001$ , Cramer's  $V = .546$ , with a lower rate of knowledge ascription in the Lotto than Zoo or Farm stories,  $ps < .001$ , and no difference in knowledge ascription between the Zoo and Farm stories,  $p = .533$ . Hence, these analyses support the conclusion that people judge that knowledge is less likely to result from probabilistic evidence than from perception.

Table 1

Percentage of participants in each experiment who attributed knowledge, affirmed the protagonist's belief was justified, and affirmed that the belief could have been false

Experiment	Knowledge	Justification	Chance of Error
Experiment 1			
Lotto	11	83	96
Zoo	78	96	71
Farm	73	100	70
Experiment 2A			
Odds	42	97	76
Perception	74	91	74
Experiment 2B			
Odds	14	86	96
Mafia	35	88	87
Experiment 3A			
Specific information	48	97	97
General information	27	98	100
Experiment 3B—Plant			
Specific information	68	100	94
General information	49	98	95
Experiment 3B—Bowl			
Specific information	46	100	90
General information	52	100	91
Experiment 3B—Beret			
Specific information	61	100	94
General information	42	98	95

### 3. Experiment 2A

The findings of the first experiment suggest that people are reluctant to ascribe knowledge for beliefs based on probabilistic evidence. The findings also show that this does not result because people view such beliefs as unjustified or because the beliefs leave open the chance of error. However, some skepticism about this interpretation is warranted because the Lotto scenario varied from the others in several respects, and not just in the kind of evidence supporting the protagonist's belief. Most notably, the scenarios differed in their "cover story" and in whether the protagonist's belief concerned the outcome of a lottery or the identity of an animal. Moreover, people might have denied knowledge in the Lotto scenario because they knew that one ticket would *definitely* win, and this might have ruled out the possibility of knowing that any particular ticket would lose. The Farm and Zoo scenarios did not include this feature (i.e., there was no reason to assume that one animal in the group was definitely a disguised mule). Hence, *these* differences, and not the kind of evidence supporting the belief, might be responsible for the differences across the stories in rates of knowledge attributions. To rule out this concern, the current experiment used scenarios which only varied in the kind of evidence available to the protagonist.

#### 3.1. Method

##### 3.1.1. Participants

Sixty-seven participants (aged 18–56,  $M = 31$  years; 24 female) were tested. Data were excluded from an additional participant who failed a comprehension question.

##### 3.1.2. Materials and procedure

Participants were randomly assigned to one of two conditions, Odds and Look. In each condition, participants read a story about a woman, Mary, who sees a dancing bear and wonders whether it is a real bear or a person in a bear suit. The scenarios differ in how Mary resolves her doubt. In the story for Odds, Mary considers the odds and concludes that it is a real bear. In the story for Looks, Mary looks again and concludes that it is a real bear. Here are the two stories, with differences between the stories indicated in square brackets.

A carnival is visiting Mary's town. To her surprise, she sees a man leading a dancing bear down the street. Mary wonders for a moment whether the bear is real, or whether perhaps it's just a person dressed in a bear suit. But then Mary [considers the odds / looks carefully] and concludes that it is a real bear. And Mary is right: it is a real bear.

After reading the story, participants answered questions similar to those in Experiment 1 (see Supporting Information for exact wordings). These included a knowledge question asking whether Mary "knows" or "only believes" it is a real bear, a justification question, and a chance-of-error question.

### 3.2. Results and discussion

Ascriptions of knowledge were higher when the protagonist's belief was based on perception (74% of responses) than when based on reasoning about odds (42% of responses),  $\chi^2(1, N = 67) = 6.66, p = .010$ , Cramer's  $V = .315$ . Binomial sign tests revealed that while participants mostly attributed knowledge for beliefs based on perception,  $N = 34, k = 9, p = .009$ , they attributed knowledge at chance levels for beliefs based on reasoning about odds,  $N = 33, k = 14, p = .487$ , n.s.

The stories did not yield varying judgments regarding whether the belief was justified,  $\chi^2(1, N = 67) = 1.00, p = .317$ , n.s., nor in whether the belief could have been incorrect,  $\chi^2(1, N = 67) = 0.04, p = .834$ , n.s. Even so, a follow-up analysis confirmed that with analysis limited to participants who ascribed justification and affirmed the chance of error ( $N = 22$  in Looks, and 24 in Odds), knowledge was more likely to be ascribed for beliefs based on perception than those based on reasoning about odds,  $\chi^2(1, N = 46) = 5.58, p = .018$ , Cramer's  $V = .348$ .

These findings are consistent with those from Experiment 1 in suggesting that participants judge that knowledge is less likely to result from reasoning about odds than to result from perception. Moreover, the findings reveal this difference using closely matched scenarios, which feature the same cover story, and in which the protagonist has the same belief. The scenarios differed only in whether the protagonist resolved her uncertainty by taking a careful look at it or by considering the odds that this belief could be false. However, the stories may have varied in the amount of information possessed by the protagonist—taking a *further* look at the bear might have provided the protagonist with new information, whereas considering the odds might not have this effect. Perhaps this informational difference influenced attributions of knowledge. To follow up on this possibility, the next experiment used scenarios that did not differ in the information available to the protagonist; the scenarios differ only in the information she bases her belief on.

## 4. Experiment 2B

This experiment follows up on the possibility that informational difference across stories influenced attributions of knowledge in Experiment 2A. To remove this difference, this experiment used scenarios that did not differ in the information available to the protagonist. Instead, the scenarios differ only in the information she bases her belief on.

### 4.1. Method

#### 4.1.1. Participants

One hundred fifty-five participants (aged 18–61,  $M = 29$  years; 59 female) were tested. Data were excluded from an additional 10 participants who failed comprehension questions.

#### 4.1.2. *Materials and procedure*

Participants were randomly assigned to one of two conditions, Odds and Mafia. In each condition, participants read a story about Ellen, who purchases a lottery ticket, but later learns that the mafia has rigged the lottery so that anyone not in the mafia has only a 1-in-10,000,000 chance of winning. In the Odds condition, she bases her belief on the fact that the odds are so poor. In the Mafia condition, she bases her belief on the fact that the Mafia rigged the lottery. Here are the two stories, with differences between the stories indicated in square brackets:

Ellen bought a ticket in this week's State Lottery. She wasn't able to watch the evening news where they reported which number won. But she does watch a special report that reveals that the State Lottery is rigged by members of the local mafia, so that there is only a 1-in-10,000,000 (one-in-ten-million) chance that anyone not in the mafia will win. Because she learns about the [1-in-10,000,000 odds / Mafia rigging the lottery], Ellen concludes that her ticket lost. And she is right: her ticket lost.

After reading the story, participants answered questions similar to those in the previous experiments (see Supporting Information for exact wordings). These included a knowledge question, a justification question, and a chance-of-error question.

#### 4.2. *Results and discussion*

Ascriptions of knowledge were higher when the protagonist's belief was based on the Mafia's involvement (35% of responses) than when based on the probabilistic evidence (14% of responses),  $\chi^2(1, N = 155) = 9.32, p = .002$ , Cramer's  $V = .245$ . Binomial sign tests revealed that participants denied knowledge in the Odds story,  $N = 80, k = 11, p < .001$ , and in the Mafia story,  $N = 75, k = 26, p = .011$ .

Judgments regarding whether the belief was justified did not vary across the stories,  $\chi^2(1, N = 155) = 0.11, p = .745$ , n.s.; however, participants were more likely to affirm the chance of error in the Odds story than in the Mafia story,  $\chi^2(1, N = 155) = 4.63, p = .042$ , Cramer's  $V = .173$ . A follow-up analysis examined participants who ascribed justification and affirmed the chance of error ( $N = 66$  in Odds, and 58 in Mafia) and found that these participants attributed knowledge more in the Mafia story than in the Odds story,  $\chi^2(1, N = 124) = 7.49, p = .006$ , Cramer's  $V = .246$ .

These findings provide further evidence that people are less likely to ascribe knowledge for belief based on probabilistic evidence compared with those based on other forms of evidence. The experiment also shows that visual perception is not the only form of evidence to result in greater knowledge ascriptions than probabilistic evidence; in the Mafia story, the protagonist based her belief on testimony about a *causal factor*, making it likely that her ticket would be a loser (i.e., the Mafia having rigged it; Harman, 1968; see the General Discussion for further discussion of causal information). The demonstration in the current experiment is dramatic because the protagonist had access to the same information and evidence in both stories; the stories only differed in

which aspect of the evidence led her to form her belief. However, it is possible that one contributing factor to this difference is that participants thought that the information about the Mafia was less expected, and hence more informative, than the information about the long odds.

Why, then, are people reluctant to ascribe knowledge for beliefs based on probabilistic evidence? The findings from the previous experiments allow us to rule out a number of explanations. It cannot be that people deny knowledge because such beliefs are unjustified—as we have seen, people affirm that beliefs based on probabilistic evidence *are* justified. Nor can it be that knowledge is denied because these beliefs leave a chance of error, failing to guarantee absolute certainty. This possibility is ruled out because people also acknowledge the chance of error when considering beliefs based on perceptual evidence, for which people do ascribe knowledge. Moreover, reluctance to ascribe knowledge for beliefs based on probabilistic evidence, relative to beliefs based on other forms of evidence, remains when considering only participants who both view the protagonist's belief as justified and affirm that there was a chance of error. Hence, some other factor (or factors) must explain people's reluctance to ascribe knowledge for beliefs based on probabilistic evidence. In the next experiment, we provide evidence that one relevant factor is the *specificity* of evidence.

## 5. Experiment 3A

Using probabilistic evidence often involves taking an “outside view” and drawing a conclusion about a particular individual from general evidence that concerns a distribution. Such judgments are contrasted by those taking an “inside view,” where people consider factors *specific* to a particular case (Kahneman & Tversky, 1982; Lagnado & Sloman, 2004; also see Griffin et al., 2007). Perception and testimony often provide information specific to a particular case—for example, information specific to a particular zebra is provided by perceiving it, or by hearing testimony about it. And in some cases, probabilistic evidence is specific to a particular case, as occurs when one considers causal information making a specific outcome very likely.

As noted in the Introduction, people may prefer taking the “inside view” over the “outside view” and might likewise neglect to draw inferences about particular cases from distributional evidence. Such tendencies might contribute to people's reluctance to ascribe knowledge. Broadly consistent with this, a study examining simulated courtroom decision-making found that statistical DNA evidence is more persuasive when framed to induce reasoning about a particular case in isolation (inside view) than when framed to view the case relative to a distribution (outside view; Koehler, 2001; also see Arkes et al., 2012). Likewise, people might be more likely to ascribe knowledge when an agent's belief is based on probabilistic evidence that concerns the specific case than on probabilistic evidence that concerns a distribution.

## 5.1. Method

### 5.1.1. Participants

One hundred and thirteen participants (aged 18–66,  $M = 34$  years; 53 female) were tested. Data were excluded from an additional 25 participants who failed comprehension questions.

### 5.1.2. Materials and procedure

Participants read a story about Luigi, who uses an electronic device to find out whether his plot of land is “terrigenic.” Participants were randomly assigned to one of two conditions, which differed in whether the device tested soil samples (specific information) or listed regional information about soil (general information). Crucially, in both stories, the device provided probabilistic evidence about whether the soil is terrigenic. As such, the conditions hold probabilistic evidence constant. Here are the two stories, with differences between the stories indicated in square brackets:

Luigi purchased a plot of land in Springville. He has an electronic device that [tests/ lists] the likelihood that soil is terrigenic. The device says there is only a 1-in-10-million chance that [Luigi’s soil/soil in Springfield] is terrigenic. So Luigi concludes that his soil is non-terrigenic. And he is right: it is non-terrigenic.

After reading the story, participants answered questions similar to those in the previous experiments (see Supporting Information for exact wordings), again including a knowledge question, justification question, and chance-of-error question.

## 5.2. Results and discussion

Ascriptions of knowledge were higher when the protagonist’s belief was based on information from the soil test (48% of responses) than when based on the list (27% of responses),  $\chi^2(1, N = 113) = 5.60, p = .018$ , Cramer’s  $V = .223$ . Binomial sign tests revealed that knowledge was attributed at chance rates when the belief was based on the soil test,  $N = 64, k = 31, p = .901$ , n.s., and that participants denied knowledge when the belief was based on the list,  $N = 49, k = 13, p = .001$ .

Judgments regarding whether the belief was justified did not vary across the stories,  $\chi^2(1, N = 113) = 0.13, p = .722$ , n.s., nor did affirmations of chance of error,  $\chi^2(1, N = 113) = 1.56, p = .212$ , n.s. Nonetheless, a follow-up analysis examined participants who ascribed justification and affirmed the chance of error ( $N = 60$  in specific-test, and 48 in general-list), and found that these participants attributed knowledge more when the belief was based on specific information than general information,  $\chi^2(1, N = 124) = 7.49, p = .006$ , Cramer’s  $V = .246$ . Hence, these findings suggest that people’s reluctance to infer knowledge from probabilistic evidence is greater when this requires drawing a conclusion about a particular individual from evidence that concerns a population (outside

view); when the evidence is specific to the particular case (inside view), people are more willing to attribute knowledge.

## 6. Experiment 3B

The previous experiment demonstrated that people are more likely to ascribe knowledge for beliefs based on probabilistic information when this information is specific to particular case (i.e., rather than when this information concerns a distribution). Because this was only demonstrated using one cover story, we next attempted to replicate this finding with other cover stories. Also, although we showed that ascriptions of knowledge could be increased, participants did not ascribe knowledge at rates exceeding chance. Hence, the findings do not show that participants are (on the whole) willing to ascribe knowledge for beliefs based on specific probabilistic evidence. In this experiment, we used cover stories intended to increase ascriptions of knowledge. In contrast to Experiment 3A, which used a cover story about an improbable outcome (i.e., 1-in-10-million chance that soil is terragenic), we instead used stories in which the outcome was probable; findings from pilot studies, conducted as part of a different series of experiments, suggested this change might lead to increased ascriptions of knowledge.

### 6.1. Method

#### 6.1.1. Participants

Three hundred and sixty-five participants (aged 18–72,  $M = 32$  years; 150 female) were tested. Data were excluded from an additional 27 participants who failed comprehension questions.

#### 6.1.2. Materials and procedure

Participants read one of three cover stories: In “plant,” Bob concludes that his spider plant contains the chemical *aracnium*; in “bowl,” Samantha concludes an ancient bowl is Minoan; and in “beret,” Janice wonders concludes that her French beret contains *lidinum* fibers. Participants reading each cover story were randomly assigned to one of two conditions, which differed in whether the protagonist drew a conclusion based on evidence specific to the particular item or whether the agent’s information was based on information about a larger population. In all cases, the evidence was probabilistic and provided 99% certainty that the conclusion was true. Here are the three cover stories, with differences between the versions in square brackets:

*Plant:* Bob wonders if his spider plant contains the chemical *aracnium*. He [conducts a test on the spider plant/consults a book on spider plants]. The [test shows/books says] that it is 99% likely that the spider plant contains *aracnium*. So Bob concludes that his spider plant contains *aracnium*. And he is right: it contains *aracnium*.

*Bowl:* Samantha is examining an ancient bowl with distinctive engravings on its handle. Experts have found that [there is a 99% chance that this bowl is Minoan/99% of very similar bowls are Minoan]. So Samantha concludes the bowl is Minoan. And she is right: It is Minoan.

*Beret:* Janice wonders whether her French beret contains lidinium fibres. She conducts [a test on the French beret/research on French berets]. The [test/research] shows that it is 99% likely for [this French beret/French berets] to contain lidinium fibres. So Janice concludes that the French beret contains lidinium fibres. And she is right: it contains lidinium fibres.

After reading the story, participants answered questions similar to those in the previous experiments (see Supporting Information for exact wordings), again including a knowledge question, justification question, and chance-of-error question. In contrast to the previous experiments, participants were not asked to rate their confidence in their answers.

## 6.2. Results and discussion

In the Plant and Beret cover stories, ascriptions of knowledge were higher when the protagonist's belief was based on information specific to the items (68% of responses in Plant, 61% in Beret) than when based on the population to which each item belonged (49% of responses in Plant, 42% in Beret),  $\chi^2_{\text{Plant}}(1, N = 126) = 4.45, p = .035$ , Cramer's  $V = .188$ ;  $\chi^2_{\text{Beret}}(1, N = 124) = 4.61, p = .032$ , Cramer's  $V = .193$ . Binomial sign tests revealed that in the Plant cover story knowledge was attributed at rates exceeding chance when the belief was based on specific information,  $N = 65, k = 21, p = .006$ , and at chance levels when the belief was based on more general information,  $N = 61, k = 30, p = 1.000$ , n.s.; in the Beret cover story, ascriptions did not depart from chance in either condition,  $ps > .100$ . In contrast, in the Bowl cover story, ascriptions of knowledge did not vary depending on whether the information was specific (46% of responses) or general (52% of responses),  $\chi^2(1, N = 115) = 0.42, p = .518$ , n.s. Binomial sign tests revealed that knowledge was attributed at chance rates in both conditions,  $ps > .600$ .

As can be seen in Table 1, participants judged that the protagonist's belief was justified, and affirmed that there was a chance of error, in both conditions in all cover stories. In no case did these judgments vary across the conditions, and follow-up analyses that included only participants who ascribed justification and affirmed the chance of error revealed the same general pattern of findings,  $\chi^2_{\text{Plant}}(1, N = 118) = 3.23, p = .071$ , Cramer's  $V = .166$ ;  $\chi^2_{\text{Bowl}}(1, N = 104) = 0.592, p = .442$ , n.s.;  $\chi^2_{\text{Beret}}(1, N = 116) = 5.81, p = .016$ , Cramer's  $V = .22$ .

In sum, findings from two of the three cover stories provide further evidence that people are more willing to ascribe knowledge to beliefs based on probabilistic evidence when the evidence is specific to the case considered, rather than when it is general and requires drawing a conclusion about a specific case from information about a population. Moreover, the findings reveal that participants sometimes ascribe knowledge at rates exceeding chance for beliefs based on specific evidence.

We are uncertain why a difference between conditions was not observed for the Bowl cover story. One possibility is that performance was impacted because the protagonist in the “specific information” version of the story had relatively indirect evidence, and the nature of the evidence was vague. Rather than basing her judgment on observations or direct tests on the bowl, she based her conclusion on the opinion of experts who had examined the bowl, and it was not specified what caused the experts to conclude that the bowl is probably Minoan. Perhaps findings would have differed if the protagonist had based her conclusion on her own examination of the bowl, or if the story had specified *how* examining the bowl led experts to conclude that it is Minoan.

It is also worth noting a potential concern with the Plant cover story. The conditions in that story varied not only in whether the probabilistic evidence was specific to the case considered, but also in whether the evidence came from a test or a book. This difference could have affected responses, because participants might the test is a more reliable source of information than the book, and source reliability impacts how people treat evidence (e.g., Hahn, Harris, & Corner, 2009). However, one reason to set aside this concern is that a similar difference in responses emerged in the Beret cover story which did not feature this same contrast in sources of evidence.

## 7. General discussion

The findings of our experiments suggest that people are often reluctant to ascribe knowledge for beliefs based on probabilistic evidence. Participants ascribed knowledge less for such beliefs than beliefs based on perceptual evidence or beliefs based on causally relevant testimony. This was even true when people considered agents that had access to both probabilistic and non-probabilistic evidence. Hence, it appears that knowledge ascriptions depend on the evidence used to form a belief, not on the evidence available to an agent.

People’s denial of knowledge for beliefs based on probabilistic evidence does not arise because such beliefs are deemed unjustified, nor because reliance on probabilistic evidence leaves open the possibility that the agent could have been led to a false belief. These explanations for the findings are ruled out because people affirmed that beliefs based on probabilistic evidence are justified, and they also acknowledged the fallibility of beliefs based on other forms of evidence, which were deemed to be instances of knowledge. As such, the findings undermine certain philosophical accounts of why agents lack knowledge in “lottery cases,” at least if it is believed that folk ascriptions bear on these accounts (Lewis, 1996; Nelkin, 2000).

It might seem puzzling that judgments about justification do not show the same pattern as judgments about knowledge do. But theories of justification are compatible with this asymmetry. For example, one approach identifies justification with a sort of *blamelessness* (Alston, 1989; Goldman, 1988: essays 4 and 5; Plantinga, 1993). On this approach, believing based on overwhelming “outside” odds renders one blameless, which would explain why such beliefs are viewed as justified. However, on this same approach,

knowledge requires more than blamelessly believing the truth. It also requires one's evidence to be properly connected to the state of affairs that makes one's belief true, and people might doubt that even overwhelming "outside" odds constitute an appropriate connection. Similarly, another approach in epistemology identifies justification with *skillful* or *competent belief-formation* (Greco, 1993; Sosa, 1991; Turri, 2012; Zagzebski, 1996). On this approach, believing in accordance with overwhelming odds is competent, which would explain why such beliefs are viewed as justified. However, on this same approach, knowledge requires more than competent belief-formation. Instead, it requires that the competent belief-formation count as the detection of the truth (Turri, in press b), and people might judge that although believing based on "outside" odds can provide true beliefs, it does not count as a way of detecting the truth.

This asymmetry between knowledge judgments and justification judgments raises a broader question about the relation between ascriptions of knowledge and other judgments we can make about others' beliefs. As mentioned in the Introduction, knowledge attributions may be especially important to our habits of flagging good sources of information and deciding whom to trust. However, knowledge attributions might not be privileged in this regard, and perhaps trust, and other epistemic decisions, depends more on other judgments, such as judgments about whether an agent's belief is justified, or judgments about our own confidence that the agent's belief is true. Or perhaps these different judgments are all important in trust and epistemic decision-making but serve different functions.

### 7.1. *Two kinds of probabilistic evidence*

We also found that knowledge ascriptions vary depending on the kind of probabilistic evidence used by an agent, and in keeping with Kahneman and Tversky's (1982) distinction between probabilistic evidence that provides an "outside" or "inside view." People were very reluctant to ascribe knowledge for conclusions about a particular case that rested on probabilistic evidence about a larger set of related cases ("outside view"); but people were more willing to ascribe knowledge when the conclusion was based on probabilistic evidence concerning the case specifically ("inside view").

As noted in the Introduction, many previous studies have also found that people base judgments on evidence specific to a particular case, while neglecting evidence that concerns a larger distribution (e.g., this tendency characterizes base-rate neglect). In most studies, participants considered *both* kinds of information together and made numerical judgments (e.g., judging the probability that a certain outcome will occur). Many factors might lead people to under-utilize distributional evidence in such studies. For example, people might neglect distributional evidence because of *interference* from specific evidence; they might have difficulty seeing the nested relation between the two kinds of evidence (e.g., Barbey & Sloman, 2007); and they might have difficulty using values given in probabilistic format (rather than natural frequency format) to make numerical judgments (e.g., Gigerenzer & Hoffrage, 1995). However, these factors cannot underlie our findings because they require that people consider both kinds of probabilistic information *together* or that people make numerical judgments—conditions that did not arise in our

studies. Hence, our findings differ from most previous studies in showing that people neglect evidence about distributions even when it is the only evidence available and when using it would be trivially easy (also see Arkes et al., 2012; Koehler, 2001; Wells, 1992). As such, the findings are not only informative about people's attributions of knowledge, but also extend knowledge of the ways that evidence from different sources is treated differently in decision making.

Why then do people neglect distributional evidence in ascribing knowledge? And why are they more willing to ascribe knowledge for probabilistic evidence that was specific to a particular case? Although our experiments do not allow us to be sure of an explanation, this difference might arise because specific evidence is viewed as relevant, or causally linked, to the protagonist's conclusion, while evidence about a distribution typically lacks such a connection. Similar explanations focusing on causal relevance, or relevance more broadly, have also been proposed to explain base-rate neglect (e.g., Bar-Hillel, 1980; Krynski & Tenenbaum, 2007; Sloman, 2003; Tversky & Kahneman, 1982; also see Taylor, Landy, & Ross, 2012; for more general effects of providing causal explanations). In the next section we spell out a variant of this account specific to ascriptions of knowledge.<sup>4</sup>

## 7.2. *Implications for accounts of knowledge ascription*

Two conflicting conclusions could be drawn from the finding that ascriptions of knowledge vary for different kinds of probabilistic evidence. First, the finding might be interpreted as contributing to a concise account of the principles underlying knowledge ascription. For instance, people might adhere to the principle that knowledge occurs only when the evidence for an agent's belief is caused by the fact that makes the agent's belief true. On such an account, an agent will know that a zebra is in the field only if *the fact that a zebra is in the field* causes the agent to have the evidence that she has (Goldman, 1967). This account predicts that people will deny knowledge in cases of "apparent evidence" (Starmans & Friedman, 2012). And it also explains why participants deny knowledge when probabilistic evidence concerns a set of cases. Consider the agent who believes his soil is not terragenic because almost none of the soil in the region is terragenic. The evidence for his belief does not stem from the fact that actually makes this belief true (i.e., his soil actually having this property). Contrast this with the case where he relies on a test of his own soil—now the evidence for his belief does stem from the fact that makes the belief true.

So perhaps a simple principled account of knowledge attribution is possible. However, other interpretations are possible. The findings suggest that in attributing knowledge to others, people are sensitive to a distinction between two kinds of evidence previously discussed in research on judgment and decision making: the distinction between evidence specific to a particular case and evidence gained by considering a larger set of cases (Griffin et al., 2007; Kahneman & Tversky, 1982; Lagnado & Sloman, 2004). This raises the possibility that people's knowledge attributions will be equally sensitive to other distinctions raised in this literature, and to other manipulations known to influence how people weigh evidence. On this view, rather than leading to a simple principled account of knowledge attribution, the current findings hint that people consider numerous factors in ascribing knowledge.

## Acknowledgments

OF's contributions to this research were supported by the Social Sciences and Humanities Research Council of Canada; JT's contributions to this research were supported by the Social Sciences and Humanities Research Council of Canada and an Early Researcher Award from the Ontario Ministry of Economic Development and Innovation.

## Notes

1. We focus on recent findings with adult participants. However, much research in developmental psychology has also examined how factors affecting how children ascribe knowledge (e.g., Keenan, Ruffman, & Olson, 1994; Nilsen, Graham, Smith, & Chambers, 2008; Pratt & Bryant, 1990; Sodian, 1988).
2. As shown by Starmans and Friedman (2012) and replicated in subsequent studies (e.g., Turri, 2013a). This has inspired further experimental work on the relationship between knowledge and luck in a genre of thought experiments known as "Gettier cases." In conflict with claims by epistemologists about what is intuitively obvious about such cases, it turns out that many people attribute knowledge in some Gettier cases (Nagel, Juan, & Mar, 2013; Starmans & Friedman, 2012; Turri, Buckwalter, & Blouw, in press).
3. Confidence ratings were included to leave open the possibility of deriving weighted knowledge and justification scores, by combining responses on the knowledge and justification questions with their corresponding confidence ratings. However, in all studies the same patterns of findings arose regardless of whether we analyzed weighted scores or responses to the knowledge and justification questions alone. Hence, for simplicity, we omit analysis of weighted scores.
4. A second possibility (pointed out by an anonymous reviewer) is that participants might infer that the protagonist has more information when described as having evidence specific to a case, and this might increase knowledge ascriptions. For example, when Bob's spider plant is 99% likely to contain aracunium, it might be assumed that spider plants in general are less likely to contain aracunium. If Bob is aware of this, then his conclusion would be based both on information specific to his spider plant and on information about spider plants in general. In contrast, Bob has only one kind of evidence when he learns that spider plants generally are 99% likely to contain aracunium. Although our findings do not rule this possibility out, the null finding in the Minoan bowl story (Experiment 3B) provide some reason to doubt it, as this account predicts that responses with that story should have varied by condition. In contrast, the causal relevance account might explain the null effect, as the specific information in that condition likely lacked causal relevance (i.e., the scenario did not explain what caused experts to conclude the bowl was probably Minoan).

## References

- Alston, W. P. (1989). *Epistemic justification: Essays in the theory of knowledge*. Ithaca, NY: Cornell University Press.
- Arkes, H. R., Shoots-Reinhard, B., & Mayes, R. S. (2012). Disjunction between probability and verdict in juror decision making. *Journal of Behavioral Decision Making*, 25, 276–294.
- Barbey, A. K., & Sloman, S. A. (2007). Base-rate respect: From ecological rationality to dual processes. *Behavioral and Brain Sciences*, 30, 241–254.
- Bar-Hillel, M. (1980). The base-rate fallacy in probability judgments. *Acta Psychologica*, 44, 211–233.
- Beebe, J. R., & Buckwalter, W. (2010). The epistemic side-effect effect. *Mind & Language*, 25, 1–25.
- Buckwalter, W. (in press). Factive verbs and protagonist projection. *Episteme*.
- Chisholm, R. (1989). *Theory of knowledge* (3rd ed.). Englewood Cliffs, NJ: Prentice Hall.
- Cohen, L. J. (1981). Subjective probability and the paradox of the gatecrasher. *Arizona State Law Journal*, 52, 627–634.
- Craig, E. (1990). *Knowledge and the state of nature: An essay in conceptual synthesis*. Oxford, England: Oxford University Press.
- Descartes, R. (1641 [1993]). *Meditations on first philosophy*. (D. Cress, trans. Indianapolis: Hackett.
- Dretske, F. (1970). Epistemic operators. *The Journal of Philosophy*, 67, 1007–1023.
- Dretske, F. (1981). *Knowledge and the flow of information*. Cambridge, MA: MIT Press.
- Fantl, J., & McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford, England: Oxford University Press.
- Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis*, 23, 121–123.
- Gigerenzer, G., & Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review*, 102, 684–704.
- Goldman, A. (1967). A causal theory of knowing. *The Journal of Philosophy*, 64, 357–372.
- Goldman, A. I. (1988). Strong and weak justification. *Philosophical Perspectives*, 2, 51–69.
- Greco, J. (1993). Virtues and vices of virtue epistemology. *JSTOR: Canadian Journal of Philosophy*, 23(3), 413–432.
- Griffin, D., Koehler, D. J., & Brenner, L. (2007). Frequency formats are a small part of the base rate story. *Behavioral and Brain Sciences*, 30, 268–269.
- Hahn, U., Harris, A. J. L., & Corner, A. (2009). Argument content and argument source: An exploration. *Informal Logic*, 29, 337–367.
- Harman, G. (1968). Knowledge, inference, and explanation. *American Philosophical Quarterly*, 5, 164–173.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford, England: Oxford University Press.
- Kahneman, D., & Tversky, A. (1982). Variants of uncertainty. *Cognition*, 11, 143–157.
- Keenan, T., Ruffman, T., & Olson, D. R. (1994). When do children begin to understand logical inference as a source of knowledge? *Cognitive Development*, 9, 331–353.
- Koehler, J. J. (2001). When are people persuaded by DNA match statistics? *Law and Human Behavior*, 25, 493–513.
- Krynski, T. R., & Tenenbaum, J. B. (2007). The role of causality in judgment under uncertainty. *Journal of Experimental Psychology: General*, 136, 430–450.
- Lagnado, D. A., & Sloman, S. A. (2004). Inside and outside probability judgment. In D. J. Koehler & N. Harvey (Eds.), *Blackwell handbook of judgment and decision making* (pp. 157–176). Malden, MA: Blackwell.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74, 549–567.
- Locke, J. (1975 [1690]). *An essay concerning human understanding*. Edited with an Introduction by Peter H. Nidditch. Oxford, England: Clarendon Press.
- Nagel, J., Juan, V. S., & Mar, R. A. (2013). Lay denial of knowledge for justified true beliefs. *Cognition*, 129, 652–661.
- Nelkin, D. (2000). The lottery paradox, knowledge, and rationality. *Philosophical Review*, 109, 373–409.
- Neta, R. (2011). A refutation of cartesian fallibilism. *Nous*, 45, 658–695.

- Nilsen, E. S., Graham, S. A., Smith, S., & Chambers, C. G. (2008). Preschoolers' sensitivity to referential ambiguity: Evidence for a dissociation between implicit understanding and explicit behavior. *Developmental Science*, *11*, 556–562.
- Nozick, R. (1981). *Philosophical explanations*. Cambridge, MA: Harvard University Press.
- Plantinga, A. (1993). *Warrant: the current debate*. Oxford, England: Oxford University Press.
- Plato. (1990 [370 BCE]). *Theaetetus*. (M. J. Levett & M. F. Burnyeat, Trans.). Indianapolis, IN: Hackett.
- 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*, 973–982.
- Schaffer, J., & Knobe, J. (2012). Contrastive knowledge surveyed. *Nous*, *46*, 675–708.
- Sloman, S. A. (2003). Causal models can be used to predict base-rate neglect. In R. Alterman & D. Hirsh (Eds.), *Proceedings of the twenty-fifth annual conference of the cognitive science society*, Volume 1 (pp. 1088–1093). Hillsdale, NJ: Erlbaum.
- Sodian, B. (1988). Children's attributions of knowledge to the listener in a referential communication task. *Child Development*, *59*, 378–385.
- Sosa, E. (1991). *Knowledge in perspective*. Cambridge, England: Cambridge University Press.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford, England: Oxford University Press.
- Starmans, C., & Friedman, O. (2012). The folk conception of knowledge. *Cognition*, *124*, 272–283.
- Sutton, J. (2007). *Without justification*. Cambridge, MA: MIT Press.
- Taylor, E. G., Landy, D. H., & Ross, B. H. (2012). The effect of explanation in simple binary decision tasks. *The Quarterly Journal of Experimental Psychology*, *65*, 1361–1375.
- Turri, J. (2011a). The express knowledge account of assertion. *Australasian Journal of Philosophy*, *89*, 37–45.
- Turri, J. (2011b). Contingent a priori knowledge. *Philosophy and Phenomenological Research*, *80*, 312–326.
- Turri, J. (2012). Is knowledge justified true belief? *Synthese*, *184*(3), 247–259. doi:10.1007/s11229-010-9773-8.
- Turri, J. (2013a). A conspicuous art: Putting Gettier to the test. *Philosophers' Imprint*, *13*, 1–16.
- Turri, J. (2013b). The test of truth: An experimental investigation of the norm of assertion. *Cognition*, *129*, 279–291.
- Turri, J. (in press a). Unreliable knowledge. *Philosophy and Phenomenological Research*.
- Turri, J. (in press b). Epistemic situationism and cognitive ability. In M. Alfano & A. Fairweather (Eds.), *Epistemic situationism*. Oxford, England: Oxford University Press.
- Turri, J., Buckwalter, W., & Blouw, P. (in press). Knowledge and luck. *Psychonomic Bulletin & Review*.
- Tversky, A., & Kahneman, D. (1982). Evidential impact of base rates. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases*. Cambridge, England: Cambridge University Press.
- Vogel, J. (1990). Are there counterexamples to the closure principle? In M. D. Roth & G. Ross (Eds.), *Doubting* (pp. 13–27). Dordrecht, the Netherlands: Kluwer Academic Publishers.
- Wells, G. L. (1992). Naked statistical evidence of liability: Is subjective probability enough? *Journal of Personality and Social Psychology*, *62*, 739–752.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford, England: Oxford University Press.
- Wright, J. C. (2010). On intuitional stability: The clear, the strong, and the paradigmatic. *Cognition*, *115*, 491–503.
- Zagzebski, L. T. (1996). *Virtues of the mind: An inquiry into the nature of virtue and the ethical foundations of knowledge*. Cambridge, England: Cambridge University Press.

### Supporting Information

Additional Supporting Information may be found in the online version of this article:

**Data S1.** Materials used in experiments.