

EPISTEMIC LUCK, KNOWLEDGE-HOW, AND INTENTIONAL ACTION

CARLOTTA PAVESE

Sage School of Philosophy, Cornell University

PAUL HENNE

Department of Philosophy, Lake Forest College

BOB BEDDOR

Department of Philosophy, National University of Singapore

Epistemologists have long believed that epistemic luck undermines propositional knowledge. Action theorists have long believed that agentive luck undermines intentional action. But is there a relationship between agentive luck and epistemic luck? While agentive luck and epistemic luck have been widely thought to be independent phenomena, we argue that agentive luck has an epistemic dimension. We present several thought experiments where epistemic luck seems to undermine both knowledge-how and intentional action and we report experimental results that corroborate these judgments. These findings have implications for the role of knowledge in a theory of intentional action and for debates about the nature of knowledge-how and the significance of knowledge representation in folk psychology.

1. Introduction

Epistemologists have long believed that epistemic luck undermines propositional knowledge (or *knowledge-that*). Consider one of Gettier's original (1963) cases: the president of the company tells Smith that Jones would get the job, and he saw Jones putting ten coins in his pocket. On this basis, Smith believes that whoever has ten coins in his pocket will get the job. However, Smith's belief is

Contact: Carlotta Pavese <cp645@cornell.edu>
Paul Henne <phenne@mx.lakeforest.edu>
Bob Beddor <rbeddor@gmail.com>

true only because—unbeknownst to him—he himself will get the job and also happens to have ten coins in his pocket. Though true and justified, Smith’s belief is too lucky to count as knowledge. For yet another standard example of epistemic luck, consider lottery cases (Williamson 2000; Hawthorne 2004; Pritchard 2005). Lottie believes, on statistical grounds, that her ticket is a loser. Though true and justified, Lottie’s belief is too lucky to qualify as knowledge. In a slogan: too lucky a belief cannot amount to knowledge. We will refer to the sort of luck that is incompatible with knowledge as *epistemic luck*.¹

Action theorists have long believed that agentive luck undermines intentional action. Consider amateur success: if a novice at darts hits the bullseye on their first try, their success does not seem intentional (Mele & Moser 1994; Malle & Knobe 1997; Knobe 2003).² Or consider a lottery loss: losing a fair lottery does not seem intentional, even if one wants to lose (Malle & Knobe 1997; Gibbons 2001). In a slogan: actions that are too lucky cannot be intentional. We will refer to the sort of luck that is incompatible with intentional action as *agentive luck*.

For the most part, researchers have assumed that these are separate phenomena. For example, many action theorists maintain that know-how or skill is an anti-luck condition on intentional action—that is, that one must know how to φ in order to intentionally φ (Heider 1958; Mele & Moser 1994; Malle & Knobe 1997; Setiya 2012). However, these theorists do not understand knowledge-how or skill as having much to do with propositional knowledge. Indeed, some explicitly deny that intentional action or knowledge-how are sensitive to the presence of epistemic luck (Poston 2009; Cath 2011; Cath 2015).

Against these views, we argue that epistemic luck and agentive luck are related in an important way. We present new cases where epistemic luck seems to undermine both knowledge-how and intentional action and we report experiments which corroborate these judgments. We then outline the implications of our discussion for various debates at the intersection of epistemology and action theory—the debate about the place for knowledge in a theory of intentional action, the debate over the nature of knowledge-how, and the debate over the significance of knowledge representation in folk psychology.

1. Epistemologists often talk of epistemic luck more generically to describe any of a number of ways in which it can be accidental, coincidental, or fortuitous that a person has a true belief. Our use of ‘epistemic luck’ is more restrictive and coincides specifically with ‘veritic luck’—the special kind of epistemic luck that most epistemologists take to be incompatible with knowledge (cf. Engel 1992; Pritchard 2005; Riggs 2007; Pritchard 2021).

2. Both Malle and Knobe (1997) and Knobe (2003) report experimental results that corroborate the claim that intentional action is incompatible with beginners’ luck. There is no denying, of course, that the novice might *intend* to hit the bullseye on their first try. After all, it is commonplace in action theory that one might intend to φ , succeed at φ -ing on that intention, yet fail to φ intentionally. Deviant causation and beginner’s luck cases are standard examples of this phenomenon.

2. Set-Up

One of the central questions in action theory concerns the nature of intentional action: what is the difference between acting intentionally and acting unintentionally? According to what we will call the *Knowledge Hypothesis*, propositional knowledge about how to achieve one's desired ends partly explains intentional action:

Knowledge Hypothesis: Intentional action is (in part) explained by knowledge about how to perform it.

By “knowledge about how to φ ,” we mean propositional knowledge concerning the means for φ -ing—a sort of knowledge that can be ascribed more fully by knowledge-that ascriptions of the form: *S* knows that a certain means is a means to φ .

Cases where agents do not act intentionally precisely because they lack relevant knowledge motivate the Knowledge Hypothesis (Gibbons 2001; Pavese 2020; 2021a; 2021b; 2021c; Pavese & Beddor 2022). For example, consider a case from Hawley (2003): Susie believes that she can annoy Joe by smoking, but in fact Joe is annoyed by Susie's tapping on her cigarette box, which she does whenever she smokes. Intuitively, while Susie succeeds in annoying Joe, she does not annoy him intentionally. A plausible explanation for this intuition is that she does not annoy him intentionally because she does not know the means of annoying him—she does not know that she can annoy him by tapping on her cigarette box.³

Despite its explanatory appeal, some critics have claimed that the Knowledge Hypothesis is too strong (Setiya 2012; Cath 2015; Moss 2020) on the basis of the intuition that one might act intentionally on a Gettiered belief—that is, in absence of knowledge. On the basis of these intuitions, it is generally thought that the Knowledge Hypothesis ought to be dismissed because it is counterintuitive. However, such criticisms are often based on controversial intuitions about a small sample of cases. In response, it is appropriate to look at a wider set of cases where, as we argue, the judgements are rather reversed—that is, epistemic luck does seem to undermine intentional action. Indeed, if epistemic luck undermines intentional action in a variety of cases, as our studies suggest, then there is reason to reassess the Knowledge Hypothesis.

3. It is an interesting question whether the Knowledge Hypothesis is plausible for basic actions as well—such as lifting one's finger—which plausibly require no means for performing them. We will only consider complex actions in this work, so we can bracket this question here—though see Pavese and Beddor (2022) for an extensive argument that the Knowledge Hypothesis extends to basic actions.

In the light of this disagreement, the main question we investigate here is:

Intentional Action Question: Can epistemic luck undermine intentional action?

A positive answer to this question would suggest that the Knowledge Hypothesis cannot be dismissed as counterintuitive and that intuitions are compatible with agentive luck being sometimes (at least in part) epistemic.⁴

In conjunction with the **Intentional Action Question**, we will investigate the following question:

Knowledge-How Question: Can epistemic luck undermine knowledge-how?

The main motivation for studying these questions together is that acting intentionally seems to require knowing how to perform the action in question (Ryle 1949; Stanley & Williamson 2001; Pavese 2021b; Cath 2015).⁵ This claim motivates our second hypothesis:

Knowledge-How Hypothesis: Intentional action is (in part) explained by knowledge-how.

By ‘knowledge-how’ we mean the sort of expertise that guides us in our skilled performance and that is standardly ascribed in English by ascriptions of the form “S knows how to φ .”

The Knowledge-How Hypothesis traces back to Ryle (1949), who motivates it by contrasting the clown with the klutz. Both of them tumble but only the clown does so intentionally. Part of the explanation for this difference seems to

4. Recent work on the relation between knowledge and action has focused on the relationship between intentional action and practical knowledge (e.g., Beddor & Pavese 2022; Pavese 2022b). Practical knowledge (in Anscombe’s 1958 sense) is the sort of knowledge of what one is doing which agents supposedly have when acting intentionally. In some recent work, Vekony, Mele, and Rose (2020) found that participants tend to ascribe intentionality even in cases where agents lack practical knowledge. For a detailed discussion of these results, see Pavese (in press). Here we just note that these results are fully consistent with our (weaker) Knowledge-that Hypothesis, according to which intentional action requires knowledge about how to achieve one’s goal. For example, according to the Knowledge-that Hypothesis, it is *necessary* for one to swim intentionally that one possess knowledge about how to successfully swim—what movements one must make, how to enter the water, how to coordinate one’s arms and one’s legs, and so on. It is compatible with the Knowledge-that Hypothesis, however, that one intentionally performs an action without necessarily knowing *that they are performing it* when they do.

5. Bengson, Moffett, and Wright (2009) report experimental results that seem to corroborate this claim.

be that the clown, but not the klutz, knows how to tumble. For further motivation, recall Hawley's (2003) case of Susie (discussed above). Susie has the ability to annoy Joe but does not have the ability to do it intentionally, plausibly because she does not know how to annoy him. The Knowledge-How Hypothesis also follows from the idea that intentional action requires skill (Heider 1958; Mele & Moser 1994; Malle & Knobe 1997; Setiya 2012) together with the plausible assumption that, if one is skilled at a task, one knows how to perform that task.⁶

Now, suppose the Knowledge-How Hypothesis is true. Then, if knowledge-how were sensitive to epistemic luck, we would expect intentional action to be sensitive to epistemic luck too. Thus, the Intentional Action Question is deeply intertwined with the Knowledge-How Question.

Before turning to our argument and results, we have two preliminary remarks about our methodology. The first concerns the metaphilosophical picture underlying this project. In this article, we will sometimes move from claims about judgments about intentional action and knowledge-how to claims about intentional action and knowledge-how themselves. Some readers might question whether these transitions from folk psychology to the phenomena themselves are justified. The relevance of folk psychology to philosophical questions is of course a thorny issue (e.g., Cappelen 2012; Williamson 2016) that we cannot hope to settle here. But we hope to give the reader a sketch of our approach.

We would like to note, first, that one main goal of this article is to present new cases that we think support the Knowledge Hypothesis. Each study involves a novel thought experiment where our own judgments align with the Knowledge Hypothesis. In this respect, our methodology is the same as that standardly used in traditional epistemology and action theory—where judgments about thought experiments are used as evidence in support or against philosophical views about the relevant phenomena. The empirical investigation we report is meant as corroborative evidence that our judgments in these thought experiments are widely shared. Indeed, we think that the current status of the debate on the Knowledge Hypothesis encourages a similar empirical investigation of our judgments. As noted, one of the main objections leveled at the Knowledge Hypothesis is that this view is at odds with judgments about cases (e.g., Setiya 2012; Cath 2015; Moss 2020). By relying on individual philosophers' judgments

6. Some may object that the Knowledge-How Hypothesis might deliver the wrong results in Davidson's (1971) carbon copier case, in which someone writes heavily on a page, intending to produce 10 legible carbon copies. According to Davidson, if they succeed in making 10 copies, they did so intentionally, even though they did not know how to make carbon copies. Taking a cue from Davidson's own remarks, however, one might respond that although the agent does not know how to make carbon copies, they know how to perform some actions that are a possible means for carbon copying, such as pressing hard enough, moving the pen thus and so, etc. Similar responses can be given to other putative counterexamples (see Shepherd & Carter 2021), where one performs an action intentionally but where luck plays a big role in performing the action.

about cases, however, it is possible that the relevant judgments are themselves dogmatic or biased on one's particular philosophical views. By undertaking an empirical investigation of non-philosophers' judgments, we explore judgments that are not (on the whole) affected by the philosophical theories that they are attempting to challenge or support. In this way, we simultaneously offer motivation for questioning this common line of objection against the Knowledge Hypothesis and offer evidence that our own judgments in support of Knowledge Hypothesis are not theoretically biased.

Our second response to this concern is that, even for those who think that ordinary judgments have little or no bearing on whether the Knowledge Hypothesis holds, it is an interesting question whether the ordinary concepts of intentional action and knowledge-how require the ordinary concept of propositional knowledge. Indeed, there is an ongoing debate on whether the ordinary concept of knowledge-how is sensitive to epistemic luck and to Gettierization (e.g., Poston 2009; Carter, Pritchard, & Shepherd 2019; Gonnerman, Mortensen, & Robbins 2018; 2021). There has also been some preliminary discussion on whether the ordinary concept of intentional action requires the ordinary concept of knowledge (Gibbons 2001; Vekony, Mele, & Rose 2020).⁷ These questions fall under the more general question as to whether mindreading and the folk psychology of action require knowledge representation (Nagel 2013; Phillips et al. 2021). Our findings are directly relevant to the mindreading question, whatever stance one takes on the more general metaphilosophical question.

Our second methodological remark is about our approach to the overarching explanation. We will present three different studies which exhibit a *general pattern*: cases involving epistemic luck decrease people's tendency to ascribe intentional action. Plausibly, the best explanation for these findings should cast light on the general pattern. So, a unified explanation of the pattern ought to be preferred to *piecemeal explanations*—that is, to a set of different explanations that explain the results in each study locally. So, while the studies may reveal vari-

7. Recent work on the relation between knowledge and action has focused on the relationship between intentional action and practical knowledge. Practical knowledge (in Anscombe's 1958 sense) is the sort of knowledge of what one is doing which agents supposedly have when acting intentionally. In some recent work, Vekony, Mele, and Rose (2020) found that participants tend to ascribe intentionality even in cases where agents lack practical knowledge. These results are fully consistent with our (weaker) Knowledge-that Hypothesis, according to which intentional action requires knowledge about how to achieve one's goal. For example, according to the Knowledge-that Hypothesis, it is *necessary* for one to swim intentionally that one possess knowledge about how to successfully swim—what movements one must make, how to enter the water, how to coordinate one's arms and one's legs, and so on. It is compatible with the Knowledge-that Hypothesis, however, that one intentionally performs an action without necessarily knowing *that they are performing it* when they do.

ous *other* patterns that differ by case—and that may be interesting and critical for future work—our focus will be on this general pattern.

One straightforward unified explanation of these findings comes from the Knowledge Hypothesis. According to it, epistemic luck undermines judgments of intentional action since judgments of intentional actions are rightly sensitive to whether knowledge is present. However, this is not the only possible unified explanation. A competing hypothesis is:

Alternative Hypothesis: Agentive luck and epistemic luck can co-occur in some cases, but they are nonetheless independent of each other.

According to the Alternative Hypothesis, knowledge-that and intentional action can be both sensitive to luck—the former to epistemic luck, the latter to agentive luck—while being explanatorily independent of each other. Thus, the Knowledge Hypothesis and Alternative Hypothesis differ in their predictions of explanatory dependency. It is crucial for our overall argument that we offer evidence against the Alternative Hypothesis. We will gradually build our case against the Alternative Hypothesis in the discussion of our studies in §§3–5, and we will pull the strings of our overall argument in §6.

3. Study 1

In this study, our goal was to test whether the sort of epistemic luck we find in classical Gettier (1963) cases undermines people’s tendency to ascribe intentional action and knowledge-how. In all of the vignettes we used, someone acts on the basis of a belief that they have about how to perform an action. In the lucky condition, the agent’s belief is Gettiered. In the control condition, a subject’s belief is not Gettiered. In each condition, we asked participants the extent to which they agreed that the agent knows-that, knows-how, and acted intentionally.

3.1. Participants

A pilot study revealed a large effect of luck on judgments of knowledge-that. As such, power analysis for a large effect revealed that we required a total of 474 participants. Expecting an exclusion rate of 1.5%, we recruited 481 participants. All participants were United States nationals, were born in and resided in the United States, spoke English as their first language, had not participated in our

pilot study, and had a 99% approval rating on *Prolific*. 3 participants reported not paying attention, so they were excluded. We analyzed data from the remaining 478 participants ($M_{age} = 34$, $SD = 12.3$, $Range_{age} = [18-79]$, 48% female).

3.2. Materials and Procedure

After consenting to participate, participants were randomly assigned to 1 of 6 conditions in a 2 (lucky, control) \times 3 (Vignette: email, address, phone) between-participants design. Each participant read one vignette (example in Table 1; see Supplementary Materials for all vignettes and dependent measures). On the same page, they were asked for their level of agreement with a statement about the agent's knowledge-that, knowledge-how, and intentionality all on a -50-50 scale [-50 = strongly disagree, 0 = neutral, 50 = strongly agree] in this fixed order. Participants were then asked for basic demographic information and to respond to one explicit attention check (See Supplementary Materials⁸).

<u>Lucky:</u>	<u>Control:</u>
<p>Alvin works at a barn supply company. Ed and Susan, Alvin's co-workers, are up for a promotion to supervisor. Kate, a typically trustworthy friend, tells Alvin that Ed got the promotion. Alvin also seems to remember that Ed has a personal email address: barns678@yahoo.com.</p>	<p>Alvin works at a barn supply company. Ed and Susan, Alvin's co-workers, are up for a promotion to supervisor. Kate, a typically trustworthy friend, tells Alvin that Susan got the promotion. Alvin also seems to remember that Susan has a personal email address: barns678@yahoo.com.</p>
<p>Alvin wants to congratulate the person who got the promotion. He believes that Ed got the promotion. So, he also comes to believe that the person who got the promotion has the email address: barns678@yahoo.com. So, Alvin writes a concise congratulatory email, addresses it to "Our New Supervisor," and then sends the email to barns678@yahoo.com, thinking that he is congratulating Ed.</p>	<p>Alvin wants to congratulate the person who got the promotion. He believes that Susan got the promotion. So, he also comes to believe that the person who got the promotion has the email address: barns678@yahoo.com. So, Alvin writes a concise congratulatory email, addresses it to "Our New Supervisor," and then sends the email to barns678@yahoo.com, thinking that he is congratulating Susan.</p>
<p>It turns out that Alvin was wrong twice over. Susan had actually got the promotion. And Susan, not Ed, had the email address barns678@yahoo.com. As a result, Susan received Alvin's email.</p>	<p>It turns out that Alvin was right twice over. Susan had actually got the promotion. And Susan had the email address barns678@yahoo.com. As a result, Susan received Alvin's email.</p>

(Contd.)

8. All supplementary materials, data, and analysis code are available at https://osf.io/s4hyd/?view_only=672950d58a7941c5bdb39684c3fcb785.

Dependent Measures:

To what extent do you agree with the following statements about the passage you just read?

Alvin knows that the person who got the promotion had the email address barns678@yahoo.com.⁹

Alvin knows how to email the person who got the promotion.

Alvin intentionally emailed the person who got the promotion.¹⁰

Table 1. Email vignette used in Experiment 1 and the dependent variables.

3.3. Results

Participants who read about the agent in the control condition agreed to a greater extent that the agent knows-that ($M = 33.41$, $SD = 22.35$, $n = 237$) than participants in the lucky condition ($M = -13.64$, $SD = 35.39$, $n = 241$) ($b = 46.97$, $SE = 2.70$, $t = 17.38$, $p < .001$, $CI_b [41.69, 52.30]$, $d = 1.59$) (Figure 1A).

Participants who read about the agent in the control condition agreed to a greater extent that the agent knows-how ($M = 39.96$, $SD = 15.75$, $n = 237$) than participants in the lucky condition ($M = 3.78$, $SD = 35.36$, $n = 241$) ($b = 36.14$, $SE = 2.47$, $t = 14.58$, $p < .001$, $CI_b [31.28, 41.01]$, $d = 1.33$) (Figure 1B).

Participants who read about the agent in the control condition agreed to a greater extent that the agent intentionally performed the action ($M = 43.06$,

9. We have one important note about our dependent measure for knowledge-that. In the representative vignette above, the statement about knowledge of the means would be: "Alvin knows that sending an email to barns678@yahoo.com is a means to email the person who got the promotion." This knowledge ascription is, however, long, clumsy and unnatural. (Indeed, the authors believe that the most natural way to ascribe knowledge of the means in English is through knowledge-how ascriptions). In order to avoid undesirable pragmatic effects due to the clumsiness of the ascription, we planned to use a more natural and simpler knowledge-that ascription, one that, in the specific case considered, would have to hold if that knowledge of the means ascription were true. As such, we asked participants the extent to which they agreed with the following statement: "Alvin knows that the person who got the promotion had the email address barns678@yahoo.com." For it to be true that Alvin knows that sending the email to barns678@yahoo.com is a means to email the person who got the promotion, it would have to be true, *given the context of the vignette*, that Alvin knows that the person who got the promotion had the email address barns678@yahoo.com. Moreover, if the latter ascription is Gettiered, so would the former.

10. In the dependent measures, the knowledge-that and knowledge-how questions are both phrased in the present tense, where the intentionality question is in the past tense. This contrast is due to the fact that, whereas the intentionality question is most naturally phrased in the past tense, the past tense of knowledge-how questions gives rise to a version of what linguists call the 'actuality entailment' (Bhatt 1999). In general, if S succeeds at φ -ing, people tend to judge that S was able to φ (in the past), even though S got lucky. Still, people tend not to agree with the claim that S has the ability to φ or that S is able to φ (present tense). Arguably, this actuality entailment extends also to past tense knowledge-how judgments. The explanation of the actuality entailment is a matter of controversy. In this paper, we bypass these issues by focusing on judgments about knowledge-how in the present tense.

$SD = 15.14$, $n = 237$) than participants in the lucky condition ($M = -9.82$, $SD = 37.72$, $n = 241$) ($b = 52.77$, $SE = 2.60$, $t = 20.24$, $p < .001$, $CI_b [47.66, 57.89]$, $d = 1.85$) (Figure 1C).¹¹

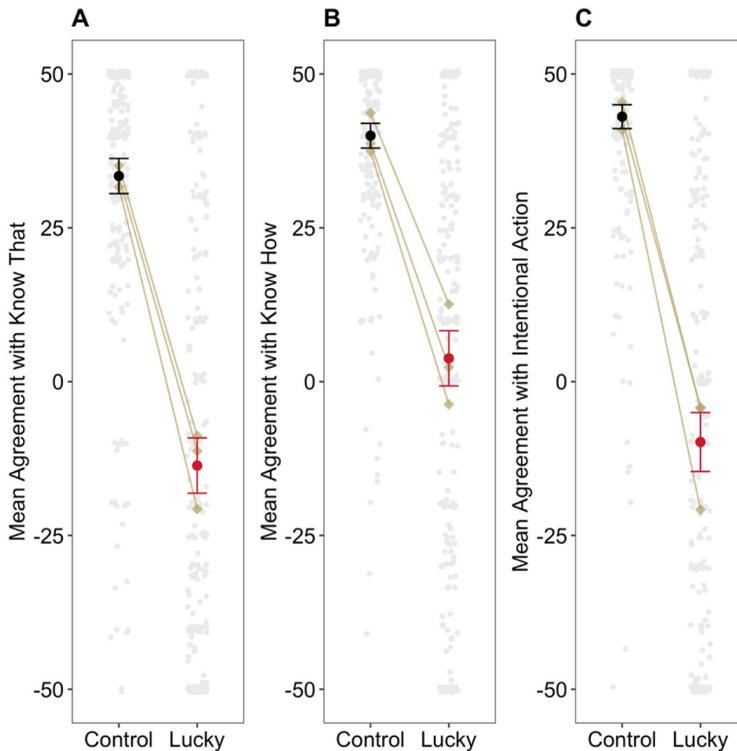


Figure 1. Mean agreement with the knowledge-that (A), knowledge-how (B), and intentional action (C) statements in Study 1. Error bars indicate 95% confidence intervals. Light gray points represent individual participant responses evenly jittered. Gold diamonds and lines represent means for individual vignettes.

We then ran a serial mediation analysis across the three vignettes with condition modeled as the independent variable, knowledge-that and then knowledge-how as serial mediators, and intentional action as the dependent variable. As theorized, the effect of condition on intentional action was serially mediated by knowledge-that and then knowledge-how (Figure 2). The indirect pathway of

¹¹. To test for ordering effects of the dependent measures, we ran an additional experiment. In Supplementary Experiment 1, we gave participants the email vignette, and we randomized the order of presentation of the dependent measures. We found the same pattern of results in the supplementary experiment, suggesting that the ordering of the dependent measures does not affect our findings. See Supplementary Materials for details: https://osf.io/s4hyd/?view_only=672950d58a7941c5bdb39684c3fcb785.

the effect of condition on intentional action *via* knowledge-that and then knowledge-how was significant ($b = -10.74$, $z = -7.37$, $p < .001$). This pathway partially accounted for the overall impact of condition on intentional action with the direct effect remaining significant ($b = -39.62$, $z = -14.37$, $p < .001$).

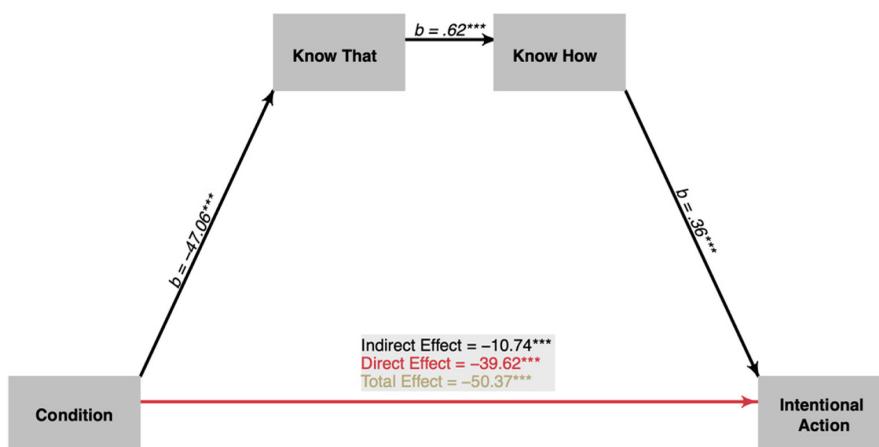


Figure 2. Mediation analysis results. *** indicates $p < .001$. Decimals rounded and then truncated to two places.

3.4. Discussion

In accordance with previous work on the effects of Gettier luck on knowledge-that judgments (Nagel, San Juan, & Mar 2013; Turri 2013; Machery et al. 2017; Turri, Buckwalter, & Blouw 2015), we found that judgments about knowledge-that are sensitive to whether the relevant belief is Gettiered: in the lucky condition relative to the control, people are less inclined to ascribe knowledge-that. In line with our hypothesis, we have also found that Gettierization has a large impact both on people's inclination to ascribe knowledge-how and on people's inclination to ascribe intentional action.¹²

12. Some readers may be concerned that participants did not completely disagree with the knowledge-how statement as the mean and the confidence intervals are right around the midpoint. While we did not plan to test this feature of our data, we have reason to think that this feature of the data does not replicate. In a supplementary experiment (see fn. 11 and Supplementary Materials), where the order of the questions was randomized, we found very low agreement with the knowledge-how statement in the lucky conditions. So, we have no evidence that this agreement around the midpoint is a robust finding. Future work might investigate this further.

One might object that the vignettes in Study 1 all involve *deviant causation*. If the vignettes in Study 1 involve some sort of deviance, one might argue that the Alternative Hypothesis is well positioned to explain these findings: since these vignettes involve *both* deviance *and* epistemic luck, the effect on judgments about intentional action might be explained not in terms of the lack of knowledge but in terms of the deviance of the performances.

In action theory, an action is deviantly caused if the causal chain that starts when the agent initiates the action and that ends with the intended outcome does not unfold through means that the agent intended because of an intervening disturbance. For instance, consider a classic example of a deviant causal chain (Adams & Mele 1989: 519). A philosopher intends to knock over his glass to distract his commentator. However, his intention upsets him so much that his hand shakes uncontrollably, striking the glass and knocking it to the floor. In this scenario, the shaking disrupts the causal chain leading to the knocking of the glass and seems to preclude the philosopher's action from being intentional. Indeed, it is widely thought that, when some outcome is deviantly caused, the agent did not intentionally bring about that outcome (e.g., Chisholm 1966a; Davidson 1973).

Study 1 does not involve an intervening disturbance, so it does not fit this common characterization of deviance. Moreover, there are several other problems with this putative alternative explanation of our findings. For one thing, it would not explain Study 1's findings involving *knowledge-how*. Even though deviance may preclude intentional action, it does not usually preclude knowledge-how. Consider again Adams and Mele (1989)'s example of the shaking philosopher above. In this scenario, the causal deviance may preclude the philosopher's action from being intentional, but the philosopher still knows how to knock over his glass. So, it is doubtful that the alleged deviance of the performance can explain the effect on the knowledge-how judgments.

Second, suppose we accept the Knowledge-How Hypothesis. This predicts that if knowledge-how is sensitive to epistemic luck, intentional action will also be sensitive to epistemic luck. So, the fact that knowledge-how judgments are sensitive to epistemic luck suggests that intentional action judgments are sensitive to epistemic luck in particular—not just any old type of luck.

Finally, the results of our serial mediation reveal that the judgments of intentional action are partially mediated by the knowledge-how and knowledge-that judgments. This result is predicted by the Knowledge Hypothesis but not predicted by the Alternative Hypothesis. Thus, when it comes to fully accounting for our findings, the Alternative Hypothesis is already at an explanatory disadvantage.

These points notwithstanding, it is worth directly testing the deviance explanation. We undertake this task in our second study.

4. Study 2

In action theory, so-called ‘lucky successes’ are not usually thought to involve deviant causation. In lottery cases, for example, the causal chain that leads to one’s winning or losing a lottery does not need to be disrupted and might happen exactly as the agent wanted it to happen (Mele & Moser 1994; Setiya 2012). Epistemologists believe that lottery cases involve a kind of epistemic luck (Williamson 2000; Hawthorne 2004), and this intuition has been confirmed experimentally (e.g., Friedman & Turri 2015). So, in Study 2 we used a lottery vignette: the agent is justified in believing on statistical grounds that they would lose, the outcome of the lottery is so undesirable that the agent wants to lose, and they lose.

4.1. Participants

We kept the sample size and qualifications the same. 4 participants reported not paying attention, so they were excluded. We analyzed data from the remaining 157 participants ($M_{age} = 35$, $SD = 13$, $Range_{age} = [18-75]$, 45% female).

4.2. Materials and Procedure

After consenting to participate, participants were randomly assigned to 1 of 2 conditions (lucky, control) in a between-participants design. Each participant read one vignette (Table 2). On the same page, they were asked for their level of agreement with a statement about the agent’s knowledge-that, knowledge-how, and intentionality on the same scale as Study 1 in fixed order. Participants were then asked for basic demographic information and to respond to the same explicit attention check as in Study 1.

4.3. Results

Participants who read about the agent in the control condition agreed to a greater extent that the agent knows that ($M = 42.12$, $SD = 11.55$, $n = 79$) than participants in the lucky condition ($M = 15.38$, $SD = 30.51$, $n = 78$) ($t(98.38) = 7.24$, $p < .001$, $d = 1.46$, $CI [1.01, 1.90]$) (Figure 3A).

Participants who read about the agent in the control condition agreed to a greater extent that the agent knows how ($M = 38.97, SD = 19.24, n = 79$) than participants in the lucky condition ($M = -6.70, SD = 29.69, n = 78$) ($t(131.73) = 11.42, p < .001, d = 1.99, CI [1.57, 2.41]$) (Figure 3B).

<p><u>Lucky:</u></p> <p>Alvin works for a large tech company. Every Christmas, the company organizes a lottery for employees, who are all required to participate. The lottery works as follows: every employee has to pick a number between 1 and 1,000,000. Just before Christmas, the winning number is randomly drawn. Any employee who picked the winning number wins an all-inclusive cruise.</p> <p>Since he hates cruises, Alvin wants to lose the lottery. Given the odds, Alvin believes that if he just picks a random number between 1 and 1,000,000, he will lose. So, Alvin picks 98,021, hoping that his number is not selected.</p> <p>Just before Christmas, the winning number is drawn: 95,500. So, by pure luck, Alvin loses the lottery and avoids the cruise.</p>	<p><u>Control:</u></p> <p>Alvin works for a large tech company. Every Christmas, the company organizes a lottery for employees, who are all required to participate. The lottery works as follows: every employee has to pick a number between 1 and 1,000,000. Just before Christmas, the winning number is randomly drawn. Any employee who picked the winning number wins an all-inclusive cruise.</p> <p>Since he hates cruises, Alvin wants to lose the lottery. Much to his relief, Alvin knows from a reliable source that the CEO of the company has rigged the lottery so that the winning number is 95,500. So, Alvin picks 98,021, confident that his number will not be selected.</p> <p>Just before Christmas, the winning number is drawn: 95,500. So, as expected, Alvin loses the lottery and avoids the cruise.</p>
<p><u>Dependent Measures:</u></p> <p>To what extent do you agree with the following statements about the passage you just read? Alvin knows that he will lose the lottery.¹³ Alvin knows how to lose the lottery. Alvin intentionally lost the lottery.</p>	

Table 2. Vignette used in Study 2 and the dependent variables.

13. As in Study 1, we planned to use the simplest and most natural relevant knowledge-that ascription. In this vignette, the action is *losing the lottery*, and the corresponding knowledge-how ascription is “Alvin knows how to lose the lottery.” Given the context of this vignette, the relevant ascription of knowledge of the means would be “Alvin knows that picking 98,021 is a means to lose the lottery.” In the context of this vignette in which Alvin has picked the losing number 98,021, it is true that Alvin knows that picking 98,021 is a means to lose the lottery only if by picking that number, Alvin knows that he will lose the lottery. Hence, for simplicity, we asked whether in the scenario in which Alvin has picked 98,021, it is true that Alvin knows that he will lose the lottery. We aim for this simplicity in all vignettes that follow.

Participants who read about the agent in the control condition agreed to a greater extent that the agent intentionally performed the action ($M = 43.65$, $SD = 18.69$, $n = 79$) relative to participants in the lucky condition ($M = 2.14$, $SD = 34.33$, $n = 78$) ($t(118.62) = 9.39$, $p < .001$, $d = 1.72$, $CI [1.30, 2.14]$) (Figure 3C).

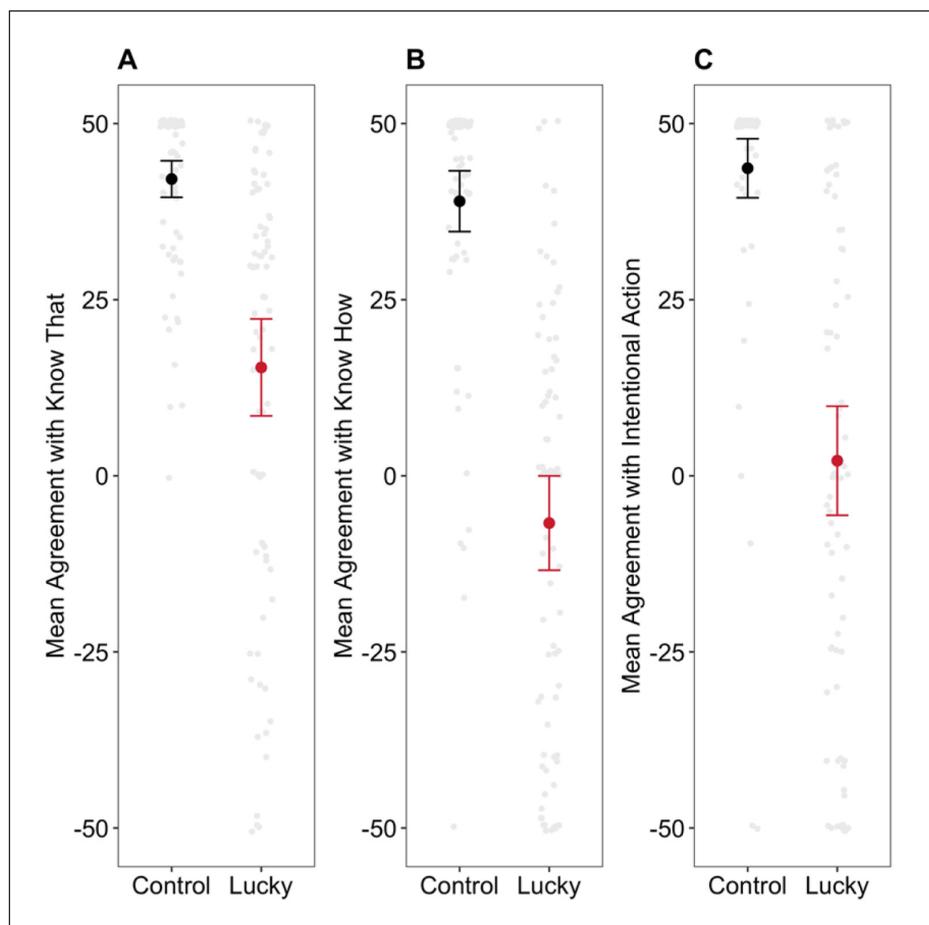


Figure 3. Mean agreement with the knowledge-that (A), knowledge-how (B), and intentional action statements in Study 2. Error bars indicate 95% confidence intervals. Light gray points represent individual participant responses evenly jittered.

We then ran a serial mediation analysis modeled in the same way as Study 1. As theorized, the effect of condition on intentional action was serially mediated by knowledge-that and then knowledge-how (Figure 4). The indirect pathway of the effect of condition on intentional action via knowledge-that and then knowledge-how was significant ($b = -15.08$, $z = -5.46$, $p < .001$). This

pathway partially accounted for the overall impact of condition on intentional action with the direct effect remaining significant ($b = -10.91$, $z = -2.95$, $p = .003$). (Figure 4).

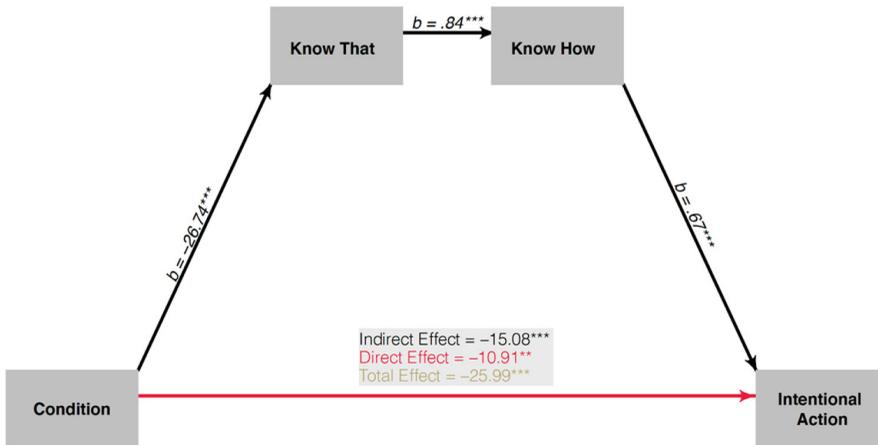


Figure 4. Mediation analysis results. *** indicates $p < .001$. ** indicates $p < .01$. Decimals rounded and then truncated to two places.

4.4. Discussion

These results replicate previous studies on lotteries (Friedman & Turri 2015). Also, in the lucky condition, this vignette does not involve deviant causation, as no intervening disturbance affects the outcome. Even though the new vignette does not involve deviant causation, we find large effects of epistemic luck on knowledge-how judgments ($d = 1.99$) and on intentional-action judgments ($d = 1.42$). As predicted, epistemic luck undermines people's tendency to ascribe intentional action even in absence of deviant causation. Moreover, the mediation effect of condition on intentional action via knowledge-that and then knowledge-how corroborates our predictions that judgments about intentional action are partially explained by judgments of knowledge-that.

5. Study 3

One might argue for a yet different explanation of the effect observed in Study 2. One might contend that losing a fair lottery is not something that an agent can

even plan—it is a *happening*, rather than a planned action—so it is not the sort of thing that can be intentional because it is not the sort of thing that can be under an agent’s control. According to one prominent view of intentional action—the plan view—an action is intentional (and controlled) only if it happens in accordance with the agent’s plans (Brand 1986). If losing a lottery is not something that the agent can plan, then the plan view might be able to explain these findings without having to invoke knowledge.

There is a straightforward argument against this explanation. The plan view of intentional action only delivers an *alternative* explanation of our findings if plans do not themselves involve knowledge; if plans are themselves epistemically laden, then this would not be a rival to the Knowledge Hypothesis. For this reason, we will focus on a non-epistemic plan view of intentional action—that is, a view on which plans do not require knowledge about how to act. One initial reason to be dissatisfied with this view is that it does not predict a mediation effect of knowledge-that and knowledge-how judgments on intentional action judgments.

In order to further test the Knowledge Hypothesis against the non-epistemic plan view, we planned a third study investigating two scenarios where the performance is clearly a planned action.

5.1. Study 3a

The first case is structurally analogous to Chisholm’s (1966b) sheep in the meadow case, itself anticipated by Dharmottara’s (770: D:4429, 9a2–3) desert mirage case. In Chisholm’s (1966b) case, an agent sees a sheep-shaped rock in the meadow and comes to believe justifiably that there is a sheep in the meadow. As it turns out, there is a sheep in the meadow, hidden behind the rock. In Dharmottara’s desert mirage case, a desert traveler thinks that there is water upon seeing a mirage from far away and when they reach the spot there is actually water under a rock. We investigated a modified version of Dharmottara’s case, where the traveler performs an action (leading a group to a location) based on their Gettiered belief.

5.1.1. Participants

We kept the sample size and qualifications the same. 1 participant reported not paying attention, so they were excluded. We analyzed data from the remaining 159 participants ($M_{age} = 32$, $SD = 11.5$, $Range_{age} = [18–67]$, 47% female).

5.1.2. *Materials and Procedure*

After consenting to participate, participants were randomly assigned to 1 of 2 conditions (lucky, control) in a between-participants design. Each participant read one vignette (Table 3). On the same page, they were asked for their level of agreement with a statement about the agent’s knowledge-that, knowledge-how, and intentionality on the same scale as Study 1 in fixed order. Participants were then asked for basic demographic information and to respond to the same explicit attention check as in Study 2.

<p><u>Lucky:</u></p> <p>Mary is a guide leading a group of travelers through the desert. After hours of travel, the group has become hopelessly lost and is in danger of running out of water.</p> <p>On the North horizon, Mary spots what appears to be water. On this basis, Mary comes to believe that there is water to the North. So, she leads the group North, even though it is nearly a mile away from the likeliest water source.</p> <p>It turns out that Mary was wrong about what she saw. As a matter of fact, Mary had spotted a mirage. But, by a sheer coincidence, there was indeed a pool of water to the North—though it could not have been seen from far away. So—by incredible luck—in leading the group North, Mary succeeds at bringing the group to water.</p>	<p><u>Control:</u></p> <p>Mary is a guide leading a group of travelers through the desert. After hours of travel, the group has become hopelessly lost and is in danger of running out of water.</p> <p>On the North horizon, Mary spots what appears to be water. On this basis, Mary comes to believe that there is water to the North. So, she leads the group North, even though it is nearly a mile away from the likeliest water source.</p> <p>It turns out that Mary was exactly right about what she saw. There was indeed a large pool of water to the North at the location where Mary had spotted it. So, in leading the group North, Mary succeeds at bringing the group to water.</p>
<p><u>Dependent Measures:</u></p> <p>To what extent do you agree with the following statements about the passage you just read?</p> <p>Mary knows that there is water to the North. Mary knows how to bring the group to water. Mary intentionally brought the group to water.</p>	

Table 3. Vignette used in Study 3a and the dependent variables.

5.1.3. *Results*

Participants who read about the agent in the control condition agreed to a greater extent that the agent knows that ($M = 9.27, SD = 27.63, n = 76$) than participants

in the lucky condition ($M = -26.32$, $SD = 25.07$, $n = 83$) ($t(151.79) = 8.48$, $p < .001$, $d = 1.38$, $CI [1.02, 1.73]$).

Participants who read about the agent in the control condition agreed to a greater extent that the agent knows how ($M = 18.11$, $SD = 24.85$, $n = 76$) than participants in the lucky condition ($M = -19.43$, $SD = 25.79$, $n = 83$) ($t(156.58) = 9.34$, $p < .001$, $d = 1.49$, $CI [1.14, 1.85]$).

Participants who read about the agent in the control condition agreed to a greater extent that the agent intentionally performed the action ($M = 28.30$, $SD = 25.54$, $n = 76$) than participants in the lucky condition ($M = -10.78$, $SD = 35.46$, $n = 83$) ($t(148.96) = 8.02$, $p < .001$, $d = 1.31$, $CI [0.96, 1.67]$).

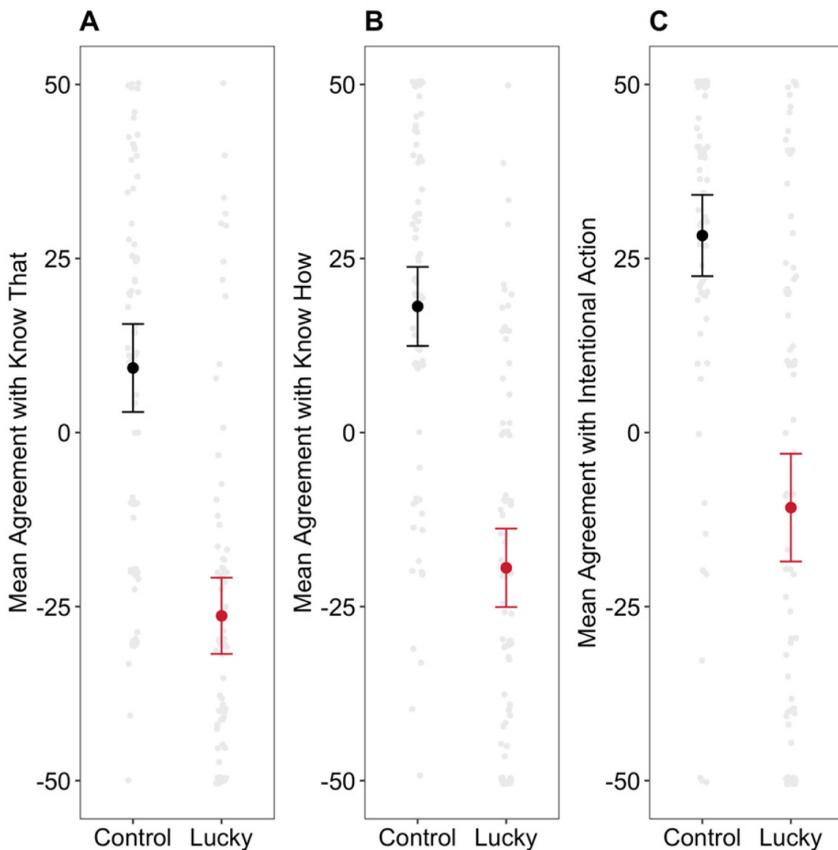


Figure 5. Mean agreement with the knowledge-that (A), knowledge-how (B), and intentional action statements in Study 3a. Error bars indicate 95% confidence intervals. Light gray points represent individual participant responses evenly jittered.

We then ran a serial mediation analysis modeled in the same way as Study 1. As theorized, the effect of condition on intentional action was serially mediated by knowledge-that and then knowledge-how (Figure 6). The indirect pathway of the effect of condition on intentional action via knowledge-that and then knowledge-how was significant ($b = -17.05$, $z = -5.57$, $p < .001$). This pathway partially accounted for the overall impact of condition on intentional action with the direct effect remaining significant ($b = -15.69$, $z = -3.33$, $p = .001$).

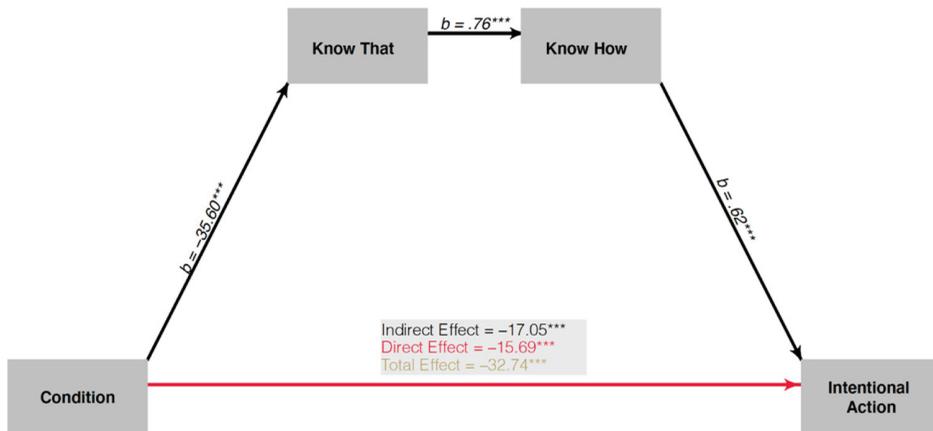


Figure 6. Mediation analysis results. *** indicates $p < .001$. Decimals rounded and then truncated to two places.

5.1.4. Discussion

With this study, we have replicated the results of Study 1 and Study 2. These results and the mediation analysis lend further support to our Knowledge Hypothesis.¹⁴ These findings are not easy to explain on the non-epistemic plan view. After all, the guide’s bringing the group to water is definitely an

14. While it is unrelated to our primary hypothesis, some readers might wonder why participants’ agreement with the knowledge-that statement is rather low in the control condition in this experiment. One explanation might be that the wording of the second paragraph in the vignette—where the water is described as “merely appearing” to be there and where its presence is implied to be unlikely since the location is described as “nearly a mile away from the likeliest water source”—might suggest that the agent is not justified in their belief because of the lower probability of the presence of water. This wording might explain the lower overall agreement. Future work might explore this further.

action performed in accordance with their plan—namely, a plan to bring the group to water by going North, in the direction where water appears to be.¹⁵

5.2. Study 3b

In this study, we used a case structurally analogous to Russell's (1948) stopped clock case. In our version of the case, a spy needs to send an important signal at 4pm, and they believe that they are sending the signal at the right time. But their belief is based on the fact that the stopped clock is usually reliable.

5.2.1. Participants

We kept the sample size and qualifications the same. 3 participants reported not paying attention, so they were excluded. We analyzed data from the remaining 157 participants ($M_{age} = 33$, $SD = 11.5$, $Range_{age} = [18-70]$, 57% female).

5.2.2. Materials and Procedure

The procedure and design for this study was the same as in the previous studies. The only difference was that each participant read a new vignette (Table 4). Participants were then asked for basic demographic information and to respond to the same explicit attention check as in Study 2.

15. Proponents of the non-epistemic plan view might protest that this is the wrong way of describing Mary's plan. Perhaps the content of her plan is something like this: *to bring the group to the location where I discovered the water to be*. And this plan arguably cannot be satisfied, given the assumption that *discovering* entails *knowing*; after all, the plan falsely presupposes that she knows water is located in the direction she is looking. However, there are two problems with this suggestion. First, it is not clear from the description of the case that she does take herself to *know* that water is located in the direction she is looking (she might instead think it is a highly plausible conjecture). Second, in order to avoid the charge of *ad hocery*, proponents of the non-epistemic plan view would need to give some reason for thinking that this is the right way of describing the content of her plan. The most obvious way to do so would be to insist plans need to be structured by one's knowledge about how to perform the action. But this would be to effectively adopt an epistemic view of plans, according to which plans require knowledge after all.

<p><u>Lucky:</u></p> <p>During World War II, Silvia, a British spy, needs to send a radio signal to her undercover contact in France. Because of how the operation is designed, it is critically important that she sends the signal at exactly 4pm on Tuesday, not a minute earlier or a minute later. For this reason, the radio transceiver comes with a clock that is usually highly reliable.</p> <p>When Silvia looks at the clock on Tuesday afternoon, it reads 4pm. Consequently, she comes to believe that it is now time to send the signal. So, at that very moment, she sends the radio signal to her undercover contact in France.</p> <p>It turns out Sylvia was wrong in relying on that clock. Unbeknownst to Sylvia, the clock's mechanism had stopped working at exactly 4pm Monday afternoon (and Sylvia has had no opportunity to inspect the clock in the meantime). However, by an incredible stroke of luck, Sylvia happened to look at the clock at 4pm, exactly 24 hours after the clock stopped working. So, by pure luck, Silvia manages to send the radio signal at exactly 4pm.</p>	<p><u>Control:</u></p> <p>During World War II, Silvia, a British spy, needs to send a radio signal to her undercover contact in France. Because of how the operation is designed, it is critically important that she sends the signal at exactly 4pm on Tuesday, not a minute earlier or a minute later. For this reason, the radio transceiver comes with a clock that is usually highly reliable.</p> <p>When Silvia looks at the clock on Tuesday afternoon, it reads 4pm. Consequently, she comes to believe that it is now time to send the signal. So, at that very moment, she sends the radio signal to her undercover contact in France.</p> <p>It turns out that Sylvia was right to rely on that clock. It was indeed exactly 4pm on Tuesday afternoon. So, by relying on that clock, Silvia manages to send the radio signal at exactly 4pm.</p>
<p><u>Dependent Measures:</u></p> <p>To what extent do you agree with the following statements about the passage you just read?</p> <p>Silvia knows that it is exactly 4pm.¹⁶</p> <p>Silvia knows how to send the signal at exactly 4pm.</p> <p>Silvia intentionally sent the signal at exactly 4pm.</p>	

Table 4. Vignette used in Study 3b and the dependent variables.

5.2.3. Results

Participants who read about the agent in the control condition agreed to a greater extent that the agent knows that ($M = 39.51$, $SD = 17.45$, $n = 79$) than participants in the lucky condition ($M = -12.15$, $SD = 33.50$, $n = 78$) ($t(115.58) = 12.09$, $p < .001$, $d = 2.25$, $CI [1.78, 2.71]$).

16. As in the other studies, we planned to use the simplest knowledge-that ascription. In this vignette, the relevant action is *sending the signal at exactly 4pm* and the relevant knowledge-how ascription is “Silvia knows how to send the signal at exactly 4pm;” the corresponding knowledge of the means ascription would be “Silvia knows that sending the signal now is a means to send the signal at exactly 4pm.” For this knowledge-that ascription to be true, *in the context of the vignette*, it would have to be true that, when she sends the signal, Silvia knows that it is exactly 4pm. Thus, we planned to use this simpler knowledge-that ascription.

There was no evidence that participants who read about the agent in the control condition agreed to a greater extent that the agent knows how ($M = 44.39$, $SD = 15.42$, $n = 79$) relative to participants in the lucky condition ($M = 40.34$, $SD = 14.80$, $n = 78$) ($t(154.88) = 1.67$, $p = .09$, $d = .27$, $CI [-0.05, 0.59]$).

Participants who read about the agent in the control condition agreed to a greater extent that the agent intentionally performed the action ($M = 44.24$, $SD = 16.02$, $n = 79$) than participants in the lucky condition ($M = 13.88$, $SD = 36.82$, $n = 78$) ($t(104.87) = 6.68$, $p < .001$, $d = 1.31$, $CI [0.88, 1.72]$).

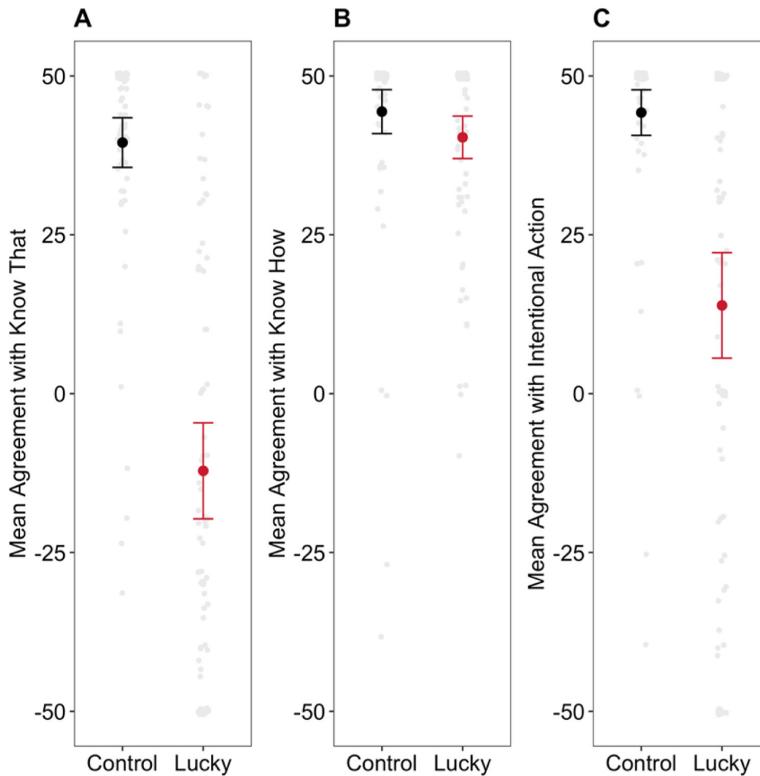


Figure 7. Mean agreement with the knowledge-that (A), knowledge-how (B), and intentional action statements in Study 3b. Error bars indicate 95% confidence intervals. Light gray points represent individual participant responses evenly jittered.

5.2.4. Discussion

In this study, epistemic luck clearly affected judgments of knowledge-that and intentionality. This effect is hard to explain on the non-epistemic plan view. The spy's success in sending the signal is a planned action. Moreover, nothing disrupts the causal path that starts on Tuesday when the spy initiates the signaling

and ends with the fulfillment of their intention, so in this sense this causal path is not deviant. Finally, despite the breaking of the clock on Monday, on Tuesday the action unfolds in accordance with the spy's plan—the plan to send the signal at the right time with her device when her clock signals 4pm.¹⁷

However, we found no evidence that epistemic luck affected judgments of knowledge-how. Given the results in the previous studies, the null effect on knowledge-how judgments was surprising.

There may, however, be a simple way to explain these unexpected results. There is a well-known difference between knowledge-how ascriptions and knowledge-that ascriptions (Pavese 2017: 377–78; Habgood-Coote 2019), which can be illustrated by contrasting the following:

- (1) Sylvia knows that it is exactly 4pm.
- (2) Sylvia knows how to send the signal at exactly 4pm.

Whereas (1) is unambiguously about Sylvia's particular situation, (2) has a generic reading, which ascribes to Sylvia knowledge of how to send the signal in normal situations. On this generic reading, (2) is equivalent to (3):

- (3) Sylvia usually knows how to send the signal at exactly 4pm.

If our participants read (2) as equivalent to (3), then it is not surprising that they were inclined to agree with (2) in the lucky condition. After all, Sylvia being a spy, and her clock being generally reliable, she usually would know how to send the signal exactly on time. Hence, we conjecture that the weak effect on knowledge-how is due to the availability of this generic reading of (2).

17. We take it that the breaking of the clock on Monday cannot plausibly interfere with the causal chain that starts when the spy initiates the signaling, since the action of sending the signal is initiated on Tuesday, whereas the breaking of the clock happens the day before. We suppose one might insist that the causal chain is deviant since it does not unfold through means that the spy intended, because not only did she intend to check a clock signaling the right time, she also intended to check a *working clock*. But this rejoinder strikes us *ad hoc*. Moreover, it risks overgeneralizing. Surely, one can perform an action intentionally even if the way one achieves one's goal does not accord *perfectly* with one's intentions. For example, I might intentionally insult someone by unleashing a string of obscenities at them, even if—contrary to my intention—they did not hear every one of the many obscenities I uttered. Of course, we do not purport to have conclusively established that there is *no* way of characterizing the relevant notion of "accordance with one's plans" on which this case would not count as having unfolded in accordance with the spy's plans; however, as we note at the outset, we take the onus to be on proponents of this alternative hypothesis to develop their account in a way that explains all of our cases without overgeneralizing.

5.3. Study 3c

To test this conjecture, we conducted a follow-up study (Study 3c), which is similar to Study 3b but controls for this ambiguity in the knowledge-how judgments (see Supplementary Materials).

5.3.1. Participants

We kept the sample size and qualifications the same. 2 participants reported not paying attention, so they were excluded. We analyzed data from the remaining 158 participants ($M_{age} = 32$, $SD = 11.8$, $Range_{age} = [18-68]$, 54% female).

5.3.2. Materials and Procedure

The procedure and design for this study was the same as in the previous studies. The only difference from Study 3b was that we clarified that the device the spy used to send the signal included a clock (see Supplementary Material for details). We then asked whether the participants agreed with (2*):

(2*) On Tuesday, Sylvia knows how to send the signal at exactly 4pm (rather than at some other time) using her device.

(2*) is arguably still ambiguous between a specific reading and a generic reading. However, the specific reading is *more* salient. First, adding ‘on Tuesday’ makes less relevant a generic reading on which the task is repeatable. Second, the qualification ‘using her device’ makes salient the change in the clock which happened overnight between Monday and Tuesday—a change that is only relevant to the truth of the specific reading of (2*). Finally, the contrast class of times is meant to make clearer that we are not simply asking whether Sylvia is generally competent at sending signals. So, we predicted an effect on knowledge-how judgments with these revised materials.

5.3.3. Results

Participants who read about the agent in the control condition agreed to a greater extent that the agent knows that ($M = 41.13$, $SD = 12.81$, $n = 79$) than participants in the lucky condition ($M = -12.84$, $SD = 33.16$, $n = 79$) ($t(100.78) = 13.49$, $p < .001$, $d = 2.69$, $CI [2.15, 3.22]$).

Participants who read about the agent in the control condition agreed to a greater extent that the agent knows how ($M = 46.49$, $SD = 9.52$, $n = 79$) than participants in the lucky condition ($M = 31.27$, $SD = 27.49$, $n = 79$) ($t(96.45) = 4.64$, $p < .001$, $d = .95$, $CI [1.52, 1.37]$).

Participants who read about the agent in the control condition agreed to a greater extent that the agent intentionally performed the action ($M = 46.40$, $SD = 13.21$, $n = 79$) than participants in the lucky condition ($M = 11.51$, $SD = 40.01$, $n = 79$) ($t(94.80) = 7.35$, $p < .001$, $d = 1.51$, $CI [1.05, 1.96]$).

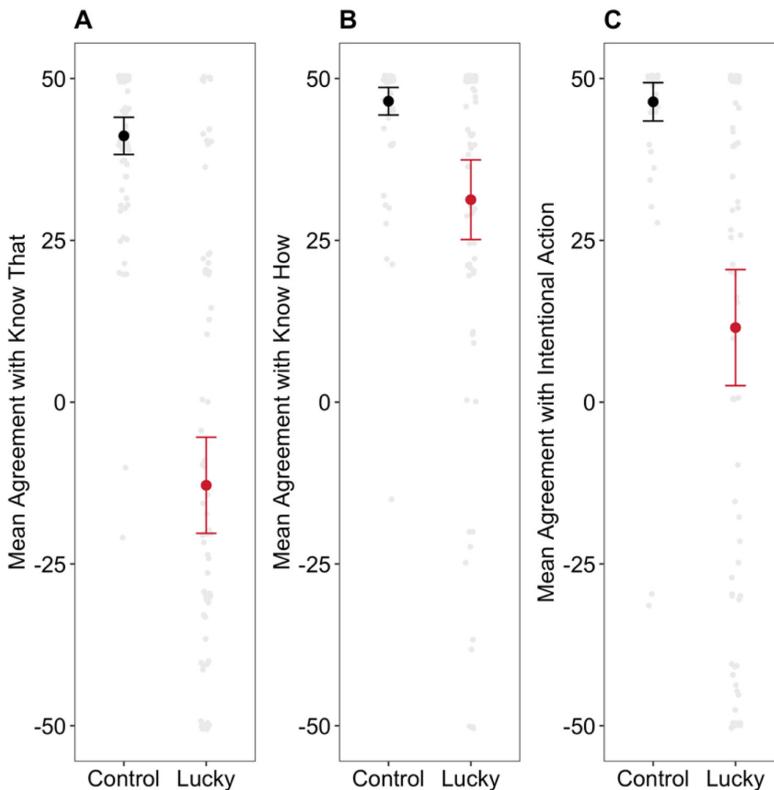


Figure 8. Mean agreement with the knowledge-that (A), knowledge-how (B), and intentional action statements in Study 3c. Error bars indicate 95% confidence intervals. Light gray points represent individual participant responses evenly jittered.

We then ran a serial mediation analysis modeled in the same way as Study 1. As theorized, the effect of condition on intentional action was serially mediated by knowledge-that and then knowledge-how (Figure 9). The indirect pathway of the effect of condition on intentional action via knowledge-that and then knowledge-how was significant ($b = -8.01$, $z = -5.70$, $p < .001$). This pathway par-

tially accounted for the overall impact of condition on intentional action with the direct effect remaining significant ($b = -25.32$, $z = -3.91$, $p < .001$).

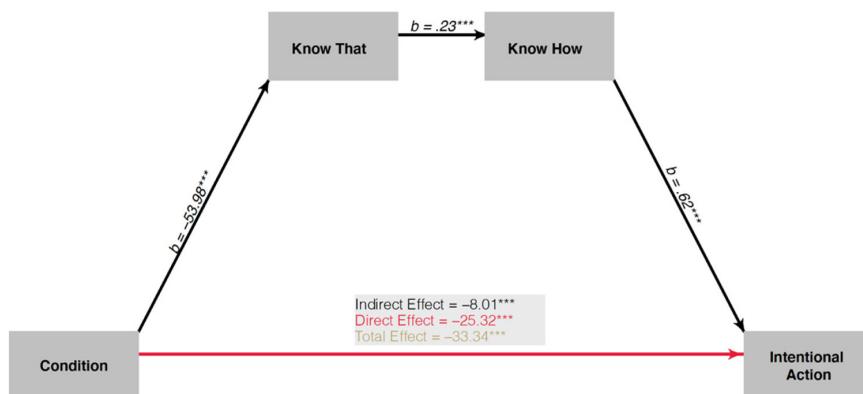


Figure 9. Mediation analysis results. *** indicates $p < .001$. Decimals rounded and then truncated to two places.

5.3.4. Discussion

This study replicated the pattern of results from the previous studies on knowledge-that and intentional action. Moreover, as predicted, we now found an effect of epistemic luck on knowledge-how judgments, suggesting that the interference of the generic reading of the knowledge-how ascription was responsible for the weak effect in Study 3b. Again, the result of the mediation analysis is predicted by Knowledge-that Hypothesis.

Notably, the effect on knowledge-how judgments is not as large as the effect on judgments of knowledge-that and intentional action.¹⁸ This finding is

18. An astute reviewer raised an interesting potential confound. The spy seems terrible at her job. She has got a critically important message to send exactly at 4 pm, yet, evidently, she only looked at the clock once, which, lucky for her, happened to be at 4 pm. The reviewer notes that if enough participants also had some such impression, the spy would strike these participants as pretty bad at her job, with no such impression arising in the control condition. A similar worry would arise for Study 3c. We note, however, that this putative confound is unlikely to have affected our findings. First, this potential confound would only explain our findings in Study 3, and it would not easily generalize to Study 1 or Study 2. Recall that we are looking for an explanation that accounts for the *general* pattern across studies. Second, note that both in the lucky condition and in the control condition, the spy only checks the clock once, so there does not seem to be any difference in the spy's competence between the two conditions, while we do see differences on our dependent measures by condition. Finally, and crucially, if the putative confound had been present, we would expect *much lower agreement* on the knowledge-how question, which on its generic reading effectively tests for the spy's general competence. Instead, what we seem to find is

expected, however, since making the specific reading more salient does not completely rule out the generic reading. Since the generic reading is still available, that might explain why the effect size for knowledge-how judgments is not as large as the effect size for knowledge-that and intentional action judgments.¹⁹

Finally, it is worth briefly discussing whether our generic-specific reading hypothesis is compatible with our earlier findings. Start with Study 2, where the agreement on the knowledge-how question is lower. Here, both the generic and specific reading of the knowledge-how ascription are arguably false. Alvin does not know how to lose this specific lottery and he presumably does not know how to lose *any* fair lottery (cf. Gibbons 2001). So, our current hypothesis nicely fits with the particularly low agreement we observed in this study. Now, consider Study 3a, where also we find low agreement on the knowledge-how ascription. In this scenario, it seems very hard to read the knowledge-how statement “Alvin knows how to bring the group to water” generically. This difficulty might be due to the fact that the specific reading—clearly available due to the presence of the definite in “bringing *the* group to water”—is made more salient by the severity of the situation. The group is “hopelessly lost,” so there is no usual situation where the agent brings the group to water. Thus, we conjecture that the severity and unusual character of the situation, together with the presence of a definite description in the prompt, might account for the lack of the generic reading in this scenario. Finally, consider Study 1. Though the agreement above the midpoint on the knowledge-how measure is surprising, we have noted (fn. 12) that this result might not be statistically significant, since it did not replicate in a supplementary experiment (See Supplementary Materials). Here, just like in Study 3a, we found low agreement on the knowledge-how measure in lucky condition. We thus speculate that the presence of a definite description in the dependent measure—that is, “the address of the person who got the promotion”—makes the false specific reading more salient, exactly like in Study 3a. In conclusion, while more investigation would be needed to fully settle this ques-

that on the generic reading of the knowledge-how statement, many still agree with the ascription to a considerable degree, showing that they did not take the spy to be generally incompetent at her job.

19. Notice that there are further interesting differences in our results that would require more investigation than we can offer here. For example, one might wonder why the mean agreement for judgments of intentional action is above the midpoint. Our hypotheses might easily explain it: if judgements of intentional action are sensitive to judgments of knowledge-how—as the Knowledge-How Hypothesis has it—and if judgments of knowledge-how are sensitive to different interpretations of the knowledge-how ascriptions, we would expect this sensitivity to affect judgments of intentional action to *some extent* too. This brings us to the further difference between the means for intentional action judgments and know-how judgments. Since we did not make any predictions about this beforehand, we choose not to explore this further difference here. But we do hope the reader will agree that this is interesting material to be explored further in future work.

tion, the overall pattern we found across studies suggest that our explanation of the recalcitrant data in Study 3b is on the right track.

6. General Discussion

Studies 1–3 display a general pattern, for in all of them, epistemic luck undermines people’s tendency to ascribe intentionality and knowledge-how. This general pattern is evidence in favor of the Knowledge-that Hypothesis over the Alternative Hypothesis.

Let us now pull the strings of our argument for this conclusion. First, we have tested a series of vignettes showing the aforementioned general pattern and in each, we have considered different incarnations of the Alternative Hypothesis—that is, possible explanations of the effect observed on intentional action judgments that differ from the Knowledge-that Hypothesis in that they are not committed to knowledge playing an explanatory role *vis à vis* intentional action. Then, in each case, we have shown that these explanations do not extend to the next vignette or have trouble accommodating the knowledge-how judgments.

So, our primary argument for the Knowledge-that Hypothesis is that different ways of construing the Alternative Hypothesis have difficulty explaining the full range of our findings. The non-deviance view of intentional action cannot account for the effects on people’s tendency to ascribe intentional action in Study 2, which does not involve causal deviance. And the non-epistemic plan view of agentive luck cannot account for the results in Study 3 where the action seems to unfold in accordance with the agent’s plans.

One might think that there is still another way of developing the Alternative Hypothesis that is in a better position to explain our results. According to a prominent view of agentive luck, *skill* is a central element for intentional action and control (e.g., Heider 1958; Shaver 1985; Mele & Moser 1994; Malle & Knobe 1997; Shepherd 2021).²⁰ Now, a skill view of agentive luck is not necessarily incompatible with the Knowledge-that Hypothesis, since skills might be thought of as involving propositional knowledge (Pavese & Beddor 2022). But, as we have seen, proponents of the skill view usually do not think of skills as having an epistemic component (cf. Mele & Moser 1994; Malle & Knobe 1997; Shepherd 2021). So, the incarnation of the Alternative Hypothesis that we are envisaging now is a non-epistemic skill view.

Now, a non-epistemic skill view might say that in Studies 1–3, the agents in the vignettes all lack the relevant skill and that already explains why participants

20. Indeed, experimental studies have shown that, *at least when the action is morally neutral*, ascriptions of intentional action are sensitive to skill possession (Malle & Knobe 1997; Nadelhoffer 2005).

are reluctant to ascribe intentional action. However, this non-epistemic skill view still faces the problem of explaining the effect of epistemic luck observed in Studies 1–3 on people’s inclination to ascribe knowledge-how.²¹ So, even this incarnation of the Alternative Hypothesis seems at an explanatory disadvantage.

Investigating the Intentional Action Question and the Knowledge-How Question together offers an overall stronger case for the Knowledge-that Hypothesis. Suppose the Knowledge-how Hypothesis is correct, and hence that intentional action is partly explained by knowledge-how. Our findings are that knowledge-how judgments are undermined by epistemic luck. Evidence for the sensitivity of knowledge-how to epistemic luck is evidence for the view that knowing how to perform some task φ at least requires propositional knowledge about how to φ —a view that we label ‘Weak Intellectualism’.²² But Weak Intellectualism together with the Knowledge-How Hypothesis entails the Knowledge-that Hypothesis.

As a final, corroborative evidence, we also executed *mediation analysis* on our experimental findings (cf. Baron & Kenny 1986; Kenny, Kashy, & Bolger 1998). In statistics and the social sciences, mediation analyses are employed in order to understand observed associations between variables in terms of an independent variable, precisely to uncover explanatory dependence relations. The Knowledge-that Hypothesis predicts that epistemic luck’s effects on judgments of intentional action are at least partially mediated by its effects on judgments of knowledge. Moreover, if Knowledge-How Hypothesis is correct, then we would expect intentional action judgments to be mediated not only by knowledge-that judgments but also by knowledge-how judgments. By contrast, the Alternative

21. The same considerations apply to a recent proposal by Quillien and German (2021) who defend a causal account of intentionality similar to Davidson’s (1963). On their ‘simple’ account, an agent did X intentionally to the extent that X was causally dependent on the degree to which the agent wanted X to happen. However, Quillien and German (2021) provide evidence that when X is the result of deviant causation, people are also disinclined to judge that it is caused by the agent’s desire. Hence, they are disinclined to judge that it is intentional. In line with this general explanation of intentionality judgments, Quillien and German (2021) might explain people’s reluctance to ascribe intentional action in Studies 1–3 in terms of people’s reluctance to judge the agents’ wanting to have caused the relevant action. Like the simple plan view and the simple skill view, however, their view gives no role to knowledge-that or knowledge-how to play in a folk theory of intentional action. So, unless it is integrated to make room for an epistemic component, Quillien and German’s (2021) view is at loss to explain both the effect of epistemic luck on knowledge-how judgments and the indirect effect of knowledge-how and knowledge-that judgments on intentional action judgments.

22. Notice that we are labeling ‘weak Intellectualism’ any view that takes knowledge-how to *at least require* some propositional knowledge. This is not the same as ‘strong intellectualism’, which *identifies* knowledge-how with propositional knowledge (see Pavese 2022a for this taxonomy). On this permissive definition of intellectualism, self-proclaimed anti-intellectualists such as Hornsby (2005), Bengson and Moffett (2011), and perhaps Habgood-Coote (2019), also count as intellectualists, for these authors allow for propositional knowledge to enter constitutively in a theory of knowledge-how. Accordingly, we take our findings concerning knowledge-how to equally support the views embraced by these authors and stronger versions of intellectualism.

Hypothesis is compatible with the effects on intentional action not being partially mediated by its effects on knowledge-that or on knowledge-how. As we saw, these predictions of the Knowledge-that Hypothesis and of Knowledge-How Hypothesis are born out.

To be clear, we do not mean to rule out that there might be yet *another* incarnation of the Alternative Hypothesis that might explain our findings in a unified fashion as thoroughly as our Knowledge Hypothesis. Our argument has been abductive, and abductive arguments are rarely conclusive. But we believe that our findings—together with the argument we have developed—raise an interesting new challenge for the opponents of the Knowledge-Hypothesis: that of offering an *alternative* unified explanation of the pattern we have observed across the three studies.

While our main argument has concerned the Knowledge-that Hypothesis, these findings have similar implications for the debate about knowledge-how. It is often claimed that knowledge-how is intuitively compatible with epistemic luck (e.g., Poston 2009; Cath 2011; Carter & Pritchard 2015). The relevant judgments are often used against the intellectualist view that knowledge-how requires propositional knowledge. But these claims had not been corroborated by empirical evidence. Two recent papers have addressed this question experimentally (Carter, Pritchard, & Shepherd 2019; Gonnerman, Mortensen, & Robbins 2021: 11333–35). It is worth briefly discussing how this experimental work relates to ours.

Carter, Pritchard, and Shepherd (2019) conducted an experimental study with the goal of assessing whether knowledge-how is sensitive to epistemic luck. In particular, Carter and colleagues presented participants with the following vignette:

Lucky Manual Charlie needs to learn how to change a lightbulb, and so he goes to the ‘how-to’ section in his local library. He finds a shelf full of identical looking books titled Home Repair. In each of these books are step-by-step instructions on the way to change a lightbulb—we’ll call the way the book describes the way ‘w’. Unbeknownst to Charlie, all the copies of Home Repair on the shelf are fakes, except for one. Pranksters have placed these copies there, and these fake copies contain mistaken step-by-step instructions on the way to change a lightbulb. Since Charlie does not know this, he reaches up and grabs the copy of Home Repair nearest to him. By sheer luck, he selects the only copy in the entire library that contains genuine and reliable step-by-step instructions for changing a lightbulb, and he reads the correct step-by-step instructions on the way to change a lightbulb. Had Charlie picked up any of the other guides—which he so easily could have—he would have believed the mistaken instructions were correct.

Carter, Pritchard, and Shepherd (2019) found that participants tend to agree with the claim that Charlie knows how to change a lightbulb and take these results to be evidence that people attribute knowledge-how in cases of epistemic luck. While Carter and colleagues did not directly test judgments of intentional action, one might well expect a similar tendency to agree with the claim that Charlie intentionally changed the lightbulb, given the assumption that there is a close connection between judgments of intentional action and knowledge-how.²³

However, there is reason to question Carter and colleagues' conclusions. They present their vignette as a case of *environmental luck*—the sort of luck present in fake barn cases (Goldman 1976; Pritchard 2005). But some philosophers have disputed whether subjects lack knowledge in fake barn cases (Gendler & Hawthorne 2005; Lycan 2006; Sosa 2007). Indeed, the extant experimental results on fake barns have been mixed, suggesting a greater willingness to ascribe knowledge in fake barn cases than in paradigmatic Gettier cases or in lottery cases.²⁴ Moreover, Carter and colleagues' own results call into question whether epistemic luck is really present in Lucky Manual: they found that participants without philosophical training generally agreed with the claim that Charlie knows *that way w* is a way of changing a lightbulb—suggesting that people are willing to attribute *both* knowledge-how and knowledge-that in this scenario.²⁵ So, all in all, it is far from clear that Lucky Manual is an authentic case of epistemic luck.

In another recent study, Gonnerman, Mortensen, and Robbins (2021: 11333–35) found that participants were more inclined to attribute knowledge-how when protagonists were described as having a justified true belief (“as having very good reason for believing”) than when they were described as having an unjustified but lucky true belief (“as having very poor reason for believing”). They noticed that this is evidence that *veritic epistemic luck* affects folk ascriptions of knowledge-how, so the ordinary concept of knowledge-how is sensitive to the presence of propositional knowledge. While we agree with their diagnosis, it is worth noting that their study does not clearly involve a Gettier case—not at least in the standard understanding of Gettier cases as involving *justified* true belief

23. This line of argument has been pushed by Cath (2015).

24. Nagel, San Juan, and Mar (2013) found that people deny knowledge in fake barns cases, and so did Wright (2010), but in the latter case the results were close to the midpoint. Colaço, Buckwalter, Stich, and Machery (2014) found that mean knowledge attributions in fake barn cases were above the midpoint, and Turri (2016; 2017) also reports finding that people tend to ascribe knowledge in fake barn cases.

25. We also note a further puzzling aspect of Carter, Pritchard, and Shepherd (2019)'s result: they found that participants were more inclined to attribute knowledge-how in the luck conditions than in the control conditions. Though very puzzling, as Carter, Pritchard, and Shepherd (2019: 717–18) themselves note, this result might have a non-epistemic explanation. Accordingly, we set this finding aside for the purpose of our discussion.

that falls short of knowledge—nor does the structure of the case exactly match that of Gettier’s original cases.

Our studies differ from Carter, Pritchard, and Shepherd’s (2019) in that we focus on uncontroversial cases of epistemic luck; our studies differ from Gonnerman, Mortensen, and Robbins’s (2021) Study 1 in that we focus on several new cases of epistemic luck which structurally mirror Gettier’s original cases (Study 1), lottery cases (Study 2), Dharmottara’s desert traveler case (Study 3a), and Russell’s stopped clock case (Study 3c). In all of these cases, we found that epistemic luck affects people’s tendency to ascribe knowledge-how. These results fit nicely with the idea, voiced by others too (Bengson, Moffett, & Wright 2009; Gonnerman, Mortensen, & Robbins 2018), that the ordinary concept of knowledge-how is at least partly intellectualist, since this position too predicts an effect of epistemic luck on ordinary ascriptions of knowledge-how. Finally, our findings speak against those who have argued that the folk are purely anti-intellectualists with respect to knowledge-how (Harmon & Horne 2016).²⁶

7. Conclusions

Our investigation of the relation between epistemic luck and agentive luck revealed that epistemic luck undermines both intentional action and knowledge-how. These results bear on two central debates in epistemology and action theory: the debate on the relation between knowledge-that and intentional action and the related debate on the relation between knowledge-that and knowledge-how. We have argued that these results raise a new challenge for the opponent of

26. Our studies might be relevant to assess another related line of argument against Intellectualism. Cath (2015) starts from the supposedly intuitive observation that intentional action is not sensitive to epistemic luck. Assuming the Knowledge-How Hypothesis, Cath (2015) concludes that in epistemic luck cases, one can still know how to perform the relevant action (cf. also Shepherd 2021: 77–81). Our findings give reason to doubt this general argumentative strategy against Intellectualism, for they show that people are often reluctant to ascribe intentional action in uncontroversial cases of epistemic luck. Some readers might wonder about Cath’s (2011; 2015) own case of alleged epistemic luck (the “Lucky Lightbulb” example). Earlier in the text, we discussed the main experimental study based on Cath’s case (Carter, Pritchard, & Shepherd 2019) and explained why their results are compatible with Intellectualism. As we noted there, participants actually tended to ascribe both knowledge that and knowledge how, suggesting that participants did not regard this as an uncontroversial case of epistemic luck. However, there are some subtle differences between Cath’s original case and the version of the vignette tested by Carter, Pritchard, and Shepherd (2019); some may wonder whether these differences will affect judgments. We think that it is not clear that Cath’s original case is a genuine case of epistemic luck either, though this point would need much more discussion than we can provide here. See Pavese and Beddor (2022) and Pavese (in press).

the Knowledge-Hypothesis: that of providing an alternative unified explanation for the general pattern we have observed in our three studies.

For those convinced by our general argument, our findings also suggest a more general lesson on how to theorize about the relation between knowledge and action. If we are right that agentive luck is sometimes epistemic, then one route forward in action theory is to use epistemologists' understanding of epistemic luck to make progress on the nature of intentional action. In this way, action theory can benefit from the insights of epistemologists.

Finally, our results have implications for work on the role of knowledge representation in folk psychology. Recently, several authors have argued that the capacity to distinguish knowledge from true belief (and justified true belief) plays a central role in our ability to understand and predict people's behavior (Williamson 2000; Gibbons 2001; Nagel 2013; Phillips et al. 2021). However, most of this work does not focus on recognizing intentional action. The results in this paper thus introduce a new dimension along which knowledge representation influences folk psychology, since they suggest that knowledge representation drives people's capacity to recognize certain actions as intentional. All in all, we hope that these results will motivate a renewed interest for the place for knowledge in the philosophy of action and in the philosophy of mind.

Acknowledgements

We would like to thank Cornell College of Arts and Sciences for funding this work through Carlotta Pavese's research account. The work was divided as follows. Carlotta Pavese is responsible for the main idea, for the structure of the argument, the effective writing of the paper, and for the supervision of every aspect of the project; Paul Henne has led the experimental part of the project, run and analyzed the experiments, and helped a lot also with structuring the main argument and editing the paper; Bob Beddor has contributed to the brainstorming, to the vignettes, and to the revisions of the paper. Particular thanks go to Joshua Knobe for providing essential advice at several stages. We also would like to thank Jennifer Nagel, Shaun Nichols, and Joshua Alexander for generous and illuminating comments on initial drafts, as well as the editors and two anonymous reviewers of *Ergo*. Finally, we are indebted for feedback to audiences at the Neurophilosophy Forum at Georgia State University, SSPP 2022, Illinois Urbana Champaign, and the Workshop on Action, Intention, and Language 2022.

References

- Adams, Frederick, and Alfred Mele (1989). The Role of Intention in Intentional Action. *Canadian Journal of Philosophy*, 19(4), 511–31.
- Anscombe, Gertrude Elizabeth Margaret (1958). *Intention*. Basil Blackwell.
- Baron, Reuben M., and David Kenny (1986). The Moderator-Mediator Variable Distinction in Social Psychological Research—Conceptual, Strategic, and Statistical Considerations. *Journal of Personality and Social Psychology*, 51(6), 1173–82.
- Beddor, Bob, and Carlotta Pavese (2022). Practical Knowledge without Luminosity. *Mind*, 131(523), 917–34.
- Bengson, John, and Marc Moffett (2011). Non-Propositional Intellectualism. In John Bengson and Marc Moffett (Eds), *Knowing How: Essays on Knowledge, Mind, and Action* (161–95). Oxford University Press.
- Bengson, John, Marc Moffett, and Jennifer Wright (2009). The Folk on Knowing How. *Philosophical Studies*, 142(3), 387–401.
- Bhatt, Rajesh (1999). Ability Modals and Their Actuality Entailments. *WCCFL 17 Proceedings*, 17, 74–87.
- Brand, Myles (1986). Intentional Actions and Plans. *Midwest Studies in Philosophy*, 10(1), 213–30.
- Cappelen, Herman (2012). *Philosophy without Intuitions*. Oxford University Press.
- Carter, Adam, and Duncan Pritchard (2015). Knowledge-How and Epistemic Luck. *Noûs*, 49(3), 440–53.
- Carter, Adam, Duncan Pritchard, and Joshua Shepherd (2019). Knowledge-How, Understanding-Why and Epistemic Luck: An Experimental Study. *Review of Philosophy and Psychology*, 10, 701–34.
- Cath, Yuri (2011). Knowing How without Knowing That. In John Bengson and Marc Moffett (Eds.), *Knowing How: Essays on Knowledge, Mind and Action* (113–35). Oxford University Press.
- Cath, Yuri (2015). Revisionary Intellectualism and Gettier. *Philosophical Studies*, 172(1), 7–27.
- Chisholm, Roderick (1966a). Freedom and Action. In Keith Lehrer (Ed.), *Freedom and Determinism* (16–23). Random House.
- Chisholm, Roderick (1966b). *Theory of Knowledge*. Prentice Hall.
- Colaço, David, Wesley Buckwalter, Stephen Stich, and Edouard Machery (2014). Epistemic Intuitions in Fake-Barn Thought Experiments. *Episteme*, 11(2), 199–212.
- Davidson, D. (1963). Actions, Reasons, and Causes. *The Journal of Philosophy*, 60(23), 685–700.
- Davidson, D. (1971). Agency. In R. Binkley, R. Bronaugh, and A. Marras (Eds.), *Agent, Action, and Reason* (1–37). University of Toronto Press.
- Davidson, Donald (1973). Freedom to Act. In *Essays on Freedom and Action* (137–56). Routledge.
- Dharmottara (770). *Explanation of (Dharmakirti's) Ascertainment of Valid Cognition*.
- Engel, Mylan (1992). Is Epistemic Luck Compatible with Knowledge? *Southern Journal of Philosophy*, 30(2), 59–75.
- Friedman, Ori, and John Turri (2015). Is Probabilistic Evidence a Source of Knowledge? *Cognitive Science*, 39, 1062–80.

- Gendler, Tamar, and John Hawthorne (2005). The Real Guide to Fake Barns: A Catalog of Gifts for Your Epistemic Enemies. *Philosophical Studies*, 124, 331–52.
- Gettier, Edmund (1963). Is Justified True Belief Knowledge? *Analysis*, 23(6), 121–23.
- Gibbons, John (2001). Knowledge in Action. *Philosophy and Phenomenological Research*, 62(3), 579–601.
- Goldman, Alvin (1976). Discrimination and Perceptual Knowledge. *Journal of Philosophy*, 73, 771–91.
- Gonnerman, Chad, Kaija Mortensen, and Jacob Robbins (2018). The Ordinary Concept of Knowledge How. *Oxford Studies in Experimental Philosophy*, 2, 104–15.
- Gonnerman, Chad, Kaija Mortensen, and Jacob Robbins (2021). *Knowing How* as a Philosophical Hybrid. *Synthese*, 199(3), 11323–54.
- Habgood-Coote, Joshua (2019). Knowledge-How, Abilities, and Questions. *Australasian Journal of Philosophy*, 97(1), 86–104.
- Harmon, Ian, and Zachary Horne (2016). Evidence for Anti-Intellectualism about Know-How from a Sentence Recognition Task. *Synthese*, 193(9), 2929–47.
- Hawley, Katherine (2003). Success and Knowledge How. *American Philosophical Quarterly*, 40(1), 19–31.
- Hawthorne, John (2004). *Knowledge and Lotteries*. Oxford University Press.
- Heider, Fritz (1958). *The Psychology of Interpersonal Relations*. Wiley.
- Hornsby, Jennifer (2005). Semantic Knowledge and Practical Knowledge. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 79, 107–45.
- Kenny, David, Deborah Kashy, and Nial Bolger (1998). Data Analysis in Social Psychology. In Daniel Gilbert, Susan Fiske, and Gardner Lindzey (Eds.), *Handbook of Social Psychology* (4th ed., Vol. 1, 233–65). McGraw-Hill.
- Knobe, Joshua (2003). Intentional Action and Side Effects in Ordinary Language. *Analysis*, 63, 190–93.
- Lycan, William (2006). On the Gettier Problem Problem. In Stephen Hetherington (Ed.), *Epistemology Futures*, Oxford University Press, New York, pp. 148–68.
- Machery, Edouard, Stephen Stich, David Rose, Amita Chatterjee, Kaori Karasawa, Noel Struchiner, Smita Sirker, Naoki Usui, and Takaaki Hashimoto (2017). Gettier across Cultures. *Noûs*, 51(3), 645–64.
- Malle, Bertram, and Joshua Knobe (1997). The Folk Concept of Intentionality. *Journal of Experimental Social Psychology*, 33, 101–21.
- Mele, Alfred, and Paul Moser (1994). Intentional Action. *Noûs*, 28(1), 39–68.
- Moss, Sarah (2020). Replies to Edgington, Pavese, and Campbell-Moore and Konek. *Analysis*, 80(2), 356–70.
- Nadelhoffer, Thomas (2005). Skill, Luck, Control, and Intentional Action. *Philosophical Psychology*, 18(3), 341–52.
- Nagel, Jennifer (2013). Knowledge as a Mental State. *Oxford Studies in Epistemology*, 4, 273–306.
- Nagel, Jennifer, Valerie San Juan, and Raymond A. Mar (2013). Lay Denial of Knowledge for Justified True Beliefs. *Cognition*, 129(3), 652–61.
- Pavese, Carlotta (2017). Know-How and Gradability. *The Philosophical Review*, 126(3), 345–83.
- Pavese, Carlotta (2020). Probabilistic Knowledge in Action. *Analysis*, 80(2), 342–56.

- Pavese, Carlotta (2021a). Knowledge, Action, and Defeasibility. In Jennifer Brown and Mona Simion (Eds.), *Reasons, Justification, and Defeaters* (177–200). Oxford University Press.
- Pavese, C. (2021b). Know-How, Action, and Luck. *Synthese*, 198(suppl. 7), 1595–617.
- Pavese, Carlotta (2021c). Knowledge and Mentality. *Philosophical Perspectives*, 35(1), 359–83.
- Pavese, Carlotta (2022a). Knowledge How. In Edward N. Zalta and Uri Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/fall2022/entries/knowledge-how/>
- Pavese, Carlotta (2022b). Practical Knowledge First. *Synthese*, 200(375), 1–18.
- Pavese, Carlotta (in press). Factive Mindreading in the Folk Psychology of Action. In Arturs Logins and Jacques-Henri Vollet (Eds.), *Putting Knowledge to Work*. Oxford University Press.
- Pavese, Carlotta, and Bob Beddor (2022). Skills as Knowledge. *Australasian Journal of Philosophy*. Advance online publication. <https://doi.org/10.1080/00048402.2022.2056753>
- Phillips, Jonathan, Wesley Buckwalter, Fiery Cushman, Ori Friedman, Alia Martin, John Turri, Laurie Santos, and Joshua Knobe (2021). Knowledge before Belief. *Behavioral and Brain Sciences*. Advance online publication. <https://doi.org/10.1017/S0140525X20000618>
- Poston, Ted (2009). Know How to Be Gettiered? *Philosophy and Phenomenological Research*, 79(3), 743–47.
- Pritchard, Duncan (2005). *Epistemic Luck*. Oxford University Press.
- Pritchard, Duncan (2021). Anti-Luck Epistemology and Pragmatic Encroachment. *Synthese*. *Syntheses*, 199, 715–29. <https://doi.org/10.1007/s11229-020-02703-2>
- Quillien, Tadeq, and Tamsin German (2021). A Simple Definition of ‘Intentionally’. *Cognition*, 214, 104806.
- Riggs, Wayne (2007). Why Epistemologists Are So Down on Their Luck. *Synthese*, 158(3), 329–44.
- Russell, Bertrand (1948). *Human Knowledge: Its Scope and Limits*. George Allen & Unwin.
- Ryle, Gilbert (1949). *The Concept of Mind*. Hutchinson & Co.
- Setiya, Kieran (2012). Knowing How. *Proceedings of the Aristotelian Society*, 112(3), 285–307.
- Shaver, Kelly G. (1985). *The Attribution of Blame*. Springer-Verlag.
- Shepherd, Joshua (2021). *The Shape of Agency: Control, Action, Skill, Knowledge*. Oxford University Press.
- Shepherd, Joshua, and Adam Carter (2021). Knowledge, Practical Knowledge, and Intentional Action. *Ergo*, 9, 556–83.
- Sosa, Ernest (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge* (Vol. 1). Oxford University Press.
- Stanley, John, and Timothy Williamson (2001). Knowing How. *Journal of Philosophy*, 98(8), 411–44.
- Turri, John (2013). A Conspicuous Art: Putting Gettier to the Test. *Philosophers’ Imprint*, 13(10), 1–16.
- Turri, John (2016). Knowledge and Assertion in “Gettier” Cases. *Philosophical Psychology*, 29(5), 759–75.
- Turri, John (2017). Knowledge Attributions in Iterated Fake Barn Cases. *Analysis*, 77(1), 104–15.

- Turri, John, Wesley Buckwalter, and Peter Blouw (2015). Knowledge and Luck. *Psychonomic Bulletin & Review*, 22(2), 378–90.
- Vekony, Romy, Alfred Mele, and David Rose (2020). Intentional Action without Knowledge. *Synthese*, 199, 1231–43.
- Williamson, Timothy (2016). Philosophical Criticisms of Experimental Philosophy. In Justin Sytsma, Wesley Buckwalter (Eds.), *A Companion to Experimental Philosophy* (22–36). John Wiley & Sons.
- Williamson, Timothy (2000). *Knowledge and Its Limits*. Oxford University Press.
- Wright, Jennifer (2010). On Intuitional Stability: The Clear, the Strong, and the Paradigmatic. *Cognition*, 115(3), 491–503.