Experimental Evidence that Knowledge Entails Justification*

Alexandra Nolte, David Rose, and John Turri

alexandramnolte@gmail.com

rose0david@gmail.com

john.turri@gmail.com

Abstract: A standard view in philosophy is that knowledge entails justification. Yet recent research suggests otherwise. We argue that this admirable and striking research suffers from an important limitation: participants were asked about knowledge but not justification. Thus it is possible that people attributed knowledge partly because they thought the belief was justified. Perhaps though, if given the opportunity, people would deny justification while still attributing knowledge. It is also possible that earlier findings were due to perspective taking. This paper reports further research that directly addresses these questions. Our findings support the hypothesis that knowledge entails justification on the ordinary view.

Keywords: Justification, Knowledge, Perspective Taking, Folk Epistemology

The view that knowledge entails justification is widely endorsed by philosophers (Ichikawa and Steup, 2012). It might trace as far back to Plato suggesting that justification is

* This is the penultimate version of a paper forthcoming in T. Lombrozo, S Nichols & J. Knobe (Eds.), *Oxford Studies in Experimental Philosophy*. Please cite the final, published version if possible. Data and materials available here: https://osf.io/3vds6/

what "tethers" true belief to knowledge (Plato, 2002). There is disagreement over the level of justification required for knowledge. For example, infallibilists claim that justification should guarantee truth, whereas fallibilists claim otherwise (BonJour, 2002, p. 38 ff). There is also disagreement over the structure of justification, with views ranging from foundationalism (e.g., Descartes, 1637) to coherentism (e.g., Davidson, 1986) and infinitism (e.g., Klein, 2007). Moreover, there is disagreement over whether justification is a property of internal states (BonJour, 1985), as evidentialists maintain (e.g., Feldman and Conee, 1985, 2001) and reliabilists deny (e.g., Goldman, 1979). But despite disagreement over the nature of justification, many epistemologists nonetheless maintain that knowledge entails justification.

Entailment from knowledge to justification features prominently in the justified true belief analysis of knowledge, is retained in many post-Gettier analyses of knowledge, and even survives in Williamson's (2002) primitivism. Much like the view that knowledge entails belief (see e.g., Rose and Schaffer, 2013; Buckwalter, Rose and Turri, 2015) and the view that knowledge entails truth (see e.g. Buckwalter, 2014), the view that knowledge entails justification is a standard assumption in contemporary epistemology.[1] However, researchers have recently argued that new evidence undermines the assumption (Sackris and Beebe, 2014; Gonnerman et al, forthcoming). Their arguments are based on new evidence, deserve attention, and should prompt a reappraisal.

Sackris and Beebe (2014) set out from Sartwell's (1991; 1992) claim that knowledge is true belief, empirically testing a range of cases suggesting that a protagonist can know some

[1] One important question concerns the nature of the entailment from knowledge to justification. It is doubtful that the entailment is analytic in the way that "All bachelors are married men" is and it is also doubtful that the entailment is either logical, such as if A=B and B=C, then A=C, or metaphysical, such as water is H_2O. Instead, and given that the ordinary concept of knowledge is what is at issue here, we will assume that the entailment is conceptual.

proposition despite lacking justification. For instance, participants were given a case featuring an individual, Sunil, who has become highly delusional. In a bout of delusion, Sunil is told by "demons" that Hilary Clinton is the current US secretary of state. (At the time the study was conducted, Clinton was the US secretary of state.) On this basis, he comes to believe that Clinton is the current secretary of state. And his belief is true. When asked whether Sunil knew—on the basis of his delusion—that Clinton was the current secretary of state, the mean response on a 7-point scale was 4.85, which is significantly above the neutral midpoint. Thus there was a central tendency to attribute knowledge based on delusion. Assuming that delusion is not viewed as providing relevant justification, this is unexpected if knowledge entails justification on the ordinary view.

We submit that, based on these results, concluding that knowledge does not entail justification would be premature. The primary concern is that participants did not rate whether the protagonist's belief is justified. It could thus be the case that knowledge ratings were high because participants thought the belief was justified. Without asking about justification, this is a possibility that cannot be foreclosed.

Another key issue is whether earlier results might be due to perspective taking. Perspective taking occurs when a subject takes up another's perspective and imagines what seems true from their point of view (Flavell, Everett, Croft, and Flavell, 1981; Surtees, Samson, and Apperly 2016). Recent evidence suggests that perspective taking comes so naturally to humans that they engage in it automatically and unconsciously, even when doing so interferes with the task at hand (Samson, Apperly, Braithwaite, Andrews, and Scott, 2010; Surtees and Apperly, 2012). Perspective taking has arguably played an important role in explaining surprising findings concerning ordinary knowledge judgments, such as the attribution of false

knowledge (Buckwalter 2014; Turri 2014) and potential cross-cultural disagreement about the relationship between knowledge and certain forms of luck (Machery, Stich, Rose, et. al., 2017a; Machery, Stich, Rose, et. al., 2017b).

Perhaps perspective taking can also help explain the finding that participants attribute knowledge in the absence of justification (Sackris and Beebe 2014). This is amplified by that fact that protagonists in the scenarios studied typically assert that they know the relevant proposition. For instance, Sunil asserts that, "I first acquired knowledge of this fact back when I was delusional." Assertion could increase the likelihood of perspective taking by making another's perspective salient and explicit. It is typical in Gettier cases (e.g., Machery, Stich, Rose, et. al., 2015), Fake Barn cases (e.g., Colaço et al., 2014), and Bank cases (Rose, Machery, Stich, et. al., 2017). People tend to defer to the mental state reports of others and so knowledge judgments may be explained by mental state deference (Turri, 2017), which could result from perspective taking. When coupled with a wide range of research suggesting that knowledge is the norm of assertion on the ordinary view (Turri, 2016a, 2017, 2018), this lends further credence to the possibility that earlier studies contained a mix of crucial ingredients that promoted perspective taking. Contrary to that suggestion, some researchers re-tested some earlier cases while attempting to inhibit perspective taking (Gonnerman et al., forthcoming). Though they found some evidence of perspective taking, they judge it to be "not all that extensive." However, like Sackris and Beebe, they did not collect judgments about justification.

So we take it to be an open question whether knowledge entails justification on the ordinary view. Our strategy for addressing this is to simply ask about both knowledge and justification. Our findings from experiment 1 conflict with previous findings challenging entailment (Sackris and Beebe 2014; Gonnerman et al., forthcoming): when asking about both

knowledge and justification, we found that justification attributions significantly exceeded knowledge attributions across a range of scenarios. This supports the view that knowledge entails justification.[2] But participants in experiment 1 were only asked whether the protagonist's belief is justified. It could be that participants interpreted justification in a broad way, to include not only epistemic norms but also moral and prudential norms. Study 2 sharpens the justification question, specifically targeting *justification by evidence*. When using this more pointed characterization, we found evidence of knowledge attribution exceeding justification attribution, albeit across a more limited range of scenarios. This speaks against entailment. However, these results could be due to perspective taking. In study 3, we address the question of whether our earlier findings might be due to perspective taking. When using response options that inhibited perspective taking, we found a very different pattern. Participants now tended to deny knowledge. We conclude that, overall, our findings support the hypothesis that knowledge entails justification.

1. Experiment 1

1.1. General methods

The following statements are true of all studies reported here. All manipulations, measures, and exclusion criteria are reported. All participants were adult residents of the United States. We recruited and tested people using an online platform of Amazon Mechanical Turk (https://www.mturk.com), TurkPrime (Litman, Robinson, and Abberbock 2017), and Qualtrics

[2] To clarify: our claim is not that the evidence *proves* entailment from knowledge to justification. Hence we use "support" throughout.

(https://www.qualtrics.com). Participants completed a brief demographic questionnaire after testing. We used R 3.5.1 for all analyses (R Core Team 2018). All stimuli and data are available through an Open Science Foundation project (https://osf.io/3vds6/). All studies were pre-registered.

1.2. Method

We decided in advance to recruit 50 participants per condition, plus a few extra as a precaution against attrition (see pre-registration).

1.2.1. Participants

Two hundred sixty-one people participated in the study. Their mean age was 35.12 years (range = 19-71, sd = 10.99), 33% (85 of 261) were female, and 95% reported native competence in English.

1.2.2. Materials and procedure

Participants were randomly assigned to one of five scenario conditions (root, accusation, cheating, dream, 8-ball). Participants read a brief scenario and responded to two test statements. The conditions differed in the scenario that participants read. In each scenario, the agent forms a belief and then participants rated whether the agent was justified in the believing the proposition, and whether the agent knows the proposition. The scenarios differed in the agent's role, the manner of belief formation, and the belief's content. The scenario factor was not of independent theoretical interest and was included to support generalization of the results beyond the specific stimuli studied here (Clark 1973; Baayen, Davidson, and Bates 2008; Judd, Westfall, and Kenny 2012). We adapted two scenarios from prior research on the relationship between justification

and knowledge attributions (Sackris and Beebe 2014), and we developed three more based on discussions in introductory epistemology courses and textbooks (e.g. BonJour (2002); Feldman (2003); Fumerton (2006)). The rationale for testing these scenarios is that they have been identified as potentially intuitive cases of knowledge without justification. Here is an example of a scenario with test statements:

> (Accusation) Andy's daughter has been accused of a serious crime. She lacks an alibi and the police have gathered strong evidence against her. But, even after considering it all, Andy feels that she must be innocent. And in fact he is correct: his daughter is innocent.

1. Andy is justified in believing that his daughter is innocent.
2. Andy knows that his daughter is innocent.

The test statements appeared together beneath the story and their order was rotated randomly. Responses were collected on a standard 6-point Likert scale, anchored with "strongly disagree" (= 1), "disagree," "slightly disagree," "slightly agree," and "strongly agree" (= 6), arranged left-to-right on the participant's screen.

1.2.3. Data analysis

Our principal research question was whether knowledge attribution would significantly exceed justification attribution, either overall or in specific scenarios. To answer this question, we conducted a linear mixed effects analysis on participant response and followed up with appropriate t-tests. The model included judgment type (within-subjects: justification, knowledge)

and participant age and sex as fixed effects. It also included random intercepts for scenario and participants nested within scenario.

1.3. Results

The linear mixed effects analysis revealed an effect of judgment on participant response (see Table 1 and Figure 1) and an unpredicted marginal effect of participant sex. A follow-up paired samples t-test revealed that mean response was higher for the justification attribution (M = 3.99, SD = 1.56) than for the knowledge attribution (M = 3.73, SD = 1.64), t(260) = 2.89, MD = 0.26 [0.08, 0.44], p = .004, d = 0.18. An independent samples t-test revealed that mean response was lower for males (M = 3.75, SD = 1.63) than for females (M = 4.08, SD = 1.52), t(355.24) = -2.28, MD = -0.33 [-0.62, -0.05], p = .023, d = -0.21. A likelihood ratio test comparing the fully specified mixed model to a comparable model without scenario revealed that scenario significantly affected participant response, log likelihood = -918.7, $\chi^2(1) = 19.91$, $p < .001$ (see Figure 1).

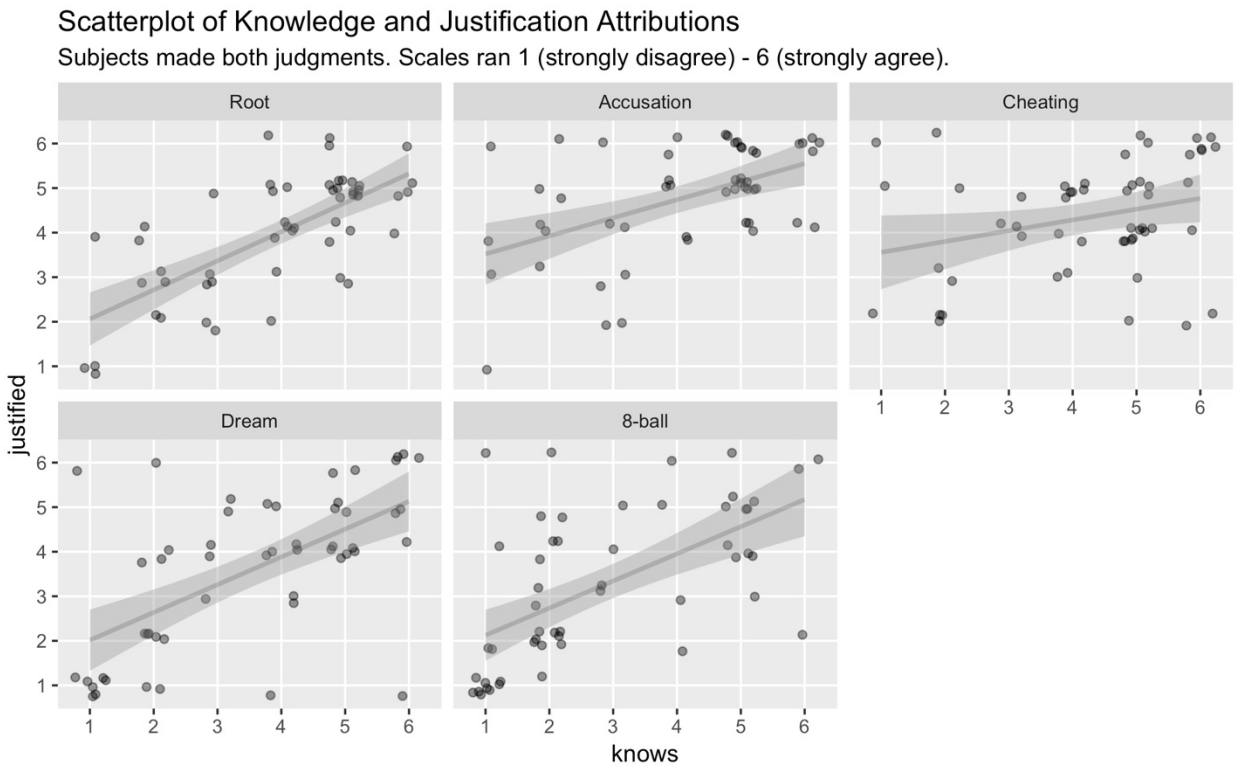

Experiment 1. Figure 1. Scatterplot of knowledge and justification attributions (within-subjects) across five scenarios (between-subjects). Scales ran 1 ("strongly disagree") – 6 ("strongly agree") with no neutral option. Dots are jittered to avoid over-plotting.

Experiment 1. Table 1. Analysis of variance for the linear mixed model's fixed effects.

	Sum of squares	Numerator df	Denominator df	F	p
Judgment	8.858	1	261.00	8.373	.004
sex	3.603	1	256.03	3.405	.066
age	0.257	1	256.72	0.243	.623

Follow-up one sample t-tests revealed that mean justification attribution was significantly above the midpoint in three scenarios but did not differ from the midpoint in the other two (see Table 2). Mean knowledge attribution was significantly above the midpoint in the same three conditions, no different from the midpoint in one, and significantly below it in another. In every condition, mean justification attribution was at least numerically higher than mean knowledge attribution (see Figure 1).

Experiment 1. Table 2. One sample t-tests comparing justification and knowledge attributions to the theoretical neutral midpoint (= 3.5) in all five scenarios.

Scenario	Judgment	n	mean	SD	df	t	p	d
Root	justified	52	3.98	1.34	51	2.60	.012	0.36
Root	knows	52	3.94	1.46	51	2.18	.034	0.30
Accusation	justified	52	4.77	1.23	51	7.44	< .001	1.03
Accusation	knows	52	4.08	1.51	51	2.76	.008	0.38
Cheating	justified	53	4.34	1.29	52	4.76	< .001	0.65
Cheating	knows	53	4.23	1.49	52	3.55	.001	0.49
Dream	justified	52	3.60	1.74	51	0.40	.692	0.06
Dream	knows	52	3.54	1.73	51	0.16	.873	0.02
8-ball	justified	52	3.25	1.70	51	-1.06	.295	-0.15
8-ball	knows	52	2.85	1.67	51	-2.82	.007	-0.39

1.4. Discussion

This experiment examined whether people attribute knowledge without justification across a range of scenarios that are claimed to be intuitive examples of knowledge without justification. If these are intuitive examples of knowledge without justification, then there should be a detectable central tendency for people's knowledge attributions to exceed their justification attributions. By contrast, if they are not intuitive examples of knowledge without justification, then there should be no such central tendency. Our findings support the latter prediction: on average, in a within-subjects design, justification attribution exceeded knowledge attribution. This same pattern was observed regardless of whether mean knowledge attribution was above or below the neutral midpoint.

One limitation of these findings concerns the wording of the justification probe, which stated simply that the agent "is justified in believing" the target proposition. However, as theorists have often noted, the proposed *justification* requirement on knowledge pertains to justification due to *evidence* or *information* available to the agent (e.g. BonJour (2002), pp. 34 ff.; Feldman (2003), p. 49; Fumerton (2006), pp. 4-5). They contrast this sort of justification with justification due to morality or prudence. Accordingly, one concern is that participants in the present study could have understood the justification probe broadly to include these sources of practical justification, rather than being restricted to justification due to evidence. The next experiment directly investigates this possibility by probing for justification in a way that specifically targets epistemic justification.[3]

2. Experiment 2

2.1. Method

We decided in advance to recruit 50 participants per condition, plus a few extra as a precaution against attrition (see pre-registration).

2.1.1. Participants

Two hundred sixty-one people participated in the study. Their mean age was 34.44 years (range = 18-68, SD = 9.27), 47% (123 of 261) were female, and 96% reported native competence in English.

[3] We acknowledge that there might be some instances of intuitive, non-evidential knowledge and that in such cases evidential justification doesn't play a role. Such cases are not our concern here. Instead our focus is on the kinds of cases that have centrally featured in the experimental literature on whether knowledge requires justification. And none of these cases feature putative instances of non-evidential knowledge.

2.1.2. Materials and procedure

Participants were randomly assigned to one of five scenario conditions (root, accusation, cheating, dream, 8-ball). The stimuli and procedures were exactly the same as in experiment 1, with the only exception being the wording of the justification probe. Whereas in experiment 1, participants rated an attribution that said the agent was "justified in believing" the target proposition, in this experiment the attribution said "the evidence justifies the agent in believing" the target proposition. Here is an illustration of the difference, focusing on the probe from the accusation scenario:

> Andy is justified in believing that his daughter is innocent. (experiment 1)
>
> The evidence justifies Andy in believing that this daughter is innocent. (experiment 2)

2.1.3. Data analysis

Our principal research question was whether knowledge attribution would significantly exceed justification attribution, either overall or in specific scenarios. To answer this question, we conducted a linear mixed effects analysis on participant response and followed up with appropriate t-tests and binomial tests on dichotomized attributions. The mixed model included judgment type (within-subjects: justification, knowledge) and participant age and sex as fixed effects. It also included random intercepts for scenario and participants nested within scenario. We dichotomized responses by counting "strongly disagree," "disagree," and "slightly disagree" as a denial (coded 0), and counting "strongly agree," "agree," and "slightly agree" as an attribution (coded 1).

2.2. Results

The linear mixed effects analysis revealed an effect of judgment on participant response (see Table 3 and Figure 2). The marginal effect of participant sex from experiment 1 did not replicate here. A follow-up paired samples t-test revealed that mean response was lower for the justification attribution (M = 3.33, SD = 1.65) than for the knowledge attribution (M = 3.8, SD = 1.6), t(260) = -3.91, MD = -0.47 [-0.7, -0.23], p = < .001, d = -0.24. A likelihood ratio test comparing the fully specified mixed model to a comparable model without scenario revealed that scenario significantly affected participant response, log likelihood = -975.48, $\chi^2(1) = 12.91$, p < .001 (see Figure 2).

Critically, paired samples t-tests revealed that mean justification attribution was significantly higher than mean knowledge attribution in the accusation and cheating conditions (see Figure 2 and Table 4). The effect size for these differences was large (Cohen's d of -0.895 and -0.946). Table 5 shows one sample t-tests comparing mean response for both attributions in

all conditions to the neutral midpoint.

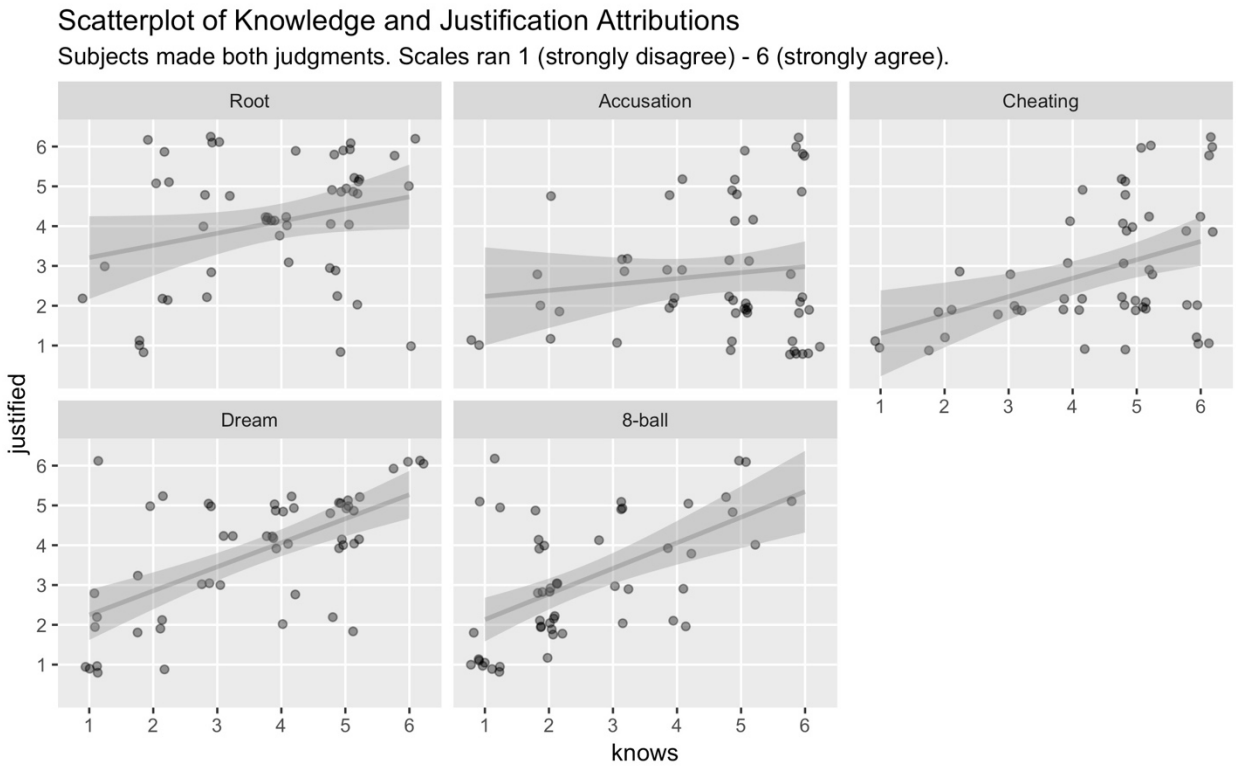

Experiment 2. Figure 2. Scatterplot of knowledge and justification attributions (within-subjects) across five scenarios (between-subjects). Scales ran 1 ("strongly disagree") – 6 ("strongly agree") with no neutral option. Dots are jittered to avoid over-plotting.

Experiment 2. Table 3. Analysis of variance for the linear mixed model's fixed effects.

	Sum of squares	Numerator df	Denominator df	F	p
Judgment	28.513	1	261.00	15.329	.000
sex	1.800	1	256.39	0.968	.326
age	0.582	1	259.06	0.313	.576

Experiment 2. Table 4. Paired sample t-tests comparing justification and knowledge attributions in all five scenarios.

Scenario	estimate	conf.low	conf.high	df	t	p	d
Root	0.192	-0.319	0.704	51	0.755	.454	0.105
Accusation	-1.830	-2.394	-1.267	52	-6.515	<.001	-0.895
Cheating	-1.529	-1.984	-1.075	50	-6.755	<.001	-0.946
Dream	0.226	-0.139	0.592	52	1.244	.219	0.171
8-ball	0.596	0.213	0.979	51	3.125	.003	0.433

Experiment 2. Table 5. One sample t-tests comparing justification and knowledge attributions to the theoretical neutral midpoint (= 3.5) in all five scenarios.

Scenario	Judgment	n	mean	SD	df	t	p	d
Root	justified	52	4.10	1.62	51	2.65	.011	0.37
Root	knows	52	3.90	1.38	51	2.12	.039	0.29
Accusation	justified	53	2.77	1.65	52	-3.21	.002	-0.44
Accusation	knows	53	4.60	1.45	52	5.56	<.001	0.76
Cheating	justified	51	2.88	1.57	50	-2.81	.007	-0.39
Cheating	knows	51	4.41	1.40	50	4.64	<.001	0.65
Dream	justified	53	3.81	1.51	52	1.50	.139	0.21
Dream	knows	53	3.58	1.57	52	0.39	.696	0.05
8-ball	justified	52	3.10	1.55	51	-1.88	.066	-0.26
8-ball	knows	52	2.50	1.34	51	-5.40	<.001	-0.75

We also examined dichotomized response options (see the data analysis section for an explanation of the coding procedure). Overall, 25% of participants attributed knowledge but denied justification, which is not different from chance rates (= .25). However, as reported above, the accusation and the cheating scenarios were importantly different from the others. Including only those two conditions for further analysis, 49% attributed knowledge but denied

justification, which significantly exceeds the percentage that attributed knowledge and justification (29%), binomial test, n = 104, k = 51, p < .001, test proportion = 0.29. Moreover, among those who attributed knowledge in these two conditions, 63% denied justification, which exceeds chance rates, binomial test, n = 81, k = 51, p = .026, test proportion = 0.5.

2.3. Discussion

This experiment again examined whether people attribute knowledge without justification across five scenarios. There was a critical difference between the justification probe used in this experiment compared to experiment 1. In experiment 1, participants rated whether the agent was "justified in believing" the target proposition, whereas in this experiment participants rated whether "the evidence justified" the agent in believing the proposition. Call this the difference between *generic justification* and *evidential justification*. If a scenario is an intuitive example of knowledge without evidential justification, then there should be a detectable central tendency for people's knowledge attributions to exceed their justification attributions. In two of the scenarios, mean justification attribution was significantly below the neutral midpoint, whereas mean knowledge attribution was significantly above the neutral midpoint. The effect size for these differences was large. A similar pattern emerged when we dichotomized the scaled attributions. Therefore, the present results support the hypothesis that knowledge does not conceptually entail evidential justification. However, in light of the findings from our first experiment, the present results also suggest that the ordinary concept of justification for beliefs is broader than evidential justification. Accordingly, although our findings thus far undermine the hypothesis that knowledge entails evidential justification, they are nevertheless consistent with the hypothesis that knowledge entails generic justification, which is broader than evidential justification.

One crucial question, however, is whether the tendency to ascribe knowledge without evidential justification is due to perspective taking. That is, perhaps participants are imagining what seems true from the protagonist's point of view. If this is not occurring, then we have a clear case of knowledge not requiring evidential justification; if it is, then perhaps knowledge does indeed entail evidential justification. Our next study addresses this question.

3. Experiment 3

3.1. Method

We decided in advance to recruit 100 participants, plus a few extra as a precaution against attrition (see pre-registration).

3.1.1. Participants

One hundred five people participated in the study. Their mean age was 39.25 years (range = 21-72, SD = 12.71), 56% (59 of 105) were female, and 91% reported native competence in English.

3.1.2. Materials and procedure

Participants read the accusation scenario from earlier experiments and responded to a single test item:

Andy _____ that his daughter is innocent. (only thinks he knows / actually does know)

3.1.3. Data analysis

Our principal research question was whether knowledge attribution would exceed chance rates. To answer this question, we conducted a binomial test against chance. We counted "only thinks he knows" as a knowledge denial (coded 0), and we counted "actually does know" as a knowledge attribution (coded 1).

3.2. Results

The rate of knowledge attribution was 24.8%, which is significantly below chance rates, binomial test, $n = 105$, $k = 26$, $p < .001$.

3.3. Discussion

This experiment examined whether the high rate of knowledge attribution in previous studies could have been due to perspective taking. The results support the conclusion that it was due to perspective taking. Two final questions are whether perspective-taking affects attributions of knowledge and justification differently, and whether inhibiting perspective-taking brings justification attribution down to at least the level of knowledge attribution across a range of scenarios. Our next, and final, study addresses these questions.

4. Experiment 4

4.1. Method

We decided in advance to recruit 50 participants per condition, plus a few extra as a precaution against attrition (see pre-registration).

4.1.1 Participants

Two hundred seventy-five people participated in the study. Their mean age was 35.76 years (range = 20-71, SD = 10.97), 40% (110 of 275) were female, and 96% reported native competence in English.

4.1.2 Materials and procedure

Participants were randomly assigned to one of the same five scenario conditions from experiment 1 (root, accusation, cheating, dream, 8-ball). The stimuli and procedures were exactly the same as in experiment 1, with the only exception being that instead of collecting responses on a 6-point

Likert scale, we used dichotomous response options that contrasted what the agent "thinks" with what is "actually" true, for both attributions. For example, here were the two attributions for the accusation scenario:

1. Andy _____ in believing that his daughter is innocent. (only thinks he is justified / actually is justified)
2. Andy _____ that his daughter is innocent. (only thinks he knows / actually does know)

4.1.3. Data analysis

Our principal research question was whether knowledge attribution would exceed justification attribution. To answer this question, we conducted a linear mixed effects analysis on participant response and followed up with appropriate proportion tests. The mixed model included judgment type (within-subjects: justification, knowledge) and participant age and sex as fixed effects. It also included random intercepts for scenario and participants nested within scenario.

4.2. Results

The linear mixed effects analysis revealed an effect of judgment on participant response (see Table 6 and Figure 3). Follow-up McNemar tests revealed that in no scenario were participants more likely to attribute knowledge than justification. In only one scenario (dream) did knowledge attribution numerically exceed justification attribution (34% compared to 27%) (see Table 7). Overall, participants were significantly more likely to attribute justification than knowledge, McNemar's test, (47% compared to 32%), $\chi^2(1) = 17.58$, p = < .001, h = 0.306.

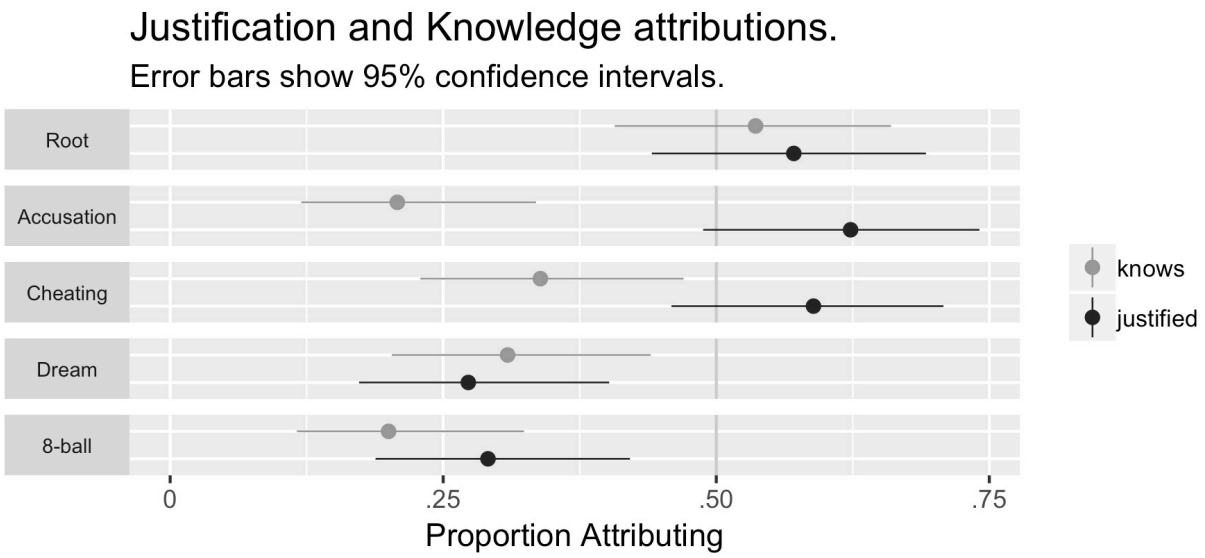

Experiment 4. Figure 3. Proportion of participants attribution knowledge and justification (within-subjects) across five scenarios (between-subjects). Error bars show 95% bootstrapped confidence intervals.

Experiment 4. Table 6. Fixed effects of the mixed linear model. Reference class for sex: male. Reference class for knows: justified.

term	estimate	standard error	Z	p
knows	-0.889	0.224	-3.976	<.001
sex	0.312	0.272	1.147	.251
age	-0.024	0.013	-1.916	.055

Experiment 4. Table 7. McNemar's tests comparing attributions of justification and knowledge in five scenarios.

Scenario	statistic	p	df	h
Root	0.056	.814	1	0.072
Accusation	15.750	<.001	1	0.873
Cheating	7.682	.006	1	0.507
Dream	0.083	.773	1	-0.080

8-ball	1.455	.228	1	0.212

4.3. Discussion

This experiment further examined whether perspective-taking might have been responsible for knowledge attribution exceeding justification attribution in earlier studies. The results support the conclusion that it was due to perspective-taking. It is worth highlighting one aspect of these results in particular. Results from experiment 2 suggested that, for instance, the accusation scenario might pose an especially strong challenge to the hypothesis that knowledge conceptually entails justification. But when we probed for attributions in a way intended to inhibit perspective-taking in the present experiment, the rate of justification attribution exceeded knowledge attribution by the largest margin in the accusation scenario specifically. Broadly similar remarks apply to the cheating scenario.

5. Conclusion

Does knowledge entail justification? Though the view that knowledge entails justification is widely held in philosophy, intriguing recent work suggests that, on the ordinary view, knowledge doesn't entail justification (Sackris and Beebe, 2014; Gonnerman et al., forthcoming). One limitation of this research is that it relied on the researchers' own interpretation of the agent's belief as not justified. Participants themselves did not rate whether the agent was justified. Thus previous findings are consistent with participants attributing knowledge partly because they also attributed justification. If so, this would then support the hypothesis that knowledge entails justification.

Accordingly, our strategy for advancing understanding of these issues was to ask about both knowledge and justification. In experiment 1, we found that when asking about both

knowledge and justification, justification attributions significantly exceeded knowledge attributions. This supports the hypothesis that knowledge entails justification. But participants in experiment 1 were only asked whether the relevant belief was "justified." Thus justification judgments might have been due to people operating with a broad notion of justification that includes epistemic, moral, and prudential norms. Experiment 2 narrowed in on epistemic justification specifically. When asked whether the "evidence justified" believing the relevant proposition, we then found evidence suggesting that knowledge may not entail justification. This supports the hypothesis that knowledge does not entail justification. We then asked whether our earlier findings might be due to perspective taking. In experiment 3, we found strong initial evidence that they were due to perspective taking. When using response options designed to inhibit perspective taking, participants tended to deny knowledge. Thus, overall, our findings support the hypothesis that knowledge entails justification.

Our results don't speak to the question of the specific view that the folk operate with concerning the nature of justification. For example, our findings do not support (or conflict with) the infallibilist claim that justification guarantees truth. Nor do they speak to the question of whether the folk operate with a specific conception of the structure of justification, such as foundationalism (e.g., Descartes, 1637), coherentism (e.g., Davidson, 1986) or infinitism (e.g., Klein, 2007). Furthermore, our findings do not shed light on whether the folk have an internalist or externalist view of justification. Though the specific view that folk operate with concerning the nature of justification remains an open question, our findings do support the widely held philosophical view that knowledge entails justification.

When asking about justification in experiment 1, the test statement asked if the protagonist was "justified." This raised the possibility that people interpreted justification

broadly to include evidential, moral, and practical considerations, the latter two of which epistemologists tend to discount as irrelevant to epistemic justification. In experiment 2, we sharpened and refined, asking whether "the evidence justifies" the belief in the proposition. Thus, it is clear there is a salient difference between asking whether a protagonist's belief is justified and whether the evidence justifies an agent in believing a proposition. Our findings demonstrate that this should be taken into consideration for further studies on justification. If one wants to probe for epistemic justification but merely asks whether a protagonist's belief is justified, then one will pick up our norms that extend beyond the epistemic realm, norms which epistemologists think are epistemically irrelevant. Asking whether the evidence justifies the relevant belief provides a more pointed way for targeting epistemic justification and suggests a more pointed way to probe for epistemic justification.

Our findings also add to the list of otherwise surprising findings on knowledge attributions explained by perspective taking. To take one example, early and influential work by Weinberg, Nichols, and Stich (2001) indicated that Westerners and East Asians viewed Gettier cases differently. The key finding: Westerners judge that protagonists in Gettier cases do not know the relevant proposition while East Asians judge that protagonists in Gettier cases do know the relevant proposition. But more recent research suggests that this cross-cultural difference is due, in part, to perspective taking (e.g., Machery, Stich, Rose, et al., 2017a; Machery, Stich, Rose, et. al., 2017b). When asked about knowledge in a way that controls for perspective taking — i.e., using "really know" and "didn't really know but only thought she knew" as response options — experimental participants across four cultures — USA, India, Japan, and Brazil — denied knowledge at rates that were both high and similar across the four groups (Machery,

Stich, Rose, et. al., 2017a). This finding has recently been extended to people across 23 countries, speaking 17 different languages (Machery, Stich, Rose, et. al., 2017b).

More generally, people are more likely to attribute knowledge when the answer options are "does know"/"doesn't know" (plain options) than when the options are "actually knows"/"only thinks he knows" (contrast options) (Cullen, 2010; Buckwalter, 2014; Turri, 2014, 2018). When using plain options, people even attribute knowledge to agents who have false beliefs. This disappears when using contrast options (Buckwalter, 2014). That's because — as John Turri (forthcoming) notes — "contrast options don't force participants to choose between responding accurately and acknowledging the agent's perspective: they enable participants to distinguish appearance from reality while minimizing the risk of appearing oblivious or insensitive."

Perspective taking isn't merely confined to epistemic judgments. It also extends to judgments about lying and assertability. In the case of lying, there are two main ways to probe for lying. One is to ask whether the protagonist "lied/didn't lie" (plain options). The other is to ask whether the protagonist "actually lied"/"only thinks he lied" (contrast options). When an agent attempts to lie but fails (the assertion is actually true) people classify it as lying when given plain options. This pattern reverses when given contrast options. This is because the latter options—the contrast options—inhibit perspective taking (Turri and Turri, 2015, forthcoming). Perspective taking plays a similar role in judgments of assertability. People are sometimes inclined to answer that a protagonist should make a false assertion supported by misleading evidence when given plain options: "should assert"/"shouldn't assert". Contrast options — "actually should assert"/"only thinks he should assert" — reverse this pattern (Turri, 2016b, 2018).

Perspective taking is clearly an issue when probing ordinary epistemic judgments. And it also extends beyond epistemology, namely, to lying and assertability. The bottom line: perspective taking can be a roadblock to genuine attribution. It infects epistemic judgments and judgments of lying and assertability. And it is pervasive. Whether it is controlled for or not can have dramatic implications for interpreting folk attribution, epistemic or otherwise. In this case, controlling for perspective taking supports the claim that knowledge does in fact, entail justification.

References

Baayen, R. H., D. J. Davidson, and D. M. Bates. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language, 59* (4), 390–412.

BonJour, L. (1985). *The Structure of Empirical Knowledge*. Cambridge, Mass.: Harvard University Press.

BonJour, L. (2002). *Epistemology: classic problems and contemporary responses*. Rowman & Littlefield.

Buckwalter, W. (2014). Factive verbs and perspective taking. *Episteme, 11*(4), 391-409.

Buckwalter, W., Rose, D., & Turri, J. (2015). Belief through thick and thin. *Noûs, 49*(4), 748-775.

Clark, Herbert H. 1973. The language-as-fixed-effect fallacy: a critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behavior, 12*, 335–59.

Colaco, D., Buckwalter, W., Stich, S., & Machery, E. (2014). Epistemic intuitions in Fake-Barn thought experiments. *Episteme, 11*(2), 199–212.

Cullen, S. 2010. Survey-Driven Romanticism. *Review of Philosophy and Psychology,* 1(2), 275–296.

Davidson, D. (1986). A Coherence Theory of Knowledge and Truth. In E. LePore (Ed.), *Truth and Interpretation* (307–319). Oxford: Blackwell.

Descartes, R. (1637). *Discourse on the Method*. (John Veitch, Trans.) Chicago: The Open Court Publishing Company.

Feldman, R. (2003). *Epistemology*. Upper Saddle River, NJ: Prentice Hall.

Feldman, R., & Conee, E. (1985). Evidentialism. *Philosophical Studies, 48*(1), 15-34.

Feldman, R., & Conee, E. (2001). Internalism Defended. *American Philosophical Quarterly, 38* (1), 1-18.

Flavell, J. H., Everett, B. A., Croft, K., & Flavell, E. R. (1981). Young children's knowledge about visual perception: further evidence for the level 1-level 2 distinction. *Developmental Psychology*, *17*(1), 99–103.

Fumerton, R. (2006). *Epistemology*. Malden, MA: Blackwell.

Goldman, A. (1979). What Is Justified Belief? In G.S. Pappas (Ed.), *Justification and Knowledge* (1-25). Dordrecht: Reidel.

Gonnerman, C., Poag, L., Redden, L., Robbins, J., & Crowley, S.(forthcoming). In Our Shoes or the Protagonist's? Knowledge, Justification, and Projection. In Lombrozo, T., Knobe, J., & Nichols, S. (Eds.), *Experimental Philosophy Volume 3*. Oxford University Press.

Holton, R. (1997). Some telling examples: A reply to Tsohatzidis. *Journal of Pragmatics, 28*, 624–628.

Ichikawa, J. J., & Steup, M. (2012). The Analysis of Knowledge. Retrieved August 1, 2015, from http://plato.stanford.edu/entries/knowledge-analysis/index.html

Judd, C. M, Westfall, J. & Kenny, D. A. (2012). Treating stimuli as a random factor in social psychology: A new and comprehensive solution to a pervasive but largely ignored problem. *Journal of Personality and Social Psychology, 103*(1), 54–69.

Klein, P. (2007). How to be an infinitist about doxastic justification. *Philosophical Studies*, *134*(1), 25-29.

Litman, L., Robinson, J., & Abberbock, T. (2017). "TurkPrime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences." *Behavioral Research Methods, 49* (2), 1–10.

Machery, E., Stich, S., Rose, D., et al (2017a). Gettier across cultures. *Nous*, 51(3), 645-664.

Machery, E., Stich, S., Rose, D., el al (2017b). The Gettier Intuition from South America to Asia. *Journal of Indian Council of Philosophical Research*, *34*(3), 517-541.

Plato. (2002). Five dialogues : Euthyphro, Apology, Crito, Meno, Phaedo. Indianapolis: Hackett Pub. Co.

R Core Team. (2018). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.

Rose, D., Machery, E., Stich, S., et. al. (2017). Nothing at stake in knowledge. *Noûs*.

Rose, D., & Schaffer, J. (2013). Knowledge entails dispositional belief. *Philosophical Studies*, *166*(1), 19-50.

Sackris, D., & Beebe J. R. (2014). Is justification necessary for knowledge? In J.R. Beebe (Ed.), *Advances in Experimental Epistemology* (175–92). New York: Bloomsbury Academic.

Sartwell, C. (1991). Knowledge is merely true belief. *American Philosophical Quarterly*, *28*, 157-65.

Sartwell, C. (1992). Why knowledge is merely true belief. *Journal of Philosophy*, *89*, 167-80.

Samson, D., Apperly, I. A., Braithwaite, J. J., Andrews, B. J., & Bodley Scott, S. E. (2010). Seeing it their way: Evidence for rapid and involuntary computation of what other people see. *Journal of Experimental Psychology: Human Perception and Performance*, *36*(5), 1255.

Surtees, A., & Apperly, I. A. (2012). Egocentrism and automatic perspective taking in children and adults. *Child Development*, *83*(2), 452–460.

Surtees, A., Samson, D., & Apperly, I. A. (2016). Unintentional perspective-taking calculates whether something is seen, but not how it is seen. *Cognition*, *148*(C), 97–105.

Turri, J. (2014). The problem of ESEE knowledge. *Ergo*, *1*(4), 101–127.

Turri, J. (2016a). Knowledge and the norm of assertion: an essay in philosophical science. Cambridge: Open Book Publishers. Retrieved from http://www.openbookpublishers.com/product/397/knowledge-and-the-norm-of-assertion--an-essay-in-philosophical-science

Turri, J. (2016b). Knowledge and assertion in “Gettier” cases. *Philosophical Psychology*, *29*(5), 759–775. http://doi.org/10.1080/09515089.2016.1154140

Turri, J. (2017). Experimental work on the norms of assertion. *Philosophy Compass*, *12*(7), e12425. http://doi.org/10.1002/ajp.22097

Turri, J. (2017). Epistemic contextualism: An idle hypothesis. *Australasian Journal of Philosophy*, 95, 141-156.

Turri, J. (2018). Revisiting norms of assertion. *Cognition*, *177*, 8–11. http://doi.org/10.1016/j.cognition.2018.03.023

Turri, J. (forthcoming). Truth, perspective and norms of assertion. *Synthese*.

Turri, A. & Turri, J. 2015. The truth about lying. *Cognition*, 138(C), 161–168.

Turri, A. & Turri, J. (forthcoming). Lying, fast and slow. *Synthese*.

Williamson, T. (2002). *Knowledge and its Limits*. Oxford: Oxford University Press on Demand.