1

Subtle Syntactic Cues Affect Intuitions about Knowledge

Methodological and Theoretical Implications for Epistemology

Zachary Horne, Arizona State University
Andrei Cimpian, New York University

Epistemology is a central topic in philosophy concerned with the nature and limits of knowledge. Epistemologists seek to answer questions about what knowledge is and whether knowledge is possible at all, as well as examine the implications of these questions for ethics, the law, and philosophy of science. Epistemologists often approach these questions by analyzing whether a given epistemic theory leads to intuitive answers on various thought experiments—hypothetical situations taken to reveal underlying facts about knowledge.

A core assumption shared by most epistemologists is that the inferences licensed by statements regarding a subject's *knowledge* (noun form) and a subject's *knowing* (verb form) are interchangeable—these syntactic forms are assumed to yield identical intuitions and thus to bear in identical ways on a given epistemic theory (e.g., BonJour, 1985; DeRose, 1992; Gettier, 1963; Hawthorne, 2004; Lewis, 1996; Stanley, 2005; Williamson, 2002). We denote this the *epistemic interchangeability assumption*. While philosophers have questioned many other fundamental epistemic assumptions—for instance, the assumption that belief is required for knowledge (e.g., Murray, Sytsma, & Livengood, 2013), or that justification, truth, and belief are jointly sufficient to account for our

intuitions about knowledge (Gettier, 1963)—the epistemic interchangeability assumption has gone unchallenged.

The psychological literature, however, suggests some reasons to be suspicious of this assumption. Many studies have found that subtle syntactic differences analogous to those between "know" (a verb) and "knowledge" (a noun) are accompanied by often substantial differences in inferences and behavior (e.g., Gelman & Heyman, 1999; Markman, 1991; Walton & Banaji, 2004). In other words, intuitions about what people "know" may be subtly but reliably different from intuitions about the "knowledge" that people have, with obvious consequences for epistemological arguments that use these syntactic frames interchangeably. Consistent with this possibility, here we present the results of six studies that raise questions about the epistemic interchangeability thesis. Our findings suggest that the epistemic state denoted by "knowledge" (noun) has stronger evidential requirements than the epistemic state denoted by "know" (verb) and its variants; the intuitions elicited via these two syntactic frames appear to diverge. Since many epistemic theories are built on linguistic examples that do not take this syntactic dimension into account, the empirical evidence presented here has significant implications for methods and theory in epistemology.

The chapter will proceed as follows: First, we illustrate the claim that linguistic intuitions about "knowing" and "knowledge" are used interchangeably in philosophical theorizing. Second, we motivate the worry that "knowing" and "knowledge" lead to different epistemic inferences by turning to psychological research on the influence of syntactic cues on people's judgments and behavior. Third, we consider how our argument fits into the recent movement within experimental philosophy that seeks to examine critically the reliability of philosophical intuitions. Finally, we show the effects of syntactic cues on epistemic inferences in six studies that examined both philosophers' intuitions (as reflected in their writings; Study 1) and non-philosophers' judgments (Studies 2–6) about these matters.

1. Sample Cases of the Epistemic Interchangeability Assumption

As mentioned above, epistemologists consider hypothetical cases to determine the parameters that affect the attribution of knowledge.

This approach to understanding knowledge proceeds by considering whether a subject "knows" that P or has "knowledge" that P in such a hypothetical case, without attention to the differences in the epistemic intuitions elicited by these syntactic frames (i.e., the epistemic interchangeability assumption). We illustrate this approach with two well-known examples.

1.1. Example #1: knowledge is sensitive to stakes

Ongoing research in epistemology suggests that factors besides justification, truth, and belief may affect people's attributions of knowledge (see Buckwalter, 2012, for a review). Consider, for example, the thesis that knowledge in *high-stakes* situations—situations in which the practical consequences of being wrong are high—is more difficult to attain than knowledge in *low-stakes* situations—situations in which the practical consequences of error are slight (see Pinillos, 2011, for a review). This thesis is partially motivated by a linguistic intuition about the application of the verb "know." For example, intuitively, we are less likely to say that Hannah "knows" that her car insurance is current if she is about to get into a car accident than if she rarely drives her car at all. Thus, on the basis of a linguistic intuition about the application of the verb "know," philosophers have argued that the concept denoted by "knowledge" is sensitive to the stakes of a situation (illustrating the epistemic interchangeability assumption).

1.2. Example #2: knowledge requires certainty

The claim that knowledge sometimes requires certainty is similarly motivated by a linguistic intuition about "know" (Hawthorne, 2004). For example, imagine that Emma has a 1 in 10 million chance of winning the lottery. Although her chances are extremely low, intuitively we cannot say that Emma "knows" she will not win the lottery until the winning numbers are announced (and they're not hers). This is so even though it is extremely unlikely she will win. Contrast this with a situation in which Emma does not enter the lottery at all—in this case, it is obvious Emma "knows" that she will not win before the winning numbers are announced. On the basis of these and similar ordinary-language intuitions, philosophers have argued that the epistemic state denoted by "knowledge" sometimes requires certainty. Here, too, the assumption is that the epistemic inferences licensed by the verb "know" apply automatically to the concept denoted by "knowledge."

These examples illustrate the common practice of arguing for or against epistemic theories on the basis of linguistic evidence that treats intuitions about "knowing" and "knowledge" as interchangeable. Next, we call this assumption into question by describing psychological evidence that subtle syntactic cues can affect our judgments.

2. Syntactic Cues Affect Inferences and Behavior

Research in developmental, cognitive, and social psychology provides substantial evidence for the claim that subtle syntactic differences can affect people's inferences and behavior. In a classic study, Gelman and Heyman (1999) investigated how linguistic labels affect children's inferences about unfamiliar agents. In their study, they either described an agent to young children using a novel category label, for instance as a "carrot-eater," or they described the same agent as someone who "eats carrots whenever she can." After describing an agent in one of these two ways, they measured children's beliefs about the stability of the agent's carrot-eating tendencies. They found that children were more likely to think that the agent in the story had an enduring and stable trait when the agent was described as a carrot-eater (i.e., the noun form) than when the agent was described as eating carrots all the time (i.e., the verb form). This result was consistent with earlier research suggesting that, relative to other syntactic categories such as verbs or adjectives, nouns tend to license stronger inferences. As further illustration of this point, consider a study by Markman and Smith (reported in Markman, 1991) in which they asked participants to compare nouns (e.g., "John is *an intellectual*") and matched adjectives (e.g., "John is *intellectual*"). Participants in this study judged that the nouns described a stronger trait than the otherwise-identical adjectives (also see Reynaert & Gelman, 2007; Semin & Fiedler, 1988; Walton & Banaji, 2004).

A cue as seemingly insignificant as a word or a phrase can also influence people's behavior. For example, Cimpian, Arce, Markman, and Dweck (2007) have shown that praising children's successes on a drawing task with statements as similar as "You are a good drawer" and "You did a good job drawing" led 4-year-old children to make different inferences about the source of their success and, as a result, to behave

differently when they subsequently encountered a setback. When they heard "You are a good drawer" (which uses a noun to describe children's skill), children were more likely to infer that they had succeeded because of a talent they possess, which then made subsequent mistakes more difficult to overcome (because such mistakes reflected negatively on this supposed talent) and made children less likely to take on challenges (because they feared they would make more mistakes).

These results demonstrate that subtle syntactic differences are accompanied by differences in people's intuitions. Our hypothesis is that the syntactic differences between "knowledge" and "know" will likewise give rise to diverging intuitions. When "knowledge" is used instead of "know," a listener may infer that a speaker is intending to make an epistemically stronger statement. (We use the term "epistemic strength" to denote a cluster of related dimensions such as reliability, certainty, and evidential support.[1]) And since these syntax-dependent intuitions are present even in 4-year-olds, it is plausible to suppose that they are pervasive and implicit—the sort of basic, introspectively opaque inferences that laypeople and professional philosophers are equally likely to make (for a similar argument, see Leslie, 2013).

We are, however, agnostic as to whether the predicted effect of syntactic form-class on epistemic intuitions is due to an accompanying difference in the semantics of the terms "know" and "knowledge" or, alternatively, to a difference in the pragmatics of these terms (e.g., perhaps "knowledge" is more formal-sounding than "know"). The psychological literature seems to suggest that syntactic differences similar to that between "know" and "knowledge" have their effects via semantics, but a pragmatics-mediated effect is congenial to our present claim as well. If we found that the syntactic distinction between "know" and "knowledge" reliably affects people's intuitions about what is being denoted, in and of itself that would cast doubt on a foundational methodological assumption in epistemology. Whether these divergent intuitions have a semantic or pragmatic source is a separate matter.

---

[1] We aggregate across these distinct dimensions because we suspect that laypeople's concept of knowledge is not fine-grained enough to draw reliable distinctions between, say, being more certain vs. having more evidence. To foreshadow our results, this suspicion was empirically supported as well, insofar as participants' judgments about reliability, certainty, etc. were nearly indistinguishable.

3. The Reliability of Intuition

We are arguing that epistemic intuitions may be affected by the subtle syntactic differences between "knowledge" and "know." The thrust of this argument fits with a recent wave of research at the intersection of cognitive science and philosophy, research that is aimed at questioning the reliability of the intuitions (e.g., Alexander, Mallon, & Weinberg, 2010; Knobe, 2015). This work, inspired by pioneering work in judgment and decision-making (see Kahneman, 2011, for a review), has suggested that intuitive judgments—even the intuitions of professional philosophers (Schwitzgebel & Cushman, 2012, 2015)—are influenced by factors that should in principle be truth-irrelevant. For example, intuitive judgments in several branches of philosophy (including ethics, epistemology, metaphysics, philosophy of mind, and philosophy of language) may be swayed by demographic and contextual factors such as socioeconomic status (e.g., Haidt, Koller, & Dias, 1993), cultural background (e.g., Machery, Mallon, Nichols, & Stich, 2004; Sytsma, Livengood, Sato, & Oguchi, 2015), or whether one is reading a concrete or abstract vignette (e.g., Nichols & Knobe, 2007). Consequently, several philosophers have argued that intuitions should not be the final arbiters of the truth of a philosophical theory. The skepticism concerning the reliability of intuition is consistent with our claim that intuitions about "knowing" and "knowledge" may diverge in ways that have significant epistemic implications.[2]

4. The Present Studies: Syntactic Cues Affect Epistemic Inferences

As reviewed above, extant research speaks to the pervasive influence of syntactic categories on people's intuitions. Moreover, ongoing research in experimental philosophy suggests that reliance on intuitive judgments—at least under certain conditions—may lead philosophers to make problematic claims. On the basis of this work, we hypothesized that the syntactic

---

[2] We do not take any of these results to be definitive evidence that intuitive judgments are unreliable. Indeed, some philosophers have raised compelling concerns about these arguments (e.g., Bengson, 2013). We are merely suggesting that our line of reasoning is consistent with this growing body of research in experimental philosophy.

differences between the noun "knowledge" and the verb "know" may lead to subtly divergent epistemic inferences. We investigated this hypothesis across six studies with diverse methodologies and sources of data.

To begin, because words denoting categories (like "knowledge") tend to be viewed as describing something stronger and more enduring than words denoting actions (like "know" and "knowing"; e.g., Gelman & Heyman, 1999), we predicted that professional philosophers who disagree about where to set the evidential bar for knowledge would use "knowledge" and "know" differentially in their published work. In particular, philosophers who argue that knowledge has a high evidential bar are predicted to use the noun "knowledge" more frequently in their papers compared to philosophers arguing that knowledge has a lower evidential bar, who might instead rely on the (weaker) verb "know" and its derivatives (Study 1). This study provides a highly naturalistic test of our hypothesis, examining syntactic differences between two classes of influential philosophical papers that are part of the contemporary epistemological canon (as judged by professional philosophers). Building on this initial, "real-world" test of our hypothesis, we then sought to test whether "knowing" and "knowledge" diverge in the intuitions they elicit in the laboratory under tightly controlled conditions. We tested this possibility by eliciting laypeople's judgments about "know" vs. "knowledge" in the context of various commonplace topics (Studies 2–5), as well as a well-known philosophical thought experiment (Study 6). The results of these six studies supported our hypothesis. Philosophers and non-philosophers alike seem to draw subtly different inferences from "know" and "knowledge"—a finding that raises doubts about the epistemic interchangeability assumption and thus has deep implications for epistemological methodology and theories.

5. Study 1: Philosophers' Papers

Prior research suggests that nouns elicit stronger inferences than otherwise-similar verbs (e.g., Cimpian et al., 2007; Gelman & Heyman, 1999). On the basis of this research, we predicted that these syntactic form-classes would be used differentially in philosophical arguments as a function of the claims being made about the nature of knowledge. Specifically, philosophers defending the view that knowledge has a high evidential bar are predicted to use "knowledge" more frequently and

"know" less frequently than philosophers defending the opposing view (namely, that knowledge does not require that a belief cannot be rationally doubted). For brevity, we term the former class of views *infallibilist* and the latter *fallibilist*, although we recognize that in doing so we are glossing over many nuances in the content of these theories.

5.1. Method

5.1.1. PARTICIPANTS, MATERIALS, AND PROCEDURE

We recruited professional philosophers to participate in our study by posting an advertisement on a widely read philosophy blog (Leiter Reports: A Philosophy Blog). The post read that we were seeking a "list of the best papers defending fallibilism about knowledge and the best papers defending infallibilism about knowledge." Participants were asked to list the top five papers defending fallibilism about knowledge and the top five papers defending infallibilism about knowledge (in counterbalanced order). This prompt was an approximate means of capturing the distinction between philosophers who think that knowledge has a high epistemic bar and those who think it has a lower epistemic bar.

We received responses from 18 professional philosophers. Of these, 72% held tenure track jobs at major research universities. Two additional philosophers filled out our survey, but they simply listed names (e.g., Kant) rather than specific papers whose text we could analyze; thus, their responses were not usable. The outcome of this process was a set of 33 distinct papers defending infallibilism ($n = 12$) or fallibilism ($n = 21$). For a full list, please see Table 1.1 in the Appendix.

5.1.2. CODING

For each paper, we first tallied (1) how often the author used "know" and its derivatives (specifically, "knows," "knew," "known," and "knowing") and (2) how often the author used "knowledge." We then calculated a "know":"knowledge" ratio by dividing the first tally by the second. Our prediction was that the "know":"knowledge" ratio would be higher for fallibilist papers (where the goal is to show that knowledge does *not* require an agent to believe a proposition in a way that completely rules out error) than for infallibilist papers (where the goal is to show that knowledge *does* require an agent to believe a proposition in a way that completely rules out error).
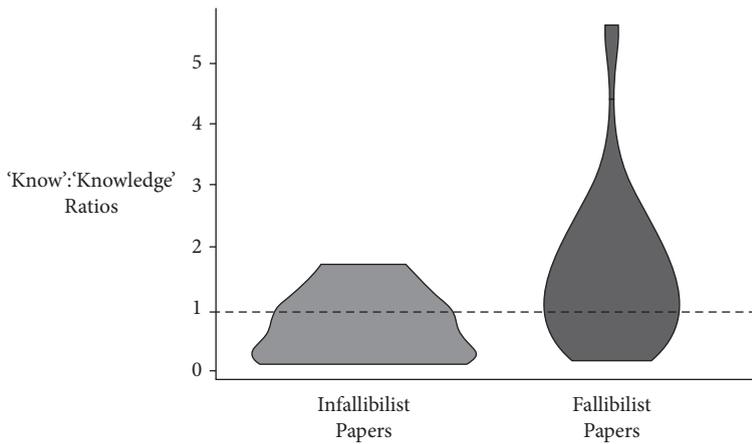
Figure 1.1 Violin plots of "know":"knowledge" ratios for infallibilist and fallibilist papers. Wider parts of the violin indicate a higher density of papers. Ratios below 1 indicate greater use of "knowledge," while ratios greater than 1 indicate greater use of "know" and derivatives. One paper from the fallibilist group (Dretske, 1971) was omitted from the figure because it had a very large "know":"knowledge" ratio.

5.2. Results and discussion

As predicted, fallibilist papers relied more on "know" and its variants than did infallibilist papers (see Figure 1.1). The median "know":"knowledge" ratio for fallibilist papers was more than twice that for infallibilist papers (1.48 vs. 0.64, respectively), Mann-Whitney $U = 58$, $z = 2.55$, $p = 0.011$.[3]

Thus, the "best papers" defending two prominent opposing epistemological theories differ in their reliance on the noun "knowledge"

[3] Supplementary analyses addressed two potential concerns: First, one paper in the fallibilist group (Burge, 2003; see Table 1.1 in the Appendix) is arguably about justification rather than knowledge per se. However, excluding this paper only made the predicted differences stronger (median "know":"knowledge" ratios = 1.51 vs. 0.64 for fallibilist vs. infallibilist papers, respectively), Mann-Whitney $U = 50$, $z = 2.73$, $p = 0.006$. Second, one paper categorized as defending fallibilism (Dretske, 1971) is on the epistemic operator "know" and might thus show a very high "know":"knowledge" ratio simply for this reason. However, the predicted effect emerged even when removing this paper from the dataset (median "know":"knowledge" ratios = 1.38 vs. 0.64 for fallibilist vs. infallibilist papers, respectively), Mann-Whitney $U = 58$, $z = 2.41$, $p = 0.015$.

(which predominates in infallibilist papers) vs. the verb "know" and its derivatives (which predominate in fallibilist papers). Whether or not the philosophers who authored these papers were aware of the different linguistic practices they were employing to advance their arguments, the fact that their use of epistemic terms varied dramatically (in a 2.3:1 ratio) as a function of the position they were defending is a strong indication that their epistemic intuitions were susceptible to the effect of subtle syntactic differences.

6. General Methods for Studies 2–6

As just discussed, our naturalistic data suggest that the syntactic differences between "knowledge" and "knowing" are importantly connected to philosophical theories. In what follows, we sought to supplement these data with several tightly controlled experimental tests of the claim that "knowledge" and "knowing" elicit diverging epistemic inferences. In the five studies reported below, we asked non-philosopher participants to compare agents who were said to know a fact and agents who were said to have knowledge of the same fact. The use of a within-subject design such as this (i.e., a design that elicits participants' intuitions about both "knowledge" and "know") was intended to capture the fact that, when crafting their arguments, philosophers have access to—and can freely compare—the intuitions that follow from the noun "knowledge" and the verb "know." Thus, a within-subject design most accurately translates the real-world phenomenon under investigation into a controlled experiment: Just as philosophers can explicitly or implicitly contrast their intuitive judgments about "knowing" vs. "knowledge," a within-subject experimental design allows participants to engage in this process as well. Aside from the shared within-subject design, however, the studies below varied on a number of dimensions such as their use of forced-choice vs. Likert-type items, the specific syntactic frames used, and the richness of the stimuli. To anticipate the results, all five studies indicated that syntactic differences influence epistemic intuitions in the predicted manner.

7. Study 2: Knowledge vs. Knowing

In Study 2, we compared people's judgments about two short phrases—one containing the term "knowledge," the other containing "knowing"—to

examine whether people associate "knowledge" (a noun) more strongly than "knowing" (a verb) with certainty and definiteness. We predicted such a difference based on the psychological literature showing that nouns license stronger inferences than verbs (e.g., Cimpian et al., 2007), as well as the finding that this syntactic distinction tracks differences in the height of the evidential bar set by different philosophical views (Study 1).

7.1. Method

7.1.1. PARTICIPANTS

Participants in Study 2 were 106 Amazon Mechanical Turk workers who were paid $0.50 for participating in the study, as well as volunteers recruited on a social media site.

7.1.2. MATERIALS AND PROCEDURE

Participants were asked to compare two phrases, "knowing that ___" (verb) and "knowledge that ___" (noun), on a number of epistemic dimensions. No propositional content was provided after these phrases (i.e., the blanks were actually left blank) to avoid contamination from specific contents. (However, subsequent experiments suggest that the specific propositional contents do not appear to affect participants' judgments.) Participants compared these two phrases on the following four dimensions that epistemologists have used in their theorizing (e.g., BonJour, 1985; DeRose, 1992):

(1) Which of these two seems to say something stronger?
(2) Which of these two seems to be more certain?
(3) Which of these two seems to be more dependent on one's circumstances? (reverse-coded)
(4) Which of these two seems to say something less definite? (reverse-coded)

For each of these questions, participants simply chose one of the options ("knowing that ___" or "knowledge that ___"). The four questions were shown in randomized order, as were the "knowing" vs. "knowledge" response options.

7.2. Results and discussion

Responses were scored as 0 or 1, with 1 indicating that participants deemed "knowledge" to be epistemically stronger than "knowing." Thus, "knowledge that ___" answers were coded as 1 for questions (1) and (2), whereas for questions (3) and (4) "knowing that ___" answers were coded as 1 instead. We averaged the responses across these questions into a composite measure representing the percentage of questions on which the participants indicated that "knowledge" denotes an epistemically stronger concept than "knowing" (Cronbach's $\alpha = 0.74$; see Table 1.2 in the Appendix for question-level results).[4] Scores above 50% on this composite indicate that "knowledge" was judged to indicate something epistemically stronger (more certain, more definite, etc.) than "knowing," whereas scores lower than 50% indicate that "knowing" is epistemically stronger. A one-sample t test against 50% revealed that, as predicted, subjects associated stronger epistemic content with "knowledge" than with "knowing" ($M = 60.8\%$, $SD = 36.5$), $t(105) = 2.88$, $p = 0.005$. Of the 106 participants, 58 (54.7%) had scores above 50% and only 33 (31.1%) had scores below 50%, $p = 0.012$ by a sign test.

In sum, participants distinguished between the epistemic implications of "knowledge" and "knowing," judging the former to say something stronger than the latter. Study 2 thus provides converging support for the claim that intuitions about "knowing" are not interchangeable with intuitions about "knowledge."

8. Study 3: "Has a Lot of Knowledge" vs. "Knows a Lot"

In the previous experiment, we asked participants to compare bare phrases: "knowing that ___" and "knowledge that ___." While these minimal stimuli were designed to avoid interference from the conceptual content associated with specific targets of knowledge or knowing, it may have also prompted participants to artificially differentiate

---

[4] Cronbach's α (Cronbach, 1951) is an estimate of the internal reliability or coherence of a scale. In essence, it is a measure of the extent to which the items in a scale correlate with one another. Cronbach's α values that are larger than 0.70 are conventionally considered acceptable. All the αs in our studies are above this threshold.

"knowing" and "knowledge" in ways that they would not do with more natural stimuli. Study 3 was conducted to address this concern. Specifically, we tested whether we would observe a difference even if participants considered complete, natural-sounding sentences using "know" vs. "knowledge."

8.1. Method

8.1.1. PARTICIPANTS

Participants were 101 Amazon Mechanical Turk workers who were paid $0.50 for participating in the study. Three additional participants were tested but excluded from the sample for having IP addresses outside of the United States ($n = 2$) or for missing a catch question ($n = 1$; see section 8.1.2).

8.1.2. MATERIALS AND PROCEDURE

Participants were asked to make epistemic judgments about the protagonists of two sentences: "Michael has a lot of knowledge" (noun) and "Christopher knows a lot" (verb). The addition of "a lot" was intended to make the stimulus sentences sound felicitous, while also leaving the content of knowing/knowledge unspecified. It is desirable to avoid specifying particular propositional content in these stimuli because such content could in principle introduce unforeseen confounds between the "knowledge" and "know" sentences. The names associated with the "knowledge" and "know" sentences were counterbalanced, as was the order of the sentences. The judgments were as follows:

(1) If you had to choose, which person is more likely to be a college professor?
(2) If you had to choose, which person do you think is more likely to be competent?
(3) If you had to choose, which person would you rather have your kids learn from?
(4) If you had to choose, which person is more likely to be misinformed? (reverse-coded)
(5) If you had to choose, which person do you think is more likely to make bad decisions? (reverse-coded)
(6) If you had to choose, which person would you trust less? (reverse-coded)

For each of these questions, participants had to choose "Michael" or "Christopher." Responses were coded as in the previous study, with 1 assigned to responses that indicated inferences from "knowledge" were epistemically stronger than inferences from "knows." Responses across the six questions were averaged into a composite percentage score (Cronbach's α = 0.81; see Table 1.3 in the Appendix for question-level analyses).

The study also included a catch question that instructed participants to select a particular response (e.g., "Michael"). This item was included to ensure participants were paying attention. One participant failed this catch item and was excluded from the sample.

8.2. Results and discussion

Participants chose the responses that indicated stronger inferences from "knowledge" than "know" on 78.4% of trials (SD = 29.2), which was significantly greater than chance (50%), $t(100)$ = 9.76, $p < 0.001$. Of the 101 participants, 75 (74.3%) had scores above 50% and only 15 (14.9%) had scores below 50%, $p < 0.001$ by a sign test.

Thus, even with natural-sounding, full-sentence stimuli, we still found that subtle syntactic differences in the language used to talk about epistemic states license substantially different inferences about the epistemic strength (certainty, reliability, etc.) of these states. These results provide further evidence for the hypothesis that syntactic differences may affect epistemic intuitions and the theories built on these intuitions.[5]

9. Study 4: Holding Propositional Content Fixed

In Study 4, we investigated three alternative interpretations of our prior results. First, one could worry that participants may have simply made different assumptions about the kinds of things one "knows a lot" vs. "has a lot of knowledge" about. For example, when someone "knows a

---

[5] Although within-subject designs are best suited to our research question, we note here that we also ran a between-subjects version of this study and replicated the observed effect ($p < 0.05$).

lot," perhaps she is assumed to have a lot of *knowledge-how* rather than a lot of *knowledge-that*; in contrast, someone who "has a lot of knowledge" may be assumed to have a lot of *knowledge-that*. Such differences could explain our findings without the need to posit that "knowledge" and "know" license different epistemic inferences. To address this issue, in Study 4 we held constant the propositional content of each item. Specifically, we asked participants to make epistemic judgments comparing knowing that P and knowledge that P, where P was a proposition such as that "avocados ripen faster in a paper bag." Holding the proposition constant allowed us to test the possibility that participants were simply filling in different epistemic content for agents who knew vs. had knowledge.

Second, perhaps participants in Study 3 interpreted "knows a lot" as being *non-factive* (i.e., non-truth-presuming; see Hazlett, 2010), whereas such interpretations were less likely for "has a lot of knowledge." To defeat non-factive readings of "know," we added the intensifiers "full"/ "fully" to the stimuli. That is, participants were asked to compare a person who "fully knows that P" with a person who "has full knowledge that P."[6] These stimuli help to block non-factive readings by strongly implying that P is true. (Ordinarily, it is infelicitous to claim that something false is "fully known.")

Finally, since some of the questions in Study 3 asked about academic-sounding activities (e.g., "which one is a college professor?"), it is possible that participants responded as they did because they superficially associated the word "knowledge" with "book smarts." (Note, though, that all six judgments in Study 3 showed the predicted difference [see Table 1.3 in the Appendix], so this possibility is somewhat unlikely.) Thus, one may wonder if the questions used in Study 3 led participants to artificially differentiate "knowledge" and "knowing." To eliminate this potential confound, in Study 4 we replaced the epistemic judgment items that referred to academic-sounding activities.

We predicted that participants would make stronger epistemic inferences from "knowledge" than from "know" even in this tightly controlled case.

---

[6] For further evidence against this alternative interpretation, see Study 6, where the actor is said to "*truly* know that P."

9.1. Method

9.1.1. PARTICIPANTS

Participants were 121 Amazon Mechanical Turk workers who were paid $0.50 for participating in this study. An additional 10 participants were tested but excluded from the sample for having IP addresses outside of the United States ($n = 2$) or for missing an attention check ($n = 8$).

9.1.2. MATERIALS AND PROCEDURE

Participants were presented with four different pairs of protagonists, one who "had full knowledge" of something and one who "fully knew" the same thing. The four pairs of sentences shown to participants were as follows:

(1a) Michael has full knowledge that you can clean stained coffee cups with baking powder.
(1b) Christopher fully knows that you can clean stained coffee cups with baking powder.
(2a) Robert fully knows that avocados ripen faster in a paper bag.
(2b) William has full knowledge that avocados ripen faster in a paper bag.
(3a) Matthew has full knowledge that apples are more effective at waking you up in the morning than bananas.
(3b) Joshua fully knows that apples are more effective at waking you up in the morning than bananas.
(4a) David fully knows that a wet rubber dish-washing glove is the quickest way to clean up pet hair.
(4b) James has full knowledge that a wet rubber dish-washing glove is the quickest way to clean up pet hair.[7]

The order of the two statements within each pair was randomized. After reading each pair, participants made the following epistemic judgments:

---

[7] We asked participants to make judgments about everyday, ordinary propositions to avoid introducing confounds that could spuriously produce the hypothesized effect. Although more abstract or unusual propositions may seem more philosophically relevant, they also may increase the likelihood of demand characteristics (see Orne, 1962). Because ordinary propositions are unlikely to produce an effect due to other unknown factors having to do with their contents, they provide a more appropriate test of our hypothesis.

(1) If you had to choose, who would you trust more on this matter?
(2) If you had to choose, who is better informed on this matter?
(3) If you had to choose, who is more likely to give you bad advice on this matter? (reverse-coded)
(4) If you had to choose, who is less confident on this matter? (reverse-coded)

Participants answered each of these questions by selecting one of the protagonists' names (e.g., Michael or Christopher). Responses were coded as in the previous studies, with 1 assigned to responses indicating that inferences from "knowledge" were epistemically stronger than inferences from "know." Responses were averaged across the four epistemic judgments to create a composite percentage score (Cronbach's α = 0.96; see Table 1.4 in the Appendix for judgment-level analyses).[8]

9.2. Results and discussion

Participants chose the responses that indicated stronger inferences from "knowledge" than from "know" on 67.6% of trials (SD = 35.5), which was significantly greater than chance (50%) on a one-sample t test, $t(120)$ = 5.45, $p < 0.001$. Of the 121 participants, 81 (66.9%) had scores above 50% and only 33 (27.3%) had scores below 50%, $p < 0.001$ by a sign test.

These results replicate those of Studies 2 and 3, and suggest that the effects we have observed so far are not driven by extraneous differences in the assumed propositional contents of "knowledge" and "know" sentences, by non-factive readings of "know," or by the use of academic-sounding activities in our epistemic judgment questions. Rather, these findings suggest that the noun "knowledge" and the verb "know" lead to different intuitions.

10. Study 5: Entailment Test

In Study 5, we used an alternative strategy to test the hypothesis that the epistemic intuitions elicited by "knowledge" and "know" are

---

[8] We note that the Cronbach's α reported here is particularly high. Thus, although one may *intuitively* think that giving bad advice, for example, is quite different than being confident about a particular proposition, the α suggests that participants are treating these items as equivalent. This supports the decision to average across items.

not interchangeable. In propositional logic, two propositions P and Q are treated as logically equivalent if P entails Q and if Q entails P. Thus, if "knowledge" and "know" elicit identical epistemic intuitions—if they can be used interchangeably—then people should agree equally with the following conditional claims:

(1) If you know that P, then you have knowledge that P, and
(2) If you have knowledge that P, then you know that P.

In contrast, if "knowledge" and "know" elicit *distinct* epistemic intuitions, then we should observe an asymmetry in people's agreement with these two conditionals—a bold prediction considering their similarity. Specifically, we predicted that people would agree with the second conditional ("If you have knowledge that P, then you know that P") more than with the first ("If you know that P, then you have knowledge that P"). In other words, people should think that the epistemic inferences licensed by "knowledge" are stronger than those licensed by "know" and thus having knowledge should entail knowing more than the other way around.

Study 5 also introduced a methodological improvement over the prior studies: The forced-choice nature of measures used in Studies 2–4 did not allow participants to indicate when they thought that "knowledge" and "know" were equally strong. For a more conservative test of our hypothesis, in Study 5 we asked participants to make their judgments about the two conditional statements using separate Likert-type scales (rather than forcing them to choose one statement as stronger).

10.1. Method

10.1.1. PARTICIPANTS

Participants were 87 Amazon Mechanical Turk workers and students enrolled in an introductory psychology course at a large Midwestern public university. Participants were paid $0.50 for participating in the study or received course credit. An additional 13 participants were tested but excluded for incorrectly answering screener questions that assessed basic competence with conditional (if–then) reasoning (see section 10.1.2). However, all effects reported here remained significant without excluding these participants ($ps < 0.05$).

10.1.2. MATERIALS AND PROCEDURE

Participants were first presented with directions explaining how to interpret conditional claims. After reading these instructions, participants were presented with two screener questions designed to test their basic understanding of conditional statements. The screener questions were:

(1) If Harry is a cat, then Harry is a mammal (correct answer: true), and
(2) If Larry is a mammal, then Larry is a cat (correct answer: false).

After answering these screener questions (which 13 participants failed), participants rated their agreement with both of the key conditional statements:

(1)   Know→Knowledge entailment:
"If you know a lot about X, then you have a lot of knowledge about X," and
(2)   Knowledge→Know entailment:
"If you have a lot of knowledge about Y, then you know a lot about Y."

The order of these statements was randomized. Participants marked their agreement on a scale from 0 ("completely disagree") to 10 ("completely agree").

10.2. Results and discussion

Our prediction was that participants would agree significantly more with the Knowledge→Know conditional claim than with the Know→Knowledge one. As predicted, participants' agreement ratings were higher for Knowledge→Know ($M = 9.23$, $SD = 1.40$) than for Know→Knowledge ($M = 8.49$, $SD = 2.25$), $t(86) = 3.13$, $p = 0.002$. Of the 87 participants, 19 (21.8%) agreed more with Knowledge→Know than with Know→Knowledge, whereas only six (6.9%) showed the opposite pattern, $p = 0.015$ by a sign test.

   Thus, although the literature has by and large assumed that "know" and "knowledge" elicit identical epistemic intuitions, our results suggest (again) that "knowledge" licenses epistemically stronger inferences than "know."

11. Study 6: Knowledge and Lotteries

In our final study, we compared people's epistemic inferences from "know" and "knowledge" in the context of a well-known thought experiment. By examining participants' judgments about a philosophical thought experiment, we were able to move closer to paradigmatic philosophical work while maintaining a level of experimental control that is difficult to achieve when examining naturalistic data (Study 1). In particular, in Study 6 we investigated the influence of the syntactic difference between "know" and "knowledge" on participants' judgments about a lottery case (Hawthorne, 2004).

In a lottery case, participants are told about a protagonist who has entered a raffle that they have an extremely small chance of winning. Intuitively, the protagonist in this case *does not know* they won't win the raffle because—despite the low probability of winning—the outcome is still uncertain (Hawthorne, 2004). However, if we imagine instead that the protagonist did not enter the raffle in the first place and is thus absolutely certain they won't win, then it is clear that the protagonist *knows* they won't win.

We predicted that intuitions about these lottery cases would differ subtly depending on whether participants are queried about the protagonist's "knowing" vs. "knowledge." When the protagonist has a small chance of winning (Some Chance condition), participants should be more likely to agree with "know" statements than with "knowledge" statements, since the latter seem to say something epistemically stronger than the former (see Studies 2–5). In contrast, when the protagonist is certain of the outcome because they did not enter the lottery (No Chance condition), participants should agree with "know" and "knowledge" statements equally strongly. This is because certainty about the outcome meets a high evidentiary threshold that should suffice for attribution of both knowing and knowledge. Thus, our theory predicts no effect of syntactic cues in this condition.

11.1. Method

11.1.1. PARTICIPANTS

Participants were 132 Amazon Mechanical Turk workers who were paid $0.50 to complete the experiment. An additional five participants were tested but excluded from the sample for having an IP address

outside of the United States (n = 3) or for missing our attention checks (n = 2).

We employed a 2 (Condition: Some Chance vs. No Chance; between subjects) × 2 (Syntactic Cues: "know" vs. "knowledge"; within-subject) design. Participants were randomly assigned to either the Some Chance condition (n = 66) or the No Chance condition (n = 68). In the Some Chance condition, participants were told that the protagonist (Emma) enters her company's raffle but has an extremely low probability of winning (1 in 100,000), and thus she forms the belief that she will not win the raffle. In the No Chance condition, participants were told that Emma has no chance of winning the company raffle (because she doesn't enter the raffle), and thus she forms the belief that she will not win the raffle.

After reading one of these vignettes, participants were asked to rate their agreement with two statements on a scale from 0 ("completely disagree") to 10 ("completely agree"):

(1) "Knows" statement:
"Emma truly knows that she won't win the company raffle."[9]
(2) "Knowledge" statement:
"Emma has true knowledge that she won't win the company raffle."

The order of these statements was randomized across subjects.

11.2. Results and discussion

We performed a 2 (Some Chance vs. No Chance) × 2 ("know" vs. "knowledge") mixed-design analysis of variance (ANOVA) to test the predicted interaction between condition and syntactic cues. This inter-action was indeed statistically significant, $F(1, 132)$ = 7.13, p = 0.009. As predicted, participants in the Some Chance condition (in which the protagonist was not certain of the lottery's outcome) displayed signifi-cantly higher agreement with the "know" statement (M = 4.20, SD = 3.54)

---

[9] "Truly" was included to further rule out the possibility that the effects we have observed were due to a non-factive reading of "knows." Although Study 4 included "fully" for the same reason, "truly" provides an even stronger means of blocking non-factive readings. It would be extremely infelicitous to claim that someone truly knows something false.

than with the "knowledge" statement ($M = 3.06$, $SD = 3.24$), $t(65) = 3.29$, $p = 0.002$. Of the 66 participants in this condition, 26 (39.4%) agreed more with the "know" than the "knowledge" statement, whereas only 11 (16.7%) showed the opposite pattern, $p = 0.002$ by a sign test. These findings support the claim that the noun "knowledge" licenses epistemically stronger inferences than the verb "knows." In contrast, participants in the No Chance condition endorsed the "know" statement ($M = 9.82$, $SD = 0.57$) and the "knowledge" statement ($M = 9.65$, $SD = 1.26$) at equally high levels, $t(67) = 1.54$, $p = 0.128$. Sixty-five of the 68 participants in this condition (95.6%) agreed with the "know" and "knowledge" statements exactly to the same extent. This is as we predicted: Certainty meets an evidentiary bar that is high enough for attribution of both knowledge and knowing.

In sum, here too we found that the subtle syntactic contrast between "know" and "knowledge" can lead to different intuitions. In this case, we observed the predicted effect in the context of a well-known thought experiment from contemporary epistemology.

12. General Discussion

Many philosophers have assumed that the epistemic inferences licensed by linguistic examples referring to "knowledge" and "knowing" are interchangeable (i.e., the epistemic interchangeability assumption). Our data suggest that this methodological assumption has not been sufficiently vetted. In six studies using a variety of methodological paradigms and sources of data (both naturalistic and laboratory-based), we provided empirical evidence that the noun "knowledge" licenses stronger epistemic inferences than the verb "know."

To summarize, we first observed "real-world" evidence for our hypothesis in professional philosophers' writing: Classic papers defending the claim that knowledge has a high evidential bar were considerably more likely to use the noun "knowledge" in building their arguments relative to papers that defended a weaker notion of knowledge; in the latter group, the verb "know" and its derivatives predominated. Second, we found experimental evidence for our hypothesis in laypeople's judgments: Across five experiments, non-philosopher participants drew stronger epistemic inferences from the noun "knowledge" than from the verb "know," which is consistent with the differential distribution

of these syntactic classes in philosophical papers defending infallibilism and fallibilism. In light of these data, epistemological arguments that are supported with evidence from ordinary-language intuitions may need to proceed cautiously.

In section 13, we discuss some of the potential implications of our findings for epistemology, and philosophy more generally.

13. Changing Conceptions of Knowledge

The view that knowledge has a high evidential bar is often associated with Descartes, who in the *Meditations* argued that if we are to provide a solid foundation for science then we must form our beliefs in a way that rules out all possible error (Descartes, 1641/2013). This view, however, has grown progressively unpopular in modern epistemology: There appear to be many straightforward cases in which our intuitions tell us that you can rightly say a subject knows proposition *P*, even though they do not believe *P* in a way that entirely rules out error. For example, it seems that you know Pluto is not made of cheese, even though your justification is defeasible. Guided by such intuitions, philosophers have recently argued that knowledge is a practical (e.g., Hawthorne, 2004; Pinillos, 2011; Stanley, 2005) or contextual (e.g., DeRose, 1992) concept rather than an "academic" one (Descartes, 1641/2013).

To speculate, the assumption that intuitions about what people "know" are interchangeable with intuitions about the concept denoted by "knowledge" may have contributed to these changing conceptions of knowledge, even if in a small way. Descartes set the evidential bar very high, seeking to remove all doubt from knowledge, because his goal was to provide a reliable foundation for science. However, as philosophers' methods of investigation started to shift toward reliance on linguistic intuitions (in the mid-20th century) and they began to use examples about what different agents can be said to "know" to inform their theories, the evidential bar for attributions of knowledge may have been inadvertently lowered and views of knowledge as practical (rather than academic) began to predominate. In other words, our data suggest that the historical shifts in theories about the nature of knowledge could be due partly to unnoticed shifts in the predominant method of investigation in epistemology rather than to

an improvement in philosophers' ability to reliably carve epistemic concepts at their joints.[10]

14. Further Theoretical Implications and Open Questions

This research contributes to the ongoing debate about the reliability of intuitions as a source of philosophical evidence. As already mentioned, work in experimental philosophy suggests that intuitions—even those of professional philosophers—are not infallible guides to the truth, as they seem to be affected by a host of factors that may be truth-irrelevant (e.g., Alexander, Mallon, & Weinberg, 2010; Schwitzgebel & Cushman, 2012, 2015; but see Bengson, 2013, for a critical response). Even if these factors did not influence intuitions, other philosophers have argued that intuitions cannot be empirically calibrated (e.g., Cummins, 1998), and so cannot serve as a reliable foundation for philosophical inquiry. Our data speak to these issues, offering some support for the concern that intuitions may be affected by as yet unnoticed factors that are not obviously connected to the relevant facts of the matter. If, as we found, subtle syntactic differences—differences that went unnoticed in the history of epistemology—can elicit systematically different epistemic intuitions, then these intuitions may need to be more closely scrutinized before being used as the foundation for an epistemic theory.

Our data also have theoretical implications for areas of philosophy beyond epistemology, such as philosophy of science, ethics, and philosophy of mind. For example, the present findings may speak to debates about what constitutes an *explanation* or an *intention* (e.g., Knobe, 2003; Waskan, Harmon, Horne, Spino, & Clevenger, 2014). As in the case of epistemology, it is possible that the literatures on these topics have been influenced by different syntactic framings: Do "explanation" (noun) and

---

[10] Of course, other factors also likely played a part in these theoretical shifts. For instance, the rise of *naturalized epistemology*, which attempts to build epistemic theories on the basis of psychological research, may have been influential as well. Psychological research has unequivocally shown that we do not have unmediated access to the world. Thus, if certainty were required for knowledge, it would appear very difficult for us to have any knowledge at all. Since many philosophers want to avoid theses that entail potentially global skepticism, they may have thus shifted to the view that knowledge cannot require certainty.

"explains" (verb) lead to the same inferences about what constitutes an explanation? How about "having the intention to act" (noun) vs. "acting intentionally" (verb)? Considering that even 4- and 5-year-old children are sensitive to similar syntactic cues (e.g., Gelman & Heyman, 1999), we think it is quite likely that the effects identified in our studies are not unique to knowledge *per se*. Instead, our data suggest a more pervasive influence that has simply gone unnoticed in philosophical theorizing.

However, several important issues are left unanswered by our studies. First, how exactly do these syntactic differences bear on "the facts of the matter" about the nature of knowledge? Although the noun "knowledge" leads people to infer more reliability and evidential support than does the verb "knowing," the present research does not unequivocally indicate just what the substance and contours of the denoted epistemic states are. Further research (both philosophical and psychological) is needed to sort out these complex matters. Second, our research leaves open the question of whether the syntactic contrast between "knowledge" and "knowing" elicits different epistemic intuitions via a semantic or a pragmatic process. Regardless, the methodological consequences are clear—linguistic statements about "knowing" and "knowledge" give rise to different epistemic intuitions which, if not accounted for, may mislead epistemological theorizing. Finally, one may ask whether our research suggests that the debate between different philosophical camps bottoms out in a merely linguistic distinction. For example, are fallibilists and infallibilists engaged in a merely verbal dispute, while perhaps being in fundamental agreement about the nature of knowledge? It should be clear that no single paper can answer all of these important questions. Nonetheless, we hope our work sparks new life into these and other ongoing epistemological debates.

15. Conclusion

We have shown that the syntactic differences between "knowledge" (a noun) and "know" (a verb) give rise to distinct intuitions, with "knowledge" licensing epistemically stronger inferences than "know." Based on these results, we argued that philosophical arguments that proceed on the assumption that the intuitions elicited by "knowledge" and "knowing" can be used interchangeably may warrant further scrutiny.

Appendix
Supplemental Results and Data

Table 1.1 The fallibilist and infallibilist papers in Study 1.

Papers	"Know": "Knowledge" Ratio
Fallibilist papers	
Anderson, C. (2014). Fallibilism and the flexibility of epistemic modals. *Philosophical Studies*, *167*(3), 597–606.	1.54
Brown, J. (2008). Knowledge and practical reason. *Philosophy Compass*, *3*(6), 1135–52.	.70
Brueckner, A. (2005). Fallibilism, underdetermination, and skepticism. *Philosophy and Phenomenological Research*, *71*(2), 384–91.	1.27
Burge, T. (2003). Perceptual entitlement. *Philosophy and Phenomenological Research*, *67*(3), 503–48.*	.17
Cohen, S. (1988). How to be a fallibilist. *Philosophical Perspectives*, *2*, 91–123.	2.17
DeRose, K. (1995). Solving the skeptical problem. *The Philosophical Review*, *104*(1), 1–52.	1.74
Douven, I., & Olders, D. (2008). Unger's argument for skepticism revisited. *Theoria*, *74*(3), 239–50.	2.62
Dretske, F. I. (1970). Epistemic operators. *The Journal of Philosophy*, *67*(24), 1007–23.	27.20
Dretske, F. (1971). Conclusive reasons. *Australasian Journal of Philosophy*, *49*(1), 1–22.	5.61
Fantl, J., & McGrath, M. (2002). Evidence, pragmatics, and justification. *Philosophical Review*, *111*, 67–94.	3.00
Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis*, *23*(6), 121–3.	2.20
Goldman, A. I. (1976). Discrimination and perceptual knowledge. *The Journal of Philosophy*, *73*(22), 771–91.	1.12
Goldman, A. I. (1979). What is justified belief? In G. S. Pappas (Ed.), *Justification and knowledge* (pp. 1–23). Dordrecht: Reidel.	1.67
Howard-Snyder, F., & Feit, N. (2003). Infallibilism and Gettier's legacy. *Philosophy and Phenomenological Research*, *66*(2), 304–27.	.35
Littlejohn, C. (2008). From E = K to scepticism? *The Philosophical Quarterly*, *58*(233), 679–84.	2.19

Pryor, J. (2000). The skeptic and the dogmatist. *Noûs*,   1.48
    34(4), 517–49.

Reed, B. (2010). A defense of stable invariantism. *Noûs*,   .77
    44(2), 224–44.

Reed, B. (2013). Fallibilism, epistemic possibility, and   .65
    epistemic agency. *Philosophical Issues*, *23*(1), 40–69.

Reed, B. (2002). How to think about fallibilism.   .61
    Philosophical Studies, *107*(2), 143–57.

Sainsbury, R. M. (1997). Easy possibilities. *Philosophy and*   1.03
    Phenomenological Research, *57*(4), 907–19.

Stanley, J. (2008). Knowledge and certainty. *Philosophical*   .66
    Issues, *18*(1), 35–57.

Infallibilist papers

Bird, A. (2005). Abductive knowledge and Holmesian   1.01
    inference. *Oxford Studies in Epistemology*, *1*, 1–31.

BonJour, L. (2010). The myth of knowledge. *Philosophical*   .14
    Perspectives, *24*(1), 57–83.

Dodd, D. (2011). Against fallibilism. *Australasian Journal*   1.72
    of Philosophy, *89*(4), 665–85.

Feit, N., & Cullison, A. (2011). When does falsehood   .38
    preclude knowledge? *Pacific Philosophical Quarterly*,
    92(3), 283–304.

Fumerton, R. (2010). Fencing out pragmatic   1.07
    encroachment. *Philosophical Perspectives*, *24*(1),
    243–53.

McDowell, J. (1995). Knowledge and the internal.   .14
    Philosophy and Phenomenological Research, *55*(4),
    877–93.

Moon, A. (2012). Warrant does entail truth. *Synthese*, *184*(3),   .12
    287–97.

Neta, R. (2011). A refutation of Cartesian fallibilism. *Noûs*,   1.01
    45(4), 658–95.

Parent, T. (2007). Infallibilism about self-knowledge.   .15
    Philosophical Studies, *133*(3), 411–24.

Stroud, B. (1984). Skepticism and the possibility of   .47
    knowledge. *The Journal of Philosophy*, *81*(10), 545–51.

Unger, P. (1971). A defense of skepticism. *The*   1.56
    Philosophical Review, *80*(2), 198–219.

Williamson, T. (1997). Knowledge as evidence. *Mind*,   .80
    106(424), 1–25.

* This paper is arguably about justification rather than knowledge. However, excluding it from our analyses only makes the predicted difference stronger.

Table 1.2 The average percentage (and standard deviation) of trials in Study 2 in which participants indicated that "knowledge that ___ " has stronger implications than "knowing that ____."

Item	Percentage (SD)	p value against chance
1. Which of these seems to say something stronger?	64.2 (48.2)	.003
2. Which of these seems to say something more certain?	51.9 (50.2)	.700
3. Which of these seems to say something more dependent on one's circumstances? (reverse-coded)	64.2 (48.2)	.003
4. Which of these seems to say something less definite? (reverse-coded)	63.2 (48.5)	.006

Note: $N = 106$. The p values derive from one-sample t tests against chance (50%).

Table 1.3 The average percentage (and standard deviation) of trials in Study 3 in which participants chose the answer indicating that "knowledge" has stronger implications than "know."

Item	Percentage (SD)	p value against chance
1. Which of these two is more likely to be a college professor?	90.1% (30.0)	< 0.001
2. Which of these two is more likely to be competent?	74.3% (43.9)	< 0.001
3. Which of these two would you rather have your kids learn from?	77.2% (42.1)	< 0.001
4. Which of these two is more likely to be misinformed? (reverse-coded)	79.2% (40.8)	< 0.001
5. Which of these two is more likely to make bad decisions? (reverse-coded)	74.3% (43.9)	< 0.001
6. Which of these two would you trust less? (reverse-coded)	75.2% (43.4)	< 0.001

Note: $N = 101$. The p values derive from one-sample t tests against chance (50%).

Table 1.4 The average percentage (and standard deviation) of trials in Study 4 in which participants chose the answer indicating that "knowledge" has stronger implications than "know."

Item	Percentage (SD)	p value against chance
1. Who would you trust more on this matter?	68.4 (39.0)	< 0.001
2. Who is better informed on this matter?	69.6 (37.7)	< 0.001
3. Who is more likely to give you bad advice on this matter? (reverse-coded)	67.1 (36.1)	< 0.001
4. Who is less confident on this matter? (reverse-coded)	65.3 (38.4)	< 0.001

Note: $N = 121$. The p values derive from one-sample t tests against chance (50%).

References

Alexander, J., Mallon, R., & Weinberg, J. M. (2010). Accentuate the negative. *Review of Philosophy and Psychology*, *1*(2), 297–314.

Bengson, J. (2013). Experimental attacks on intuitions and answers. *Philosophy and Phenomenological Research*, *86*(3), 495–532.

BonJour, L. (1985). *The structure of empirical knowledge*. Cambridge, MA: Harvard University Press.

Buckwalter, W. (2012). Non-traditional factors in judgments about knowledge. *Philosophy Compass*, *7*(4), 278–89.

Buckwalter, W. (2014). Intuition fail: Philosophical activity and the limits of expertise. *Philosophy and Phenomenological Research*. doi: 10.1111/phpr.12147.

Buckwalter, W., & Schaffer, J. (2015). Knowledge, stakes, and mistakes. *Nous*, *49*(2), 201–34.

Cimpian, A., Arce, H. M. C., Markman, E. M., & Dweck, C. S. (2007). Subtle linguistic cues affect children's motivation. *Psychological Science*, *18*(4), 314–16.

Cronbach, L. J. (1951). Coefficient alpha and the internal structure of tests. *Psychometrika*, *16*(3), 297–334.

Cummins, R. C. (1998). Reflection on reflective equilibrium. In M. DePaul & W. Ramsey (Eds.), *Rethinking intuition: The psychology of intuition and its role in philosophical inquiry* (pp. 113–28). Lanham, MD: Rowman & Littlefield.

DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, *52*(4), 913–29.

Descartes, R. (1641/2013). *Meditations on first philosophy*. Cottingham, J. (trans.). New York: Cambridge University Press.

Gelman, S. A., & Heyman, G. D. (1999). Carrot-eaters and creature-believers: The effects of lexicalization on children's inferences about social categories. *Psychological Science*, *10*(6), 489–93.

Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis*, *23*(6), 121–3.

Haidt, J., Koller, S. H., & Dias, M. G. (1993). Affect, culture, and morality, or is it wrong to eat your dog? *Journal of Personality and Social Psychology*, *65*(4), 613–28.

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

Hazlett, A. (2010). The myth of factive verbs. *Philosophy and Phenomenological Research*, *80*(3), 497–522.

Kahneman, D. (2011). *Thinking, fast and slow*. New York: Macmillan.

Knobe, J. (2003). Intentional action and side effects in ordinary language. *Analysis*, *63*(279), 190–4.

Knobe, J. (2015). Philosophers are doing something different now: Quantitative data. *Cognition*, *135*, 36–8.

Leslie, S. J. (2013). Essence and natural kinds: When science meets preschooler intuition. *Oxford Studies in Epistemology*, *4*, 108–66.

Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, *74*(4), 549–67.

Machery, E., Mallon, R., Nichols, S., & Stich, S. P. (2004). Semantics, cross-cultural style. *Cognition*, *92*(3), B1–B12.

Markman, E. M. (1991). *Categorization and naming in children: Problems of induction*. Cambridge, MA: MIT Press.

Murray, D., Sytsma, J., & Livengood, J. (2013). God knows, but does God believe? *Philosophical Studies*, *166*(1), 83–107.

Nichols, S., & Knobe, J. (2007). Moral responsibility and determinism: The cognitive science of folk intuitions. *Nous*, *41*(4), 663–85.

Orne, M. T. (1962). On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. *American Psychologist*, *17*(11), 776–83.

Pinillos, N. Á. (2011). Some recent work in experimental epistemology. *Philosophy Compass*, *6*(10), 675–88.

Reynaert, C. C., & Gelman, S. A. (2007). The influence of language form and conventional wording on judgments of illness. *Journal of Psycholinguistic Research*, *36*(4), 273–95.

Ryle, G. (1945). Knowing how and knowing that: The presidential address. *Proceedings of the Aristotelian Society*, *46*, 1–16.

Schwitzgebel, E., & Cushman, F. (2012). Expertise in moral reasoning? Order effects on moral judgment in professional philosophers and non-philosophers. *Mind & Language*, *27*(2), 135–53.

Schwitzgebel, E., & Cushman, F. (2015). Philosophers' biased judgments persist despite training, expertise and reflection. *Cognition*, *141*, 127–37.

Semin, G. R., & Fiedler, K. (1988). The cognitive functions of linguistic categories in describing persons: Social cognition and language. *Journal of Personality and Social Psychology*, *54*(4), 558–68.

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

Stanley, J. (2011). *Know how*. New York: Oxford University Press.

Stanley, J., & Williamson, T. (2001). Knowing how. *The Journal of Philosophy*, *98*(8), 411–44.

Sytsma, J., Livengood, J., Sato, R., & Oguchi, M. (2015). Reference in the land of the rising sun: A cross-cultural study on the reference of proper names. *Review of Philosophy and Psychology*, *6*(2), 213–30.

Walton, G. M., & Banaji, M. R. (2004). Being what you say: The effect of essentialist linguistic labels on preferences. *Social Cognition*, *22*(2), 193–213.

Waskan, J., Harmon, I., Horne, Z., Spino, J., & Clevenger, J. (2014). Explanatory anti-psychologism overturned by lay and scientific case classifications. *Synthese*, *191*(5), 1013–35.

Williamson, T. (2002). *Knowledge and its limits*. New York: Oxford University Press.