

Experimental Evidence Supporting Anti-intellectualism About Knowledge

Ángel Pinillos and Shawn Simpson¹

1 Introduction

According to the traditional conception of knowledge, knowledge is a purely intellectual concept. In this chapter, we give some evidence against the orthodox view. We think that knowledge is, in part, a practical concept. To be specific, whether someone who believes *P* also counts as knowing *P* may depend on more than just the evidence available to her, or other intellectual features of her situation. It may also depend on practical facts, such as the cost of being wrong about *P*.

Anti-intellectualism about knowledge (AIK) has recently been defended by a number of authors.² These authors have pursued various strategies for defending the thesis. In this chapter we pursue a strategy that is in some ways old and deeply entrenched, but in other ways, new. Where we stay with tradition is in holding that judgments about hypothetical (or actual) cases give us powerful evidence for first-order theses about knowledge and other fundamental concepts. Where we move away from tradition is in how we collect and analyze these judgments. We rely here on the judgments of people who have not prejudged the issue.

2 General methodology

According to the traditional way of utilizing judgments about cases, a philosopher will make a judgment about a hypothetical case and use the

content of the judgment as a premise in an argument.³ Making the judgment is a private or semiprivate affair. The philosopher may or may not consult other people, including peers, on the matter (but this further consulting, if it happens, is often informal and unsystematic—for example, the philosopher may or may not consider variations of the thought experiment before settling on the judgment they will use in theorizing). Now, focusing on judgments about knowledge, there is a reasonable suspicion that the traditional method of collecting and deploying intuitions may be less reliable than previously thought. First, some philosophers have admitted that the judgments of philosophers may very well be biased. Goldman (2007, p. 15) writes:

philosophers are leery about trusting the intuitions of other philosophical analysts who have promoted general accounts of the analysandum (e.g. knowledge of justification). Commitment to their own favored account can distort their intuitions, even with respect to their own (pre-theoretical) concept.

The possibility of biased judgments (which Goldman is calling “intuitions”) even among experts is also echoed by Schaffer (2006, p. 90), who admits that his own judgments about knowledge may be biased: “Perhaps my intuitions are unusual, and no doubt are theoretically biased.” These self-reports raise worries about the traditional practice of harnessing judgments.

Second, empirical research on folk judgments gives some evidence that judgments about knowledge may vary with features that may seem to some to be irrelevant to the subject matter under investigation. For example, some research supports the idea that certain judgments about knowledge vary across gender and cultural lines.⁴ We should not always take for granted then that philosophers themselves are immune from these effects.⁵ We believe there is a legitimate worry that the traditional method of harnessing judgments may sometimes yield results that are more a reflection of the philosopher’s background or idiosyncratic features than of knowledge. This is then another reason to be cautious about the traditional practice of collecting intuitions.

Despite these and related issues, we do not think philosophers should give up on the traditional method of collecting and deploying judgment about cases. What we think is warranted, however, is a pluralism in methods—where

the new methods we devise aim to get around some of the shortcomings of traditional philosophical methodology. This doesn't mean that the new methods are better than the old methods. In fact, they may even be less reliable overall. But even if this is the case, it does not mean that it cannot give us philosophically important information in certain cases. It seems to us that, for certain instances, we can benefit from pursuing multiple strategies.

In this chapter we employ the tools of experimental philosophy and the behavioral sciences. We seek to collect judgments about knowledge in controlled experiments not just from one or two individuals, but from hundreds of people who have no vested interest in this or that theory in epistemology. We are thereby more likely to avoid the bias Alvin Goldman and Jonathan Schaffer warned us about. We should also be able to identify idiosyncratic judgments. What is more, if we discover patterns of judgments that vary along demographic dimensions, we can use this data to make a more informed judgment about the quality of the judgments in question.

A further benefit of the method we employ is that we can now consider a broader range of judgments. For example, we can attempt to elicit reactions which take the form of numerical responses to some questions about hypothetical scenarios. Individual responses, even from experts, may not all converge on a unique numerical answer. They may admit of a degree of inter- or even intrapersonal variability. Nevertheless, we may discover philosophically informative tendencies toward certain types of responses. Whether there is such a tendency (as opposed to just a lot of noise) is something that can be discerned with the help of modern statistical techniques, and likely cannot be discerned by the use of traditional armchair methods alone. We will see examples of this approach in this chapter.

Of course, the experimental method will have drawbacks. Experimenters often rely on philosophically naive populations (as we do in this chapter). Although this method has the advantages we saw above, it is still problematic in situations where the judgments under investigation concern a concept which requires advanced training to master, or where the hypothetical scenario concerns difficult or technical subject matter. But not all cases are like that. In this chapter, for example, we seek judgments concerning a widely used and nontechnical concept, "knowledge," and make use of simple hypothetical scenarios. Moreover, we use various checks to ensure our subjects are

responding to what is being asked of them. Thus, we think that the responses of our subjects can be used in philosophical analysis in a similar way that judgments are used in traditional philosophical theorizing.

3 Our hypothesis and Anti-intellectualism about knowledge

Suppose you have an unimportant haircut appointment today at 3 p.m. You quickly glance at your calendar on your way out and form the belief that it is at 3 p.m. This may suffice for you to know the appointment is at 3 p.m. Now modify the case so that the appointment is a matter of life and death. In that case, if you form the same belief in the same way, would you now know that the appointment is at 3 p.m.? It may seem to some that quickly glancing at the calendar in the second case is not enough to attain knowledge. If this is right, then AIK may be true, since, apparently, the only relevant difference between the cases are practical and nonintellectual features of the situation.

Let us define AIK as the thesis that whether someone who believes P also counts as knowing P may sometimes depend on practical features concerning the agent's relation to P, including what is at stake for the agent. We need an easier way of talking about this so we will often just describe the thesis as saying that knowledge is sometimes *sensitive to practical interests*. We will also include in the definition a claim about the direction of the effect. In cases where stakes matter, for example, higher stakes raise the bar for knowledge. As Weatherson (2012) has pointed out, the thesis has an existential form. It doesn't claim that for every difference in practical interest facts, there will be a corresponding difference in knowledge facts. Rather, it just claims that there are cases in which differences in practical interest facts correspond to differences in knowledge facts.

As mentioned at the outset, we defend AIK by appealing to judgments about cases. But, we are adopting the experimental method for collecting and analyzing those intuitions. The hypothesis we seek evidence for is (H) (where "sensitive to practical interests" is given the gloss from above):

(H) Folk attributions of knowledge are sometimes sensitive to practical interests.

(H) by itself does not deductively entail AIK. But (H) does entail AIK if we add the assumption that the attributions of knowledge mentioned in (H) semantically express propositions that are true or true with respect to the hypothetical cases presented to the folk. Unless we accept broad skepticism about our ordinary capacity to attribute mental states to others, this assumption should be *prima facie* accepted—at least for the simple cases we will be discussing. The *prima facie* assumption, of course, can be defeated. For example, the folk could be making performance errors, their attributions may be tracking conversational implicatures that diverge from the semantically expressed propositions, or perhaps the attributions could be explained by other theories in epistemology. These and related worries, if actual, will sever the connection between (H) and AIK. These worries are infinite and there is no way that we can address all of them. So our strategy will be to provide some new evidence for (H) and at the same time give reasons to think that some of the worries are not so worrisome. This will be enough to give new support for AIK.

4 Prior work

As of the writing of this manuscript, there have been over half a dozen papers reporting on experiments that are directly relevant to assessing (H). The “first wave” reported results that disfavor (H).⁶ The second wave reported results that challenge the first-wave papers and are positive for (H).⁷ Finally, two further papers challenge the second-wave results. In a later section we will address in detail the challenge to the second-wave results.⁸ But for now, we want to make some general remarks about the first- and second-wave studies.

4.1 Negative results

The negative first-wave papers report on experiments, most of which follow a certain pattern. The researchers presented subjects with pairs of vignettes that differ only in what is at stake (practical interests) for the protagonist (low vs. high stakes). The subjects were then asked to record how much they agreed with a particular knowledge claim ascribed to the protagonist. In general it was found that there were no statistically significant differences in the mean

level of agreement with the knowledge attribution across the level of stakes. This was thought to constitute some evidence against (H).

One may worry that since (H) is an existential thesis, the fact that differences in reactions to knowledge attributions were not discovered for some cases does not impugn (H). However, this defense of (H) is weak, since the researchers used vignettes similar to those that defenders of AIK claimed illustrated their thesis. But, there are other worries with these experiments. We will focus on two issues, the “Awareness of Stakes” problem and the “Same Evidence” problem.

The “Awareness of Stakes” Problem: As mentioned above, all the experiments in the negative papers involved comparing subjects’ responses to a high-stakes scenario with ones typically involving a low-stakes situation. However, with the exception of one probe, all the high-stakes vignettes depicted a protagonist who was aware that the stakes were high. This may have a distorting influence in that subjects may expect the protagonist in the high-stakes situation to be anxious or less confident than his low-stakes counterpart (who is in turn aware of his low-stakes situation). This is enough to create a possible confound in the experiment, since the probes in question may differ not just with respect to stakes. They may also differ in perceived levels of anxiety, confidence, and whatever else may be involved when a person is aware that the stakes are high. In sum, awareness of stakes on the part of the protagonist may weaken the evidentiary force of these experiments against our hypothesis (H).⁹ We note, moreover, that these problems need not just arise in the experimental setting. They may also arise when philosophers construct thought experiments for themselves.¹⁰

The “Same Evidence” problem: There is a general worry about keeping the perceived evidence available to the protagonist constant across the low- and high-stakes vignettes (recall that most of these experiments are between-subject studies). For example, in Buckwalter (2010), subjects are told about a protagonist who claims to know that a bank will be open on Saturday based on the evidence that the protagonist was at the bank last week. But this simple statement describing the evidence possessed can be interpreted in different ways. For example, the protagonist might be thought to have asked a worker about the bank hours, or perhaps he asked another customer, or he might have just quickly glanced at the hours posted. This issue is exacerbated by the fact that in most of the probes, the protagonist is aware of the stakes *and*

sincerely claims that he knows the relevant proposition. Supposing knowledge is sensitive to stakes and subjects are aware of this, subjects must think that the evidence the protagonist in the high-stakes situation has is sufficient to meet the threshold for knowledge in a high-stakes case (as opposed to the threshold for a low stakes case)—after all, the protagonist knows the stakes are high yet still knows the relevant proposition. If this is right, then the perceived evidence would not be the same across the high- and low-stakes scenarios. The confound then may weaken the evidentiary force of these experiments against our hypothesis (H).¹¹

4.2 Positive results

Pinillos (2012) reports on some new experiments that aim to minimize the two problems above. First, he developed a type of probe he dubbed “evidence seeking,” where the evidence available to the protagonist of a vignette is not fixed by the experimenter. Instead, subjects are asked their opinion about how much evidence the protagonist would need to collect before he counts as knowing. The hypothesis (H) would gain support if subjects in the high-stakes conditions tend to say that more evidence is needed to know than subjects in the low-stakes condition. This type of probe then would alleviate, somewhat, some of the issues surrounding the “Same Evidence” problem. Second, in one of his studies, the high-stakes vignette was “ignorant” in the sense that the protagonist in question was not aware of the high stakes. This type of probe alleviates, somewhat, some of the issues surrounding the “Awareness of Stakes” problem. Overall, Pinillos’ results support (H).

Sripada and Stanley (2012) reported on experiments that also support (H). They used the direct method from the first-wave experiments where subjects are asked about their level of agreement with an attribution of knowledge of the form “X knows P” (concerning a given vignette). However, their experiments were carefully designed to avoid some worries from earlier studies. For example, the protagonists of the vignettes are not aware of what is at stake, nor do they self-attribute knowledge.

We do not want here to engage in a detailed discussion of prior work. We do wish to point out, however, that there is an asymmetry in the data collected so far. The data that have been taken to support (H), including the work of

Pinillos (2012) and Sripada and Stanley (2012), involve experiments where statistical significance was found (at the standard p level of 0.05) and hence the null hypothesis (no stakes effect) was rejected. In contrast, the experiments from the first wave, which are taken to count against (H), all invoke null results. That is, they involve experiments where statistical significance was not found, and so the null hypothesis (no stakes effect) fails to be rejected. It is well known that a null result does not, in general, support the null hypothesis nearly to the same degree that a significance finding disconfirms it. In many cases, it may hardly support the null hypothesis at all. Hence, the researcher who wants to use statistical insignificance to defend the null hypothesis (no stakes effects) needs to do quite a bit more if she wants to make her point. This is especially so in our case where there are other statistically significant findings that tend to disconfirm the null hypothesis. (Of course, we cannot criticize the first-wave authors for not taking into account data from the second-wave studies. Our point mainly concerns what we should conclude after the two waves, and not so much how we should criticize the first-wave studies.)

There are a variety of statistical techniques that can be used by the researcher wishing to confirm the null hypothesis.¹² In particular, one may report the chances of making a Type II error. That is, the chances of accepting the null hypothesis (no stakes effect) in the case that the null hypothesis is false. The higher this probability, the less prone we should be to accept the null hypothesis based on statistically insignificant results. Without this information, which can be gotten through a power analysis, it is highly unclear to what extent nonsignificant findings can support the null hypothesis (no stakes effects).¹³ Stevens (2007, p. 111) makes the point:

Researchers not sufficiently sensitive to the power problem may interpret non-significant results from studies as demonstrating that “treatments” made no difference. In fact, however, it may be that treatments did make a difference, but that the researchers had poor power for detecting the differences. The poor power may result from small sample size and/or from small effect size. The danger of low power studies is that they may stifle or cut off further research in an area where effects do exist, but perhaps are more subtle (as in personality, social, or clinical psychology).¹⁴

Following this advice, we can raise a concern about drawing strong conclusions from the first-wave results (not that the authors drew strong conclusions—they

were appropriately cautious). Although all three papers use null results to cast doubt on (H), only Feltz and Zarpentine's papers directly address power. They worry (correctly) that their studies did not have enough statistical power to detect a stakes effect. As a corrective, they collapse the data sets, thus getting a larger sample size. Indeed, when they consider the aggregate (454 in high-stakes and 184 in low-stakes participants), Feltz and Zarpentine do discover a statistically significant difference between high- and low-stakes conditions (a fact that has not been often reported in the subsequent literature). But they claim this does not support (H) since the effect size is too small¹⁵:

the effect size approaches triviality at .01. The small effect size indicates that the variance in knowledge attributions explained by the stakes, however real, is very small. . . . the results of this analysis suggest that the practical facts in these situations do not qualitatively change knowledge attributions and they are not likely to be a fundamental or important feature of our ordinary knowledge attributions. (p. 17)

We raise two worries about this response. First, let us grant that the effect size might be small, but we don't think we can conclude that the effect of practical interests is "not likely to be a fundamental or important feature of ordinary knowledge attributions." For the purposes of epistemology and hence philosophy, we do not require the effects of practical interests to be psychologically important. Philosophical importance is not psychological importance. A subtle psychological effect is still an effect. And so may be enough to support the thesis that knowledge is sensitive to practical interests. However, we grant that a small effect may weaken the case for (H).

Second, many of the studies that Feltz and Zarpentine aggregate have the problems I mentioned at the outset ("awareness of stakes" and "same evidence" problems). Hence, the fact that statistically significant differences were found in these studies (even with a small effect), despite the problems mentioned, should be encouraging to the defender (H).

5 Our studies

In this chapter, we follow Pinillos (2012) in developing new evidence-seeking probes. However, we go well beyond that work in a number of ways, including

(1) testing to see whether some other pragmatic features, beyond stakes, play a role in knowledge ascriptions. In particular, we probe whether the probability of what could go wrong if the agent is mistaken plays a role in knowledge ascriptions; (2) we develop new “ignorant” probes in which both the high- and low-stakes protagonists are unaware of what is at stake for them; and (3) we develop a new way to test folk attitudes concerning the connection between knowledge and action, a connection that plays a central role in Anti-intellectualist theories of knowledge.

5.1 Study 1: Evidence-seeking experiments, stakes, and probabilities

We describe two “evidence-seeking” experiments aimed to test (H). In both experiments we ask subjects their opinions about how much evidence is required before a fictional agent counts as knowing some proposition P. We predict that as the practical interests become more pressing for the fictional agents, subjects’ responses will reflect a more stringent evidentiary condition on knowledge. Experiment 1 tests whether folk attributions of knowledge are sensitive to stakes. Experiment 2 tests whether they are sensitive to stakes and the likelihood that potential negative repercussions will actually happen. In addition, Experiment 2 attempts to control for a further confound concerning ignorance of stakes/probabilities.

5.2 Experiment 1 (Water purifier)

Method

A total of 141 subjects from Amazon Turk were paid 15 cents (US) each to take exactly one of two surveys. The two surveys (conditions) are dubbed “Low Stakes” and “High Stakes.” They both concern a protagonist, Brian, who is installing a water purifier at home because he does not like the taste of the tap water. However, in High Stakes, it is also the case that the water supply has been poisoned. If Brian fails to assemble the water purifier properly, he and his family might die. However, Brian is ignorant of these high stakes. For both vignettes, subjects are told that Brian has gone to the other room to get the water purifier instructions off the internet. He copies them down on a piece of

paper and heads to the water faucet. All subjects then were asked to respond to the following prompt:

(Opinion question) Suppose Brian goes back and compares his entire written copy to the instructions online, and he can do this as many times as he wants. After how many comparisons will Brian know he has written them down correctly? Please write your answer in the box below. This should be a whole number. (Note: If you think Brian already knows, write “0.” If you think he’ll never know, no matter how many times he checks, write “never”)_____.

Participants were given 10 minutes to respond. As explained above, in accordance with (H), we expect that the numerical answers for High Stakes will be higher than the numerical answers to Low Stakes.

Results

We discarded 47 surveys because the participants either failed to follow instructions, failed a reading comprehension check or did not write a numerical response to the main prompt. The reading comprehension checks were placed before the target question and included a question to see if subjects were aware of what was at stake. We computed statistics for $N = 94$ subjects. The results are as follows: Low Stakes ($N = 46$, $m = 0.72$, $sd = 0.72$), High Stakes ($N = 48$, $m = 1.29$, $sd = 1.254$). The difference in means was statistically significant $t(75.54) = -2.70$, $p < 0.01$. Cohen’s $d = 0.54$ (this is a medium-size effect).

Discussion

These results support (H). Subjects seem to think it takes more evidence to know something when the stakes are high (compared with when stakes are low). This experiment improves on previous studies. First, the fact that the experiment is “evidence seeking” minimizes the “Same Evidence” problem. This is because instead of recording a subject’s agreement with a knowledge claim—a claim we think is in many cases likely influenced by prior judgments concerning the amount of evidence possessed by the protagonist—we record a subject’s judgment concerning how much evidence would actually be needed for such a claim. Second, the fact that the protagonist in high stakes is ignorant of what is at stake for him makes the “Awareness of Stakes” problem less acute.

In particular, there is less reason to think that the higher responses to high stakes are due to the subjects assuming that the protagonist is anxious or less confident because of his awareness of the stakes.

5.3 Experiment 2 (Airplane)

Experiment 2 goes beyond the first one in two main ways. First, an important feature of the first experiment is that in High Stakes, but not in Low Stakes, the protagonist is ignorant of what is at stake for him. That is, he's not aware of all the relevant factors determining the stakes of his situation. Therefore the conditions differ in more than just what is at stake for the protagonist. They differ in whether there is ignorance. This difference could be the source of a confound. We eliminate it in Experiment 2.

Second, recall that Experiment 1 tests whether folk attributions are sensitive to stakes, a dimension of practical interests. However, one's practical interests surrounding a belief may depend on more than just the costs of being wrong. It may also involve the probability of the possible costs coming to fruition. The second experiment then considers not just what happens when the cost of being wrong varies, but also tests what happens when we change the probability of the possible costs being realized.¹⁶ The first prediction is that as the cost of being wrong goes up, subjects will raise the bar for knowledge. The second prediction is that as we raise the probability of the thing that may go wrong, subjects will also raise the bar for knowledge. Both predictions are in accordance with (H).

Method

We constructed four "Airplane" vignettes, each corresponding to a condition. They are about an airline steward, Jessie, who is assigned to find a name on a roster of 200 passengers before a flight. In every case, Jessie thinks the name he is being asked to look up belongs to someone that is supposed be bumped up to first class and it wouldn't matter much if Jessie failed to find the name. In every case, Jessie looks through the roster just once and comes to think that the name is not on the list. Subjects are told that in fact the name is not on the list.

The four cases all differ only in some feature concerning Jessie's practical interests and do not differ in intellectual features of the situation. In particular,

they differ in what is at stake—what event would happen if Jessie made a mistake. But they also differ on the probability of that event happening. In high stakes–high probability (HSHP), if Jessie is wrong, there is a high probability that the person on the list, a criminal, would hijack the plane. In high stakes–low probability (HSLP), if Jessie is wrong, the probability is low (but existing) that the criminal would hijack the plane. In low stakes–high probability (LSHP), if Jessie is wrong, there is a high probability that the person on the list, a nice guy, would accept the invitation to go to first class. In LSLP, there is a low probability that the nice guy would accept the invitation to go to first class. Again, Jessie is unaware that the name belongs to a nice guy/hijacker and unaware of the probability that he would go to first class/hijack the plane.

The four surveys totaling 305 were randomly distributed to volunteer students taking introductory courses at Arizona State University. The main prompt was as follows:

We are now interested in your opinion about what it would take for Jessie to know that the name is not on the roster (the name of the nice guy/hijacker). Recall that according to the story, Jessie has already surveyed the entire roster once. How many more times do you think Jessie needs to survey the entire roster before he knows the name is not on the list (enter a whole number: 0,1,2,3, . . . etc. or write “never” if you think Jessie will never know) _____.

The surveys contained reading comprehension checks. The reading comprehension checks occurred before the target question and they included questions checking to see if subjects knew what was at stake (and the probabilities). The subjects were given approximately 10 minutes to complete the entire survey.¹⁷

Results

We discarded 75 surveys because participants either failed to follow instructions, failed a reading comprehension check or wrote “never” as a response to the main prompt.¹⁸ We computed statistics for $N = 230$ subjects. LSLP ($N = 50$, $m = 1.6$, $sd = 0.969$), LSHP ($N = 58$, $m = 1.76$, $sd = 0.823$), HSLP ($N = 61$, $m = 1.93$, $sd = 0.944$), and HSHP ($N = 61$, $m = 2.15$, $sd = 1.152$) (Figure 1.1).

A one-way ANOVA reveals that there were statistically significant differences between the responses across the conditions. $F(3,224) = 3.211$, $p < 0.05$.

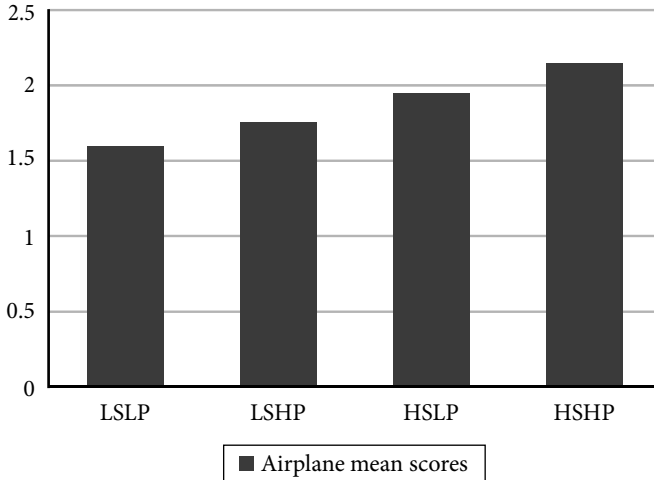


Figure 1.1 Airplane mean scores.

This suggests that practical interests play a role in knowledge ascriptions. Tukey's post hoc reveals statistically significant differences only between the LSLP and HSHP conditions, $p < 0.05$, Cohen's $d = 0.52$ (large effect size).

A two-way ANOVA was used for a factorial analysis. Comparing all the low-stakes cases against the high-stakes cases reveals a statistically significant main effect for stakes $F(1,224) = 7.639$, $p < 0.01$. Cohen's $d = 0.38$ (moderate effect size). No statistically significant main effect was found for probability. And no interaction effect was found.

Discussion

As mentioned above, there was a significant main effect for stakes. This supports (H): that folk attributions of knowledge are sensitive to practical interests. However, there is no significant main effect for probability. Nor were there statistically significant differences between LSLP and LSHP or between HSLP and HSHP. So we failed to get good evidence that probability plays a role in knowledge attributions. However, two facts suggest that further research may discover such a role. First, the mean scores reveal that for each of the stakes conditions (low and high), the higher probability condition corresponds to a higher mean response. Second, as mentioned above, Tukey's post hoc analysis for the four conditions revealed a significant difference only between LSLP

and HSHP. These conditions differ in both stakes *and* probability. This suggests that probability may be playing some role though not one that we were able to detect directly.¹⁹

6 Study 2: Agreement/disagreement with knowledge ascriptions

The previous experiments made use of evidence-seeking probes. In this study, we attempt a different type of experiment to ensure that the prior results were not merely products of the way we asked the questions. We construct pairs of conditions concerning potential knowers (protagonists) of some proposition P. The conditions differ in what is at stake for the protagonists. Subjects are asked about the extent to which they agree or disagree with a statement that says that the protagonist knows P. As with Experiment 2 from Study 1, in both the low-stakes and high-stakes conditions, the protagonist is ignorant of what is at stake. Moreover, to ensure the vignettes represent the protagonist as having the same beliefs and matched for accuracy, the protagonist is always mistaken about what is at stake: but it is always the same mistake.

Method

We developed three pairs of vignettes (Coin, Air, and Bridge) each with two conditions for what is at stake for the protagonist (Low and High). The surveys were taken by Amazon Turk workers in the United States. We present the Coin vignettes in full (see the Web appendix for the others):

Coin Low Stakes: Peter is a college student who has entered a contest sponsored by a local bank. His task is to count the coins in a jar. The jar contains 134 coins. Peter mistakenly thinks the contest prize is one hundred dollars. In fact, the prize is just a pair of movie passes for this weekend. Peter wouldn't want them, however, since he is leaving town this weekend. So nothing bad would happen if Peter doesn't win the contest. After counting the coins just once, Peter concludes there are 134 coins in the jar. His friend, who also thinks the prize is one hundred dollars says to Peter "you only counted once, even if there are in fact 134 coins in the jar, you don't know there are 134 coins in the jar."²⁰ You should count them again."

Coin High Stakes: Peter is a college student who has entered a contest sponsored by a local bank. His task is to count the coins in a jar. The jar

contains 134 coins. Peter mistakenly thinks the contest prize is one hundred dollars. In fact, the prize is \$10,000 which Peter really needs. He would use the money to help pay for a life-saving operation for his mother who is sick and cannot afford healthcare. So the stakes are high for Peter since if he doesn't win the contest, his mother could die. After counting the coins just once, Peter concludes there are 134 coins in the jar. His friend, who also thinks the prize is one hundred dollars says to Peter "you only counted once, even if there are in fact 134 coins in the jar, you don't know there are 134 coins in the jar. You should count them again."

Along with some comprehension questions, subjects were presented with the following prompt followed by a 7-point Likert scale (0–6) where 6 is "strongly agree" and 3 is "neutral":

Besides giving Peter advice about what he should do, Peter's friend also said that Peter doesn't know something. He said that since Peter only counted the coins once, Peter doesn't know that there are 134 coins in the jar (even if it turns out there are 134 coins in the jar). We are interested in your opinion about this. To what extent do you agree with the following statement:

"PETER KNOWS THERE ARE 134 COINS IN THE JAR"

In accordance with (H), we predicted a higher level of agreement with the knowledge statement in low stakes than in high stakes. We also ran two other probes, Air and Bridge, with 5-point Likert scales (0–4) where 4 is "strongly agree" and 2 is "neutral" (See Web Appendix). We made similar predictions.

For these experiments, we also added a "normative" question before the knowledge prompt. For example, in the COIN case we asked whether the subject thought Peter should count the pennies again. The main reason we do this is to help prevent subjects from thinking that the knowledge question implicates a question about what Peter should do (since we just asked this question, we would be violating a principle about relevance if we asked for it again immediately after). We think this feature of the experiment, together with the fact that the knowledge prompt is designed to help subjects focus on the concept of "knowledge," can help alleviate possible worries about implicatures.

Results

For the statistical analyses, we excluded those who failed comprehension checks or failed to follow instructions. Table 1.1 displays the main results.

Table 1.1 Study 2

	Low stakes	High stakes
Coin (Likert: 0–6)	N = 87, m = 3.68, sd = 1.80	N = 78, m = 3.06, sd = 1.76
Air (Likert 0–4)	N = 25, m = 2.16, sd = 1.03	N = 30, m = 2.03, sd = 1.0
Bridge (Likert 0–4)	N = 28, m = 2.32, sd = 1.16	N = 31, m = 1.71, sd = 1.13

In accordance with (H), the mean level of agreement with the knowledge attribution is lower in the high-stakes conditions than in the low-stakes conditions. The differences reached statistical significance for the Coin vignettes $t(161.78) = 2.23$, $p = 0.023$, $d = 0.35$ (small effect size) and the Bridge vignettes $t(57) = 2.053$, $d = 0.53$ (medium effect size). But not the Air Vignettes.²¹

Discussion

The results from this experiment support (H). Mean level of agreement with the knowledge attribution differed across the stakes conditions. Significance was reached for two out of the three experiments, suggesting that stakes play a role in knowledge attribution in accordance with (H). Moreover, we get corroborating data for (H) which does not rely on “evidence-seeking” experiments. This is significant since the “evidence-seeking” results conflict with earlier work and so it might be thought that the former results are simply due to an artifact of the experiment.

7 Study 3: Knowledge reason principle

So far, we have provided some evidence that folk attributions of knowledge are sensitive to practical interests (H). Our perspective is that a good explanation of this purported fact is that knowledge is in fact sensitive to practical interests. Now, it has been noted that AIK plausibly follows from deeper principles connecting knowledge and action. Given that people seem to be using knowledge as if the notion were sensitive to practical interests, we might wonder whether they implicitly accept any one of those principles. One such

principle is the Reason-Knowledge principle defended by Hawthorne and Stanley (2008):

(RKP) Where one's choice is p-dependent, it is appropriate to treat the proposition that p as a reason for acting if and only if you know that p.

The principle connects knowledge with a normative claim about action. The connection is intimate since the principle states an equivalence between the notions (when the possible object of knowledge is relevant for the action).

Pinillos (2012) uncovered some evidence that people implicitly accept the principle. What was discovered was that people seemed to replace a question about knowledge with a normative question about action, questions that are deemed equivalent according to the principle. It was also discovered that people responded to questions about knowledge with the same answers other people gave to normative questions about action, where again, the questions are equivalent according to RKP.

To be sure, this sort of evidence does not in any way establish that the folk implicitly accept RKP. But it makes the idea more plausible. We add further to the evidence in Study 3. Experiment 1 of this study reports results that support RKP.

It should be flagged, however, that RKP is a strong principle. It is much stronger than is required to make plausible that knowledge is sensitive to practical interests AIK.²² Consider the weaker principle, ACTION:

(ACTION) (Where P is relevant for action) If X knows P, then it is appropriate for X to act on P.

Experiment 2 will report on some evidence which suggests that the folk accept ACTION.

7.1 Study 3: Experiment 1 (Do People Accept RKP?)

Method

For this experiment, we test whether people implicitly accept RKP. To do this, we revisit Study 1, Experiment 2 (Air). Recall that in all four conditions (LSLP, LSHP, HSHP, HSLP), we asked subjects an evidence-seeking question: "How many more times do you think Jessie needs to survey the entire roster before he knows the name is not on the list _____." Now, right before we asked

this “knowledge” question, we asked subjects a normative question about the relevant action in the vignette:

Recall that according to the story, Jessie has already surveyed the entire roster once. At least how many more times do you think he should survey the entire roster looking for the name? _____.

If RKP is true, this normative question is equivalent (given the context) to the knowledge question. To be sure, if Jesse needs to survey the roster four times before he knows the name is not on the list, then (according to RKP) it wouldn't be appropriate for Jesse to use the belief that the name is not on the list in the relevant action until he surveys the list four times. But what is the relevant action? Given the story and Jesse's task, the relevant action is reporting that the name is not on the list. That is, it wouldn't be appropriate for Jesse to report that the name is not on the list until he surveys the roster four times (we are assuming that there are no other reasons available that Jesse could use to justify making the report). At the very least then, Jesse should survey the roster four times (response to the normative prompt). That is, if people accept RKP, then responses to the normative prompt should correlate with responses to the knowledge prompt.

Results

Across all four conditions of the Air experiment ($N = 228$), there is a statistically significant correlation between the knowledge prompt responses ($m = 1.87$, $sd = 0.996$) and the normative prompt responses ($m = 1.75$, $sd = 0.994$); $r = 0.708$, $p < 0.001$. Moreover, the effect size is large. This is some evidence that subjects are treating the normative and knowledge prompts in the same way and also evidence that subjects implicitly accept RKP. As we have seen, this adds plausibility to (H).

7.2 Study 3: Experiment 2 (Do People Accept ACTION?)

Method

For this experiment, we test whether people implicitly accept the principle: ACTION. In order to do this, we revisit the results from the Low-Stakes and High-Stakes conditions for the Coin, Air, and Bridge probes (Study 2). Recall that in these probes we sought responses concerning the level of agreement

with a “knowledge” prompt. In COIN, for example, we asked about the statement “Peter knows there are 134 coins in the jar.” Now right before we asked this question we also asked a normative question. In the COIN probe, for example, we asked our subjects whether they thought that Peter should count the pennies again.²³ Here, they only had three options: NO, NEUTRAL, and YES. Coding the NO and NEUTRAL in one category and coding the YES in a second category, we can compare the answers to the knowledge prompt across these two groups. Similar comparisons are made for AIR and BRIDGE.

Focusing on the COIN case, if people tend to implicitly accept the ACTION principle, we should see that subjects in the “yes, should count the coins again” category are less likely (compared with the other group) to agree with the knowledge statement from the prompt above. Similar predictions are made about AIR and BRIDGE.

Results

As predicted, we found evidence that people accept the ACTION principle. In the COIN case (7-point Likert from 6 was “strongly agree” and 3 was neutral). Among those who answered YES SHOULD COUNT AGAIN, the mean response was 3.1 ($sd = 1.69$). Among those who answered NO/NEUTRAL, the mean response was 3.7 ($sd = 1.9$). The difference was statistically significant $t(163) = 1.91, p = 0.029$ (one-tailed), $d = 0.3$. A similar result was found for AIR (5-point Likert where 4 is “Strongly Agree” and 2 is “Neutral”): YES SHOULD CHECK ROSTER AGAIN ($M = 1.96, sd = 0.908$), NO ($M = 2.88, sd = 1.246$), $t(53) = 2.5, p = 0.016, d = 0.95$. And, again, the same result holds for BRIDGE (5-point Likert as with AIR): YES SHOULD JUST CROSS THE BRIDGE ($M = 2.32, sd = 1.06$), NO ($M = 1.56, sd = 1.19$), $t(57) = 2.58, d = 0.68$. These experiments suggest that those subjects who thought that the protagonist did not have enough information to act on P were less likely to agree that the protagonist knows P. This indirectly supports the idea that people accept ACTION.

7.3 Discussion of both experiments

The foregoing results give some evidence that people implicitly accept RKP and ACTION. These principles have been used by researchers as major premises to justify AIK. A number of issues arise which we address in turn.

First, there is no question that the level of support that these data provide for the claims that people accept RKP and ACTION is not very strong. A number of worries come to mind. In the general discussion below we consider an important objection developed by Buckwalter and Schaffer.

Second, although (as defenders of AIK argue) RKP and ACTION may entail (assuming fallibilism) that knowledge is sensitive to practical interests AIK, it does not follow that *folk acceptance* of those principles entails (H), that is, folk attributions of knowledge are sometimes sensitive to practical interests. However, a person who accepts those principles (and fallibilism) but is such that her attributions of knowledge are not sensitive to stakes is less rational than a person who accepts the principles and is such that her attributions of knowledge are sensitive to stakes. Barring some special reason to think otherwise, we should think a person is more rational rather than less rational. Hence in this case, because we seem to lack those special reasons, evidence that the folk accept the principles gives us some further evidence that folk attributions of knowledge are sometimes sensitive to practical interests.

Third, suppose it turns out that although the folk implicitly accept principles like RKP and ACTION, (H) is nonetheless false. The implicit folk acceptance of the principles may actually directly support AIK, bypassing (H), since the implicit acceptance of those principles will support those principles themselves (in the way that folk belief often supports philosophical claims). And arguably, those principles lead to AIK. Insofar as those principles are thought to be natural, the discovery that they are implicitly accepted by the folk adds to their plausibility.²⁴

In sum, we have gathered some new evidence in favor of the notion that people implicitly accept principles used by theorists to support AIK. We indicated how these data can support (H) and how even if not, they may still support AIK.

8 Buckwalter and Schaffer objection

Buckwalter and Schaffer (MS) have raised an important worry about the evidence-seeking experiments we use here and that also appear in Pinillos (2012). They think that subjects' responses to the evidence-seeking

knowledge prompts reveal very little about their use of “knowledge.” And so the evidence-seeking experiments cannot tell us much about whether ordinary uses of knowledge are sensitive to practical interests (H).

Buckwalter and Schaffer’s criticisms of evidence-seeking experiments can be broken down into three parts. First, they show that replacing the word “knows” in the evidence-seeking experiments with other attitude verbs like “believes,” “hopes,” and “guesses” yields responses that are not different (as measured by statistical significance) from the original answers to the “knows” question. They conclude that the original responses are not really about knowledge. Second, they provide a positive account of how people are interpreting the evidence-seeking prompts—this interpretation is not about knowledge. Finally, they run another experiment that they argue should yield a certain result if the evidence-seeking experiments were really testing folk uses of “knowledge,” but in fact, they don’t get that result.

These are valuable criticisms. Before we address them in detail, we should keep in mind that the experimental evidence in favor of (H) goes beyond the evidence-seeking experiments. Experiments in this chapter as well as Pinillos (2012) and Sripada and Stanley (2012) involve nonevidence-seeking experiments which support (H).

8.1 Beliefs, guesses, and hopes

In an attempt to give support for the claim that ordinary knowledge attributions are sensitive to practical interests (H), Pinillos (2012) gave participants a pair of vignettes about a student Peter who will soon turn in a paper for an English class. In one condition, the cost of having a single typo is high (he could fail the class) and in another case, the cost is low (the teacher won’t care much). Pinillos then asked participants the following evidence-seeking question:

“How many times do you think Peter has to proofread his paper before he knows that there are no typos? times.”

Pinillos discovered that the median answer for the low- and high-stakes condition were 2 and 5, respectively. The difference was statistically significant. Moreover, Pinillos ran a few variations on this experiment including some that involved ignorance of what is at stake and others that involved a situation

where the protagonist (Peter) formed the belief that there are no typos right from the start. Pinillos also ran another case involving a person having to count pennies in a jar for a contest. In all these cases, he detected a stakes effect. Pinillos then goes on to conclude (together with further evidence) that this evidence supports AIK.

Buckwalter and Schaffer (2013) and Buckwalter (this volume) doubt that these data support Anti-intellectualism because they think these results say very little about how people deploy the concept of knowledge. The main argument they give for this claim comes from some follow-up experiments where they replicate Pinillos' experiments but also where "knows" in the prompt is replaced with "believes," "guesses," and "hopes" (These questions will be referred to as the "Knows," "Belief," "Guess," and "Hopes" questions, respectively):

"How many times do you think Peter has to proofread his paper before he believes/guesses/hopes that there are no typos? _____ times."

There are two main findings here. First, they find a stakes effect for each of these constructions. Second, they find no statistically significant differences between responses to these constructions and responses to the original Knows question.²⁵ Buckwalter and Schaffer conclude that the effect of stakes has nothing to do with people's perceptions of knowledge. They go on to offer a hypothesis about why this may be.

As Buckwalter and Schaffer discuss, it is possible to explain the results for "believes" and "guesses" in a manner that does not impugn the evidence-seeking experiments.²⁶ Concerning belief, we can follow Williamson (2000, 47) in thinking that knowledge is the norm for belief. For instance, subjects may accept that one should believe P only if they know P. If so, responses to the Belief and Knows prompts should be similar.

Concerning guesses, there are uses of "guess" that are sensitive to the guesser's practical interests. Note first that guesses need not be wholly blind. Consider, for example, the colloquial definition of a scientific hypothesis as an "educated guess." How much evidence will a scientist gather before she makes an educated guess? Plausibly, this may depend on what is at stake—the costs of being wrong. Hence some types of guessing may be thought to be sensitive to practical interests.

Consider also various “guessing” games. Country Fairs often have a “guess the weight of the pumpkin” contest (or sometimes an animal, like a pig). The instructions for such contests are to guess the weight of the object in front of them. Of course, participants may wish to think hard about the answer and try to use various strategies for getting the question right especially if the prize is really big—hence the stake sensitivity of guessing. If it is relatively easy for the subject to know (if only they gathered a little bit more information), then the extra effort taken before they are willing to guess might just be enough for them to put them in a position to know. In the Typo vignette, it is not very difficult for Peter to know that there are no typos (by proofreading x times, say), so we would expect Peter to proofread x times even before he is willing to guess there are no typos. In sum, we don’t find it very surprising that agents would give the same answer to the Guess and Belief questions.

The defense we just gave does not seem to be available for “hopes,” however. We do not see how correct responses to the Hopes question could be the same as correct responses to the Knows question. Nonetheless, we think these data can be explained in a way that is not problematic for us (or for the thesis advanced in Pinillos (2012)). To see this, we will consider and evaluate three possible explanations of Buckwalter and Schaffer’s data. The first is a thesis about the intelligibility of the Hopes prompt. The second is their preferred explanation of the data. The third is the explanation we prefer. We think our explanation vindicates the evidence-seeking probes.

First, one might worry that the Hopes evidence-seeking question (“How many times do you think Peter has to proofread his paper before he hopes that there are no typos? _____ times.”) is too difficult to understand. We don’t normally expect that hoping for something should depend on how much evidence we gather, so the question seems to presuppose something odd. According to this hypothesis, responses to the Hopes question do not reflect speaker competence since subjects do not properly understand it. This hypothesis, however, is not supported by the data we collected. We ran an experiment where we presented subjects with the Hopes question and other subjects with the Knowledge question. We then asked them to report the extent to which they understood the question (We used a Likert Scale with endpoints at 1 = not understanding at all, and 4 = fully understanding). What we discovered was that there were no differences in the reported level of

understanding between the questions and that the level of understanding was very much near the “fully understanding” ceiling.²⁷ So this first hypothesis is not supported empirically. Apparently, subjects have no problem understanding the Hopes question. Though perhaps we should be cautious concerning self-reports of comprehension.

The second hypothesis we consider is the one put forward by Buckwalter and Schaffer (2013). Here’s their explanation about how subjects are interpreting the Typo probes involving the evidence-seeking question (“How many times do you think Peter has to proofread his paper before he believes/guesses/hopes that there are no typos? _____ times.”):

... Pinillos is . . . seeing a stakes effect on the modal element “has” embedded in his complex probes. This “has” is most naturally read as a deontic modal (a “deontic modal” is one interpreted normatively). For instance, Typo probe know is most naturally reading as asking how many times Peter should, given his goals in life and the circumstances he finds himself in, proofread his paper before he knows that there are no typos. It is entirely unsurprising to find a stakes effect on deontic modality. Everyone can agree that the practical consequences of error matter when it comes to considering what one needs to do, normatively speaking. As the stakes get higher, Peter should proofread his paper more times, period. A fortiori he should proofread his paper more times before forming any attitude. Knowledge plays no role whatsoever in the matter, which is why replacing “know” with “believe” or “guess” or “hope” makes no difference whatsoever to the data. (p. . . .)

Buckwalter and Schaffer think that the stakes effect found is attributed to the modal “has” and it has little to do with the particular attitude mentioned (“knows,” “hopes,” etc.). According to this interpretation, the stakes effect discovered only reveals that subjects will think Peter should check more times for typos before forming any attitudes when the stakes are high (compared with low stakes). It doesn’t reveal much about people’s use of “knowledge” per se. Let us call this interpretation of the evidence-seeking question, “the modal reading.” We now turn to an investigation of the modal interpretation hypothesis.

To test Buckwalter and Schaffer’s hypothesis, we ran an experiment where we presented a set of subjects with the low-stakes Typo vignette and another set, the high-stakes Typo vignette. We then asked all subjects both the Knows and Hopes questions concurrently.²⁸ The idea here is that if the responses to

the Knows and Hopes prompts were different for an individual, then that individual is not giving the questions the modal reading (on the modal reading, the Hopes and Knows questions should be given the same answers). Furthermore, if the stakes effect still persists, then we can conclude that the stakes effect is not to be explained by subjects giving the evidence-seeking questions the modal reading. Hence, Buckwalter and Schaffer's modal interpretation hypothesis would be on the wrong track.

Our idea is supported by the data. Placing the Knows and Hopes questions side by side reveals that subjects give different answers to these questions. When the Hopes question is presented first ($N = 40$), the mean response (for both low and high stakes) for the Hopes question is 1.5 ($SD = 1.038$) and the mean response for the Knows question is 3.7 ($SD = 1.09$). A paired-samples t-test reveals a statistically significant difference between these responses $t(39) = 12$, $p < 0.01$ (We excluded five outliers with responses above 9.). Essentially the same result holds when the Knows question is presented first ($N = 29$). Hopes mean = 1.97 ($SD = 0.9$), Knows mean = 3.93 ($SD = 1.4$). Paired sample t-test: $t(28) = 10.06$, $p < 0.01$ (We exclude five outliers with responses above 9.). Moreover, the stakes effect still persists for the Knows question (though not for Hopes). The mean for the Knows low-stakes response was 3.31 ($N = 39$, $SD = 0.86$). The mean for the Knows high-stakes response was 4.42 ($N = 31$, $SD = 1.36$). The difference was statistically significant: $t(48.35) = 3.95$, $p < 0.01$. The mean for the Hopes low-stakes response was 1.54 ($N = 39$, $SD = 0.96$). The mean for the Hopes high-stakes response was 1.9 ($N = 39$, $SD = 1.0$). The difference was not statistically significant: $t(68) = 1.53$, $p = 0.153$. There were no ordering effects for the Knows questions. The mean for the Knows questions when they were presented first was 3.93 ($SD = 1.43$) and 3.7 ($SD = 1.09$) when they were presented second, $N = 69$, $t(67) = 0.75$, $p = 0.45$. There were no ordering effects for the Hopes question. The mean for the Hopes questions when they were presented first was 1.5 ($SD = 1.03$) and 1.97 ($SD = 0.9$) when they were presented second, $N = 69$, $t(67) = 1.9$, $p = 0.057$.²⁹

We now discuss our preferred explanation of the puzzling data Buckwalter and Schaffer uncovered. We believe that the stakes effect that they discovered for the Hopes prompt and the fact that the Hopes responses did not differ from the Knows responses is due to the "Anchoring and Adjustment" effect (Tversky and Kahneman 1974).³⁰

The effect, which researchers have called “strikingly pervasive and robust,” is exhibited when humans attempt to make an estimate under uncertainty such as the one required for the evidence-seeking experiments.³¹ Their responses will tend to gravitate toward an anchor point.³² For example, when subjects are asked to estimate the year George Washington was elected US. President (1788), they will first anchor on a known number such as 1776—the year of independence. Their estimate will then tend toward the anchor. In a study by Epley and Gilovich (2006), subjects’ responses to this question averaged at 1779.67 even though the range of possible answers (given by subjects taken from the same population) is 1777–1784. Hence, the response mean skewed toward the anchor number. Similar findings were replicated for a number of different questions.

We think the evidence-seeking experiments applied to “hope” are naturally seen as revealing the anchoring and adjustment bias. When subjects are presented with the Typo vignette and the evidence-seeking question, it is natural for them to anchor on the number of times Peter should proofread his paper before he turns it in. The responses here will naturally be sensitive to stakes. This is because the question is just one about rational decisions which, under classical decision theory, is sensitive to the costs of being wrong (stakes). Pinillos (2012) shows that the answer to this question was the same as the answer to the Knows question.³³

According to Epley and Gilovich’s (2006) processing model, subjects’ responses will gravitate toward the anchor as a result of a satisficing process: Subjects will form a series of hypotheses about possible responses starting from the anchor and moving away, but terminating when a plausible response is found. We believe that when subjects are presented with the Hopes question, they are anchoring, satisficing, and terminating soon at or near the anchor. This explains Buckwalter and Schaffer’s data. We do not think that subjects are adjusting enough. Subjects are anchoring to the natural answer to the normative question “how many times should Peter proofread” and not adjusting enough to properly answer the Hopes question. To be clear, the reason we think subjects are making a performance error is that their answers seem obviously mistaken. And the reason why we think the error arises from insufficient adjustment is that this seems to be the best explanation of their mistake.

Although we think that subjects presented with the Hopes question in the Buckwalter and Schaffer study are not adjusting enough, we do not believe that the same bias is at play when subjects are presented with the Knows prompt. Epley and Gilovich (2006) show that motivation or willingness to spend mental effort can correct for the anchoring bias in their satisficing model.³⁴ If there were an anchoring bias involved in answering the knowledge question, we would expect less reflective and more reflective participants to give different responses to the Knows prompt. But Pinillos (2012) reported that reflectiveness did not affect answers to the Knows prompts. So we do not think that the anchoring and adjustment bias is at play for the Knows questions.

We now consider Buckwalter and Schaffer's final critique of the evidence-seeking experiments. They developed a pair of modified typo vignettes, "two reads," with essentially the same content except that it is stipulated that Peter has already checked for typos twice. Subjects are then asked for their level of agreement with the statement: "Peter knows that there are no typos in his paper." Recall now that in the original high-stakes evidence-seeking experiment subjects said, on average, that Peter had to proofread five times before he knows there are no typos (in low stakes, subjects said, on average, that Peter had to proofread two times). Given this, we would expect in "two reads" that the level of agreement with the knowledge statement would differ across conditions. And we would also expect people to disagree with the statement in the high-stakes condition. However, this is not what was discovered. In a between-subjects study, Buckwalter and Schaffer (2013) discovered that people tended to agree with this statement for both high- and low-stakes conditions. And more importantly, there are no statistically significant differences across the conditions. This result is trouble for the original evidence-seeking experiments. Buckwalter and Schaffer take this result to further confirm their hypothesis that the evidence-seeking probes are not telling us anything about how people deploy "knowledge."

We will consider two responses which put some pressure on Buckwalter and Schaffer's criticism. First, we think that their vignettes trigger an implicature that Peter is satisfied with his proofreading after proofreading twice. Here's their exact wording: "Peter is naturally a pretty good speller, plus he has a dictionary with him which he has already used to check the paper carefully, twice over."

According to Grice's maxim of Quantity, conversational participants should be as informative as required. If in the story, Peter in fact counts three times, it would be a violation of Quantity if the story just said he counted twice. For this reason we think it is plausible that subjects read the story as implicating that Peter counted exactly twice. Now if subjects think that counting twice is enough to satisfy Peter, then they would likely think that this must be enough for him to know there are no typos. After all, Peter is aware of what is at stake and he is in the best position to judge how carefully he was proofreading. Subjects may interpret the high-stakes vignette as involving two *very careful reads* of the story, and this may be enough for Peter to know there are no typos. If this is right, Buckwalter and Schaffer's experiment does not impugn the evidence-seeking probes.

For the second response we revisit the anchoring effect. The crucial difference between the Two-Counts probe and the evidence-seeking experiments is that in the former, the anchor (which is *two counts*) is provided by the experimenter. This is important. Research indicates that the satisficing process, which leads to insufficient adjusting, best applies to cases which involve self-generated anchors.³⁵ When the anchor is given, the process may be one of selective accessibility.³⁶ According to this processing model, agents will test a hypothesis naturally given by the anchor. This process of testing the hypothesis will cause subjects to selectively access information about the target that is consistent with the hypothesis. And so make it more likely to agree with the hypothesis. For example, Mussweiler et al. (2004) report on an experiment where subjects are asked to judge whether the percentage of African nations who are members of the United Nations are higher or lower than 65 percent. When they do so "they selectively retrieve knowledge from memory that is consistent with this assumption (e.g., 'Africa is a huge continent,' 'There are more African nations than I can keep in mind,' etc.)" (p. 192). Hence their estimation of the relevant percentage will be closer to 65 percent than it would otherwise be since they will access information about Africa consistent with the given anchor number.

We think a similar bias may be at play in Two Counts. Agents consider the hypothesis that Peter knows there are no typos after proofreading just twice. Doing so causes them to construe the vignette so that in fact the proofreads are of sufficient quality to engender knowledge after just proofreading twice.

We think that either the Gricean or the Anchoring mechanisms are at play in the Two-Counts probe. If this is right, then we can meet Buckwalter and Schaffer's criticisms. Evidence-seeking probes still give us good reason to accept (H) and AIK.

9 Conclusion

In this chapter we reported on some new experiments that support (H), folk attributions of knowledge are sometimes sensitive to practical interests; and AIK, a substantive thesis in epistemology. We use evidence-seeking experiments to show this as well as more traditional measures. In addition, we give some evidence that people accept principles like ACTION and RKP which dovetail with AIK. Finally, we address some objections raised in the literature against evidence-seeking experiments and find the objections to be problematic. As a final point, we add that (contrary to initial impressions perhaps) there is a sense in which our method resembles traditional methods in philosophy. We rely on judgments about cases to get at important philosophical questions. Where we depart from tradition is that we use an experimental approach to collecting and analyzing those judgments.

Notes

- 1 Pinillos, at Arizona State University, is the primary author. Simpson is at the Graduate Center of the City University of New York.
- 2 See Hawthorne (2003), Stanley (2005), Fantl and McGrath (2002, 2007), Pinillos (2012) and Weatherson (2012).
- 3 We take the premise in question to not be, in general, a psychologized entity: the premise is the content of the judgment, not the fact that people possess the intuition.
- 4 Weinberg et al. (2001) and Buckwalter and Stich (2010).
- 5 See Schwitzgebel and Cushman (2012) for evidence that philosophers are not immune to order effects.
- 6 Buckwalter (2010), Feltz and Zarpentine (2010), May et al. (2010). Phelan (forthcoming) gives evidence that practical interests don't play a role in

attributions of evidence, but doesn't address knowledge. The term "first wave" is used by Buckwalter and Schaffer (MS) in referring to these first experiments.

- 7 Pinillos (2012) and Sripada and Stanley (2012).
- 8 Buckwalter (this volume) and Buckwalter and Schaffer (MS).
- 9 A version of this problem for the experimental work is raised by DeRose (2011).
- 10 Nagel (2010).
- 11 A version of the "Awareness of Stakes" problems can be found in Pinillos (2011, 2012) and Sripada and Stanley (2012).
- 12 It may be, as James Beebe pointed out to us, that the failure to detect significance is due to the survey materials themselves (as opposed to having an insufficient number of subjects).
- 13 The probability is $(1 - \text{power})$.
- 14 Stevens (2011). *Intermediate Statistics: A Modern Approach*.
- 15 Another reason they give for not claiming (H) has been supported is that the responses to low- and high-stakes vignettes are both on the agreement side of neutral (with neutral being 4). So we don't have the switch from knowing to nonknowing responses. We do not find this reason very convincing. If the goal is to show that ordinary knowledge attributions are sensitive to practical interests, we just need to find a difference in responses across the conditions in the predicted direction. We do not also need the stronger claim that the responses must average to a certain range. In fact, a second paper from the first wave by May et al. also discovers statistically significant differences between low- and high-stakes responses. But since the responses do not switch from knowing to nonknowing, they do not take these results to support the stake sensitivity of ordinary knowledge ascriptions. Our criticism applies here as well. In fact, we think that since two of out the three first-wave studies which are taken to be trouble for (H) actually find statistically significant differences between responses to low- and high-stakes scenario, we don't think that the first-wave studies turn out to be very troubling for (H) after all.
- 16 Of course, the probability of the costs being realized is distinct from the epistemic probability of the belief being true.
- 17 In addition, right before the knowledge prompt, we asked a normative question "how many times should Jessie check the roster." We ask this question right before the knowledge prompt so that people won't take the knowledge prompt to implicate this normative prompt. We think people won't make the implicature because it would violate the principle of relevance. Moreover, this normative question will serve a further purpose concerning the connection between knowledge and action. Study 3 addresses this issue.

- 18 Writing “never” does not necessarily reveal an error on the part of subjects. It may indeed reveal a skeptical attitude toward the possibility of knowledge. Giving people an option to write “never” prevents people from writing arbitrarily large numbers, as was found in Pinillos (2012).
- 19 Note that these are ignorant probes so that the probability of the possible harm is not known by the subject. It is instead known by the attributor. If probabilities are understood subjectively, an interesting question would arise as to whether the possible sensitivity to probabilities in ignorant cases must be understood as supporting attributor sensitivity (contextualism), as opposed to subject sensitivity. We do not explore this issue here.
- 20 We insert a knowledge denial here so that we make it natural for subjects to evaluate Peter’s response to this challenge.
- 21 Air: $t(53) = 0.462$, $p = 0.646$.
- 22 Fantl and McGrath (2009).
- 23 The normative prompt for COIN was this: “Peter’s friend told Peter he should count the coins again. Do you think that Peter should count the coins again?” The normative prompt for AIR was “Do you think Jason should look through the entire roster at least one more time?” The normative prompt for BRIDGE was “Do you think that John should just cross the bridge.” For BRIDGE and AIR, subjects were not given a “Neutral” option in responding, it was just a binary “Yes” and “No.”
- 24 Although in Experiment 2 those who think that Peter should count the coins again are less likely to say he knows that there are 134 coins in the jar (thereby supporting folk acceptance of ACTION), we must note that 25 percent of our participants thought that (1) Peter should count the coins again and (2) nonetheless still think that Peter knows that there are 134 coins in the jar. All of those people are failing to act in accordance with ACTION. Since this number does not approach a majority, we cannot say that it gives us good evidence against ACTION. Yet, the number is high enough to merit further investigation.
- 25 We replicated their studies for “Hopes.”
- 26 The discussion was prompted by a response from Pinillos in the Experimental Philosophy blog and in personal communication. They do not endorse this response.
- 27 Low-Stakes Hopes ($M = 3.81$, $SD = 0.402$), Low-Stakes Knowledge ($M = 3.95$, $SD = 0.218$), $t(30.828) = 1.43$, $p = 0.163$. High-Stakes Hopes ($M = 3.85$, $SD = 0.489$), High-Stakes Knowledge ($M = 3.95$, $SD = 0.229$), $t(37) = 0.788$, $p = 0.433$.

- 28 We instructed participants to read both questions before writing down their answers.
- 29 One may reply that in this experiment only the Knows question gets the modal reading but the hopes question doesn't. This would explain why subjects give different responses to these questions. It would also explain why there is a stakes effect for Knows but not for Hopes. This response is ad hoc. We see no reason to think that subjects would give the modal reading to one but not the other question, in the cases when they are given side by side. Note that there were no ordering effects (see prior footnote) so one can't argue that order forces a particular reading.
- 30 Tversky, A., and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124–30.
- 31 Mussweiler, T., Englich, B., and Strack, F. (2004) Anchoring effect. In R. Pohl (ed.), *Cognitive illusions: A handbook of fallacies and biases in thinking, judgement, and memory* (pp. 183–200). London, UK: Psychology Press.
- 32 This anchor point may or may not be self-generated. Epley and Gilovich (2006) "The Anchoring-and-Adjustment Heuristic. Why the Adjustments Are Insufficient." Nicholas Epley and Thomas Gilovich; *Psychological Science*, 2006, 17(4), pp. 311–18. They argue that self-generated anchors give rise to a satisficing mechanism. More on this below.
- 33 This was verified for the "coin counting" experiment. It wasn't tested for the Typo cases.
- 34 Epley and Gilovich (2006) show that participants who are under the influence of alcohol score lower in the "Need For Cognition" test and those who are under heavy cognitive load are more likely to display the bias. The authors also argue that this phenomenon only holds for self-generated anchors but not for provided anchors. The bulk of research on anchoring involves experimenter-provided anchors.
- 35 See Epley and Gilovich (2006) and Mussweiler, T., Englich, B., and Strack, F. (2004) Anchoring effect. In R. Pohl (ed.), *Cognitive illusions: A handbook of fallacies and biases in thinking, judgement, and memory* (pp. 183–200). London, UK: Psychology Press.
- 36 Mussweiler, (1997). A selective accessibility model of anchoring: Linking the anchoring heuristic to hypothesis-consistent testing and semantic priming. *Psychologia Universalis* (Vol. 11), Lengerich, Germany: Pabst; Mussweiler and Strack, (1999a) Hypothesis-consistent testing and semantic priming in the anchoring paradigm: A selective accessibility model. *Journal of Experimental*

Social Psychology, 35, 136–64. Mussweiler and Strack, (1999b), Comparing is believing: A selective accessibility model of judgmental anchoring. In W. Stroebe and M. Hewstone (eds), *European review of social psychology* (Vol. 10, pp. 135–67). Chichester, UK: Wiley. Strack and Mussweiler, (1997). Explaining the enigmatic anchoring effect: Mechanisms of selective accessibility. *Journal of Personality and Social Psychology*, 73, 437–46.

References

- Buckwalter, W. (2010), “Knowledge isn’t closed on saturday: A study in ordinary language.” *Review of Philosophy and Psychology*, 1, 395–406.
- Buckwalter, W. and Schaffer, J. (2013), “Knowledge, stakes, and mistakes.” *Noûs*, 47(3).
- Buckwalter, W. and Stich, S. (2010), “Gender and philosophical intuition.” Available at SSRN: <http://ssrn.com/abstract=1683066>.
- DeRose, K. (2011), “Contextualism, contrastivism, and x-phi surveys.” *Philosophical Studies*, 156, 81–110.
- Epley, N. and Gilovich, T. (2006), “The anchoring and adjustment heuristic: Why adjustments are insufficient.” *Psychological Science*, 17, 311–18.
- Fantl, J. and McGrath, M. (2002), “Evidence, pragmatics, and justification.” *Philosophical Review*, 111, 67–94.
- . (2007), “On pragmatic encroachment in epistemology.” *Philosophy and Phenomenological Research*, 75, 558–89.
- . (2009), *Knowledge in an Uncertain World*. New York: Oxford University Press.
- Feltz, A. and Zarpentine, C. (2010), “Do you know more when it matters less?” *Philosophical Psychology*, 23, 683–706.
- Goldman, A. (2007), “Philosophical intuitions: Their target, their source and their epistemic status.” *Grazer Philosophische Studien*, 74, 1–26.
- Hawthorne, J. (2003), *Knowledge and Lotteries*. Oxford University Press: Oxford.
- Hawthorne, J. and Stanley, J. (2008), “Knowledge and action.” *Journal of Philosophy*, 105, 571–90.
- May, J., Sinnott-Armstrong, W., Hull, J. G. and Zimmerman, A. (2010), Practical interests, relevant alternatives, and knowledge attributions: An empirical study. *Review of Philosophy and Psychology*, 1, 265–73.
- Mussweiler, T. (1997), *Psychologia Universalis: Vol. 11. A Selective Accessibility Model of Anchoring: Linking the Anchoring Heuristic to Hypothesis-Consistent Testing and Semantic Priming*. Lengerich (Germany): Pabst.

- Mussweiler, T. and Strack, F. (1999a), "Hypothesis-consistent testing and semantic priming in the anchoring paradigm: A selective accessibility model." *Journal of Experimental Social Psychology*, 35, 136–64.
- (1999b), "Comparing is believing: A selective accessibility model of judgmental anchoring," in W. Stroebe and M. Hewstone (eds), *European Review of Social Psychology*, Vol. 10. Chichester, UK: Wiley, pp. 135–67.
- Mussweiler, T., Englich, B. and Strack, F. (2004), "Anchoring effect," in R. Pohl (ed.), *Cognitive Illusions: A Handbook of Fallacies and Biases in Thinking, Judgement, and Memory*. London: Psychology Press, pp. 183–200.
- Nagel, J. (2007), "Epistemic intuitions." *Philosophy Compass*, 2, 792–819.
- (2010), "Knowledge ascriptions and the psychological consequences of thinking about error." *Philosophical Quarterly*, 60, 286–306.
- Phelan, M. (forthcoming), "Evidence that stakes don't matter for evidence." *Philosophical Psychology*.
- Pinillos, N. Á. (2011), "Some recent work in experimental epistemology." *Philosophy Compass*, 6, 675–88.
- (2012), "Knowledge, experiments and practical interests," in J. Brown and M. Gerken (eds), *Knowledge Ascriptions*. Oxford: Oxford University Press.
- Schaffer, J. (2006), "The irrelevance of the subject: Against subject sensitive invariantism." *Philosophical Studies*, 127, 87–107.
- Schwitzgebel, E. and Cushman, F. (2012), "Expertise in moral reasoning? Order effects on moral judgment in professional philosophers and non-philosophers." *Mind & Language*, 27, 135–53.
- Sripada, C. and Stanley, J. (2012), "Empirical tests of interest-relative invariantism." *Episteme*, 9, 3–26.
- Stanley, J. (2005), *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Stevens, J. (2007), *Intermediate Statistics: A Modern Approach* (2nd edn). Mahwah, NJ: Erlbaum.
- (2011). *Intermediate Statistics: A Modern Approach* (3rd edn). Mahwah, NJ: Erlbaum.
- Strack, F. and Mussweiler, T. (1997), "Explaining the enigmatic anchoring effect: Mechanisms of selective accessibility." *Journal of Personality and Social Psychology*, 73, 437–46.
- Tversky, A. and Kahneman, D. (1974), "Judgment under uncertainty: Heuristics and biases." *Science*, 185, 1124–31.
- Weatherson, B. (2012), "Knowledge, bets and interests," in J. Brown and M. Gerken (eds), *Knowledge Ascriptions*. Oxford: Oxford University Press, pp. 75–103.
- Weinberg, J., Nichols, S. and Stich, S. (2001), "Normativity and epistemic intuitions." *Philosophical Topics*, 29, 429–60.
- Williamson, T., (2000), *Knowledge and its Limits*. Oxford: Oxford University Press.