

# Factive Mindreading in the Folk Psychology of Action

Carlotta Pavese<sup>1</sup>

(This is the penultimate version)

forthcoming in Arturs Logins and Jacques-Henri Vollet (eds.) *Putting Knowledge to Work: New Directions for Knowledge-First Epistemology*, Oxford University Press

## Abstract

In the recent literature, several authors have argued that the capacity to track factive mental states plays a central role in explaining our ability to understand and predict people’s behavior (Nagel 2013; Nagel 2017; Phillips & Norby 2019; Phillips *et al.* 2020; Westra & Nagel 2021). The topic of this chapter is whether this capacity also enters into an explanation of our ability to track skilled and intentional actions.

### 1. Introduction

A recent and growing literature suggests that, since early age, humans possess the capacity to track factive mental states—the capacity for *factive mindreading*—and that this capacity plays an important role in our mindreading practices (Perner, 1993; Nagel, 2013; Nagel, 2017; Phillips & Norby, 2019; Phillips *et al.*, 2020; Westra & Nagel, 2021).<sup>2</sup> Infants can track perceptual knowledge in others since they are six-months old (Phillips *et al.* 2020). Pre-school children learn factive epistemic vocabulary (“know”, “knowledge”) before they learn how to correctly use non-factive epistemic vocabulary (“think”, “thought”, “belief”) (Nagel, 2013; Phillips *et al.*, 2020). Successful attribution of

---

<sup>1</sup> I am grateful to Arturs Logins and an anonymous referee for comments that have improved this chapter.

<sup>2</sup> Indeed, recent work has made a persuasive case that primates are also capable of factive mindreading and they employ it when predicting their mates’ behavior (Horschler, Santos, and MacLean, 2019; Phillips *et al.*, 2021; Williamson, this volume). There is also some evidence that some primates are capable of recognizing rational and goal-directed action (e.g., Povinelli *et al.* 1998; Wood *et al.* 2007; though see Tomasello 2000 for a dissenting voice). On the assumption that nonhuman primates *can* recognize intentional action, it is an interesting question, though one that I will not be able to explore here, whether their recognizing intentional action involves factive mindreading.

factive mental states often emerges when children are aged 3, well before successful attribution of mere belief states (Nagel 2013; Phillips & Norby 2019; Phillips et al. 2021).

Distinctive of factive mindreading is that it does not involve *decoupling*—the cognitive process whereby a subject represents the content of another agent’s mental state in a way that conflicts with the content of the subject’s primary and action-guiding representation of the world (Westra & Nagel, 2021: 2-3). Because factive mindreading does not require decoupling, it offers a less costly and all in all more efficient cognitive device for a variety of cognitive purposes, from predicting the behavior of others to finding the best communicative strategies in conversational settings (Phillips et al. 2021; Westra & Nagel 2021).

Thus far, the discussion on factive mindreading has proceeded independently of the study of the *folk psychology of action*—of our capacity to track intentional and skilled actions. Yet, the folk psychology of action have been studied extensively in the last three decades: a huge amount of experimental work has been devoted to the question as to how people track intentional action, specifically in connection with the concept of responsibility (e.g., Malle and Knobe, 1997; Knobe, 2003a; 2003b; Knobe and Burra, 2006; Nadelhoffer, 2004; 2005; Jacob et al., 2012; Guglielmo et al., 2009; Guglielmo & Malle, 2010). Some of this work has even focused on how people track skilled action (e.g., Malle & Knobe, 1997; Nadelhoffer, 2004; 2005). This literature suggests that adult human beings are quite proficient in tracking intentional and skilled behavior, suggesting that this ability—which has been shown to be stable across different populations and interculturally robust (Knobe and Burra, 2006; Knobe, in press)—is a core component of social cognition. Yet, this research has not tackled the question as to whether factive mindreading plays any role *vis a vis* our ability to track intentional and skilled action. Indeed, the role of factive mindreading in the folk psychology of action has thus far remained rather unexplored.

And yet the same considerations of expediency and efficiency that makes factive mindreading an helpful cognitive device for predicting the behavior of others would presumably make it also an effective device for tracking intentional and skilled behavior.

On this hypothesis, the easiness and reliability of factive mindreading would contribute to explaining the remarkable easiness and reliability with which we can track intentional and skilled action. Indeed, given how reliable and efficient we are at tracking intentional and skilled action, and given how pervasive factive mindreading seems to be in social cognition, it would be quite surprising if factive mindreading was found to play no role in the folk psychology of action.

Here is how the technical terms will be used going forward. Let ‘factive mindreading’ refer to the ability to track factive mental states in ourselves and in others. A *factive mental state* is one that obtains only if its content is a true proposition. Clear examples of factive mental states are *understanding that*, *being aware that*, *seeing that*, *remembering that*, and *knowing that*. If every factive mental state involves knowledge (Williamson, 2000), then factive mindreading boils down to the ability for knowledge representation. If not every factive mental state involves knowledge, then factive mindreading is a more general ability than the capacity for tracking knowledge. Let ‘social cognition’ refer to the way in which we process and use information in social contexts to predict our behavior and that of others. Let ‘folk psychology’ refer to the general ability that humans have to predict and understand others’ behavior through mindreading (whether factive or not). Let ‘folk psychology of action’ refer to the specific ability to recognize and ascribe intentional and skilled action.

We are now in position to state my goals with some precision. This chapter aims to single out, to clarify, and to explore the question: what role does factive mindreading play in the folk psychology of action? This question is part of the more general factive mindreading program—the exciting program in the theory of mind that aims at understanding the role of factive mindreading in folk psychology and in social cognition. A natural continuation of the factive mindreading program is to look at whether factive mindreading plays a role not only when it comes to predicting people’s behavior but also when it comes to recognizing and ascribing intentional and skilled action. This is the sense in which I want to ask about the role of factive mindreading in the folk psychology of action. Of course, the study of factive mindreading is in its early stages; moreover, the

evidence that I will overview is preliminary. My main aim in this chapter is to raise the central questions, to assess *prima facie* reasons for skepticism, to sketch the map of the terrain, and to individuate general issues that might spring further investigation.

Here is the plan. §2 overviews knowledge-first approaches to the philosophy of action, which motivated some of the experimental work surveyed in this paper. §3 critically discusses some recent empirical work that might be taken to suggest, erroneously I think, that people are not factive mindreaders when ascribing intentional and skillful action. §4 overviews some experimental work about the role of know-how in the folk psychology of action and highlights its relevance for the question at hand; §5 overviews some experimental work that provides positive evidence for the role of propositional knowledge in the folk psychology of action. §6 identifies open questions for further investigation.

## 2. Knowledge-First Action Theory

The idea that factive mental states such as knowledge or awareness enter constitutively in a theory of intentional action has a rich history: it has roots in Aristotle, and it plays a starring role in the work of Elizabeth Anscombe (e.g., Anscombe, 1958; Hampshire, 1959; Gibbons, 2001; Thompson, 2011; Rodl, 2011; Pavese, 2018; 2020; 2021; 2022; Kneer, 2021; Beddor & Pavese, 2022; Pavese & Beddor, 2022). This section overviews general trends in the literature on the role of knowledge and factive mental states in a theory of intentional action, including recent knowledge-first approaches to action theory (Gibbons, 2001; Pavese, 2013; 2016; 2018; 2020; 2021a; 2021b; 2022).

There are many different ways in which knowledge might enter into the theory of action. According to Anscombe, practical knowledge enters as a requirement on intentional action—intentional action requires knowing what one is doing when doing it. All indicates that she understood practical knowledge in terms of propositional knowledge—as knowing *that something is the case* (Schwenkler, 2019; Pavese, 2022: §3). Many have followed Anscombe on this insight. Even those who have challenged the idea that a practical knowledge condition applies to every instance of intentional action

tend to retain *some* version of a practical knowledge condition. For example, Setiya (2008) questions whether basic actions require practical knowledge but defends a version of a practical knowledge condition for non-basic complex actions. Others argue that at least *self-awareness*, rather than knowledge, is necessary for intentional action (cf. Rodl, 2007) but they do understand self-awareness factively such that, if one is self-aware that one is F-ing, then one is F-ing (cf. Schwenkler, 2019).<sup>3</sup> Paul (2009) argues against a practical knowledge condition on habitual and automatic action but deems it plausible for deliberate conscious actions. Others retort that practical knowledge needs not be consciously entertained at the time of the performance and yet might still play a role in explaining its intentionality (Setiya 2008; Pavese 2022).

Assuming that skilled action is always or typically intentional (Pavese, 2013; 2016; Pavese & Beddor, 2022; Pavese 2022), one might expect practical knowledge to also enter into a theory of skilled action. Consider Pollock's first 'drip' painting, *Free Form* (1946). According to an accredited reconstruction, Pollock began by painting an entire canvas red and then added black and white tangles and pools of paint by flinging and dripping diluted oil paint from a brush or stick. Suppose my 5 year-old niece, armed with a canvas, a brush, and white, black, and red oil colors, started dripping color on the canvas randomly and without following a plan. While it is, presumably, not impossible for my niece to drip oil on a canvas in such a way that the outcome turns out to look exactly like Pollock's *Free Form*, if she did manage to reproduce, by pure luck, the exact same colors and forms on a canvas, the difference between her so doing and Pollock's performance would be no less noteworthy. Indeed, even if their bodily movements turn out to be exactly alike, Pollock's execution would be skilled throughout; my niece's would instead be accidental, fortuitous, and unskilled. We would describe this difference by saying that, while Pollock knew, or was aware of, what he was doing, my niece did not. Examples such as this motivate thinking that factive mental states such as practical knowledge ought to enter in a theory of skilled action too.

---

<sup>3</sup> The distinction between awareness and knowledge is often left unclear, however, and might be due to an unnecessarily doxastic conception of knowledge (cf Williamson, this volume).

While many have recognized a role for practical knowledge in a theory of intentional or skilled action, a parallel debate on the nature of know-how has recognized a role for know-how to play in action theory too. Ryle (1949) argued that what distinguishes skilled and intentional action from unskilled unintentional action is whether the action manifests know-how. As an example, Ryle (1949) compared a clown who tumbles on purpose, and whose tumbling is skilled, to a klutz who does not tumble on purpose, and whose tumbling is not skilled. Several authors have followed Ryle on this point. Stanley and Williamson (2001:442-3) claim that intentional actions are employments of know-how. Setiya (2012) argues that intentional action is guided by either knowing how to perform the action itself or knowing how to perform some other action that is a means to perform it. Gibbons (2001) argues that the concept of knowledge, rather than true beliefs, is required for intentional action, on the grounds e.g., that one cannot intentionally lose or win a fair lottery, on the plausible assumption that one cannot know how to win or lose a fair lottery. Indeed, this philosophical intuition that intentional action is closely connected to know-how or skills has been confirmed to be shared widely by non-philosophers too (Bengson et al., 2009). In a representative case, Irina, a novice figure skater, decides to try a complex jump called the Salchow. Irina is seriously mistaken about how to perform a Salchow. However, Irina also suffers from a severe neurological abnormality that often makes her perform actions that are different than those she thinks she is performing. By a fluke, whenever Irina attempts to do a Salchow (in accordance with her misconceptions), her neurological condition causes her to perform the correct sequence of moves, so she ends up successfully performing a Salchow. Intuitively, Irina does not skillfully or intentionally perform the Salchow nor does she know how to perform it.

The connection between know-how and intentional action offers a new line of argument for the centrality of factive mental states in a theory of intentional and skilled action. On the basis of this insight, Pavese (2013; 2018; 2020; 2021ab; 2022) has developed a systematic knowledge-first account of intentional and skilled action. Start

with standard action theory, according to which one intentionally Fs when one has a plan to F, where a plan to F is a knowledge state that specifies the means to F:

(Intentionality/knowledge): If s intentionally Fs, then there are some means  $m_1, \dots, m_n$  to F such that s knows that  $m_1, \dots, m_n$  are means for oneself to F.

Now consider:

(Know-how/Intentionality): If s intentionally Fs, s knows how to F.

Suppose that (Intentionality/Knowledge) is true so that the intentionality of an action is to be explained at least in part in terms of propositional knowledge. Then by (Know-how/Intentionality) and (Intentionality/Knowledge), we get that, if one intentionally Fs, one both knows how to F and one has propositional knowledge of some means to F:

(Know-how, Intentionality, Knowledge): If s intentionally Fs, s both knows how to F and for some means  $m_1, \dots, m_n$ , s knows that means  $m_1, \dots, m_n$  are means for oneself to F.

As Pavese (2020) points out, a picture of know-how as propositional knowledge-involving provides the best explanation for why (Know-How, Intentionality, Knowledge) should hold. According to this explanation, (Know-How, Intentionality, Knowledge) is true not just because of a coincidental aligning of propositional knowledge and know-how in intentional action. Rather, its truth is grounded on the very nature of know-how, as involving propositional knowledge of the means for performing an action. On this knowledge-first account of action, intentional action is itself *knowledgeable*.

Here is yet another theoretical argument for thinking that know-how is key to understanding intentional action (Pavese, forthcoming). Intentional actions are *creditable*

to their agents. If a politician's policy ends up benefiting the environment but that outcome was not sought out and entirely coincidental, we would not say that the politician is to be credited for benefiting the environment. Indeed, intentionality seems a necessary condition on creditability. On the other hand, actions manifesting know-how are *paradigmatically* creditable to their agents (cf., e.g., Sosa 2007). Consider a sprinter that scores a time record unhelped by luck or the wind, just by virtue of their know-how. In this case, the athlete should be credited for that success. On the other hand, if the success was due to luck or to an excessively strong wind, the time record should not be awarded to them. If intentional action is an employment of (some degree of) know-how—we would have it explained why intentional actions are creditable to their agents: in virtue of their being the employment of (some degree of) know-how.<sup>4</sup>

Similar considerations apply from intentional action to skilled action, given the tight relation between skills and know-how. The idea that skills are closely connected to know-how has considerable pre-theoretic appeal. Ordinary speakers seem to use the terms, “know-how” and “skills,” interchangeably, as do many philosophers and scientists (Ryle, 1949; Heider 1958; Dreyfus, 2008; Setiya, 2012; Pavese, 2016; Cath, 2020). Skills are routinely ascribed by ascriptions of know-how even in those languages that lack a dedicated word for skills, such as Italian and French (Pavese 2016). Indeed, it sounds incoherent to affirm that someone is skilled at  $\phi$ -ing while denying that they know how to  $\phi$ . For example, it is weird to affirm that Mary is a skilled swimmer but she does not know how to swim or that Mark does not know how to make risotto but is nonetheless skilled at it.

This plethora of considerations provide some evidence that skills entail know-how. Does know-how entail skills? This is more controversial. Consider Mark, a bumbling chef. Admittedly, it seems coherent to say “Mark knows how to make risotto, but I would not say he is *skilled* at it.” However, this can be explained on the assumption that

---

<sup>4</sup> The claim that intentional action manifests know-how or skill (at least to some degree) does not commit one to taking intentional action to always be a skilled action, since an action might manifest *some* degree of know-how or skill *without being skilled*—without satisfying the standards for skillfulness in a context. See Pavese and Beddor (2022) on this point.



gradable adjectives such as “skilled” require that their argument exceeds a certain threshold (Pavese 2016). While one may know how to  $\varphi$  without counting as skilled, it is nonetheless true that knowing how to  $\varphi$  *sufficiently well* entails being skilled at  $\varphi$ -ing. Note that it *does* seem incoherent to say “Mark knows how to make risotto *very well*, but I would not say he is skilled at it.”

We have thus found no good reason to abandon the idea that skills are a species of know-how. Indeed, plausibly, skills are know-how—as Ryle (1949) first alleged—and they differ from some abilities, such as instincts, in that they are learned and from yet other abilities, such as habits, in that they are acquired through deliberate practice (Pavese forthcoming). Suppose that S skillfully  $\varphi$ s. Presumably, that means that S’s success in  $\varphi$ -ing is guided by S’s skill at  $\varphi$ -ing. If skills are a species of know-how, it also follows that S’s  $\varphi$ -ing is guided by S’s knowledge of how to  $\varphi$ . If so, then there is room for a view on which skilled action is guided by a factive mental state about the means for acting.

To sum up: in several current strands in action theory, people have recognized a variety of roles for factive mental states to play in a theory of intentional and skilled action. Some take practical knowledge to be a factive mental state central to a theory of intention-in-action and of intentional action. Knowledge-first accounts of intentional action link intentional action with know-how and know-how with propositional knowledge about the means. Finally, given the link between skills and know-how, factive mental states might have a role to play in a theory of skills and skilled action too.

The discussion in this section has been limited to *philosophical theories of skilled and intentional action*. However, it would be surprising if such theories dramatically diverged from the *folk psychology of action*. After all, such philosophical theories are routinely tested against judgments that are supposed to be widely shared or even commonsensical. So an emerging knowledge-first landscape in the theory of action offers some initial motivations to investigate the question whether factive mindreading plays a role in the folk psychology of action too.

This said, it is also important to notice that, although knowledge-first approaches to the theory of action encourage thinking that we are factive mindreaders when it comes to ascribing intentional and skilled action, a view according to which we are factive mind readers in the folk psychology of action does *not* commit one to a knowledge-first approach to the theory of action. For it might be that the capacity to ascribe factive mental states plays a role in our capacity to track intentional action in a variety of different contexts, even though it does not play such a role in *every* context. Moreover, one ought not rule out *a priori* the possibility that factive mental states other than knowledge—such as, for example, awareness—play a more starring role in folk-psychology than knowledge itself.

With these complexities in mind, let us begin by looking at some preliminary reasons for doubting a substantive role for factive mind reading in the folk psychology of action.

### 3. Reasons for Skepticism

In recent experimental work, Vekony et al. (2021) have tested whether people's judgments of intentional and skilled action require practical knowledge in Anscombe's sense. The authors' primary intent is to put pressure on the *philosophical* theory of intentional action on which practical knowledge is a necessary condition for intentional action. However, these results are also relevant to whether factive mindreading plays a role in the folk psychology of action, since they concern people's judgments about intentional action and whether these judgements go together with judgments of practical knowledge and awareness—both plausibly factive mental states. It is worth looking at the experiments in some detail, since these are the only published experimental studies that might seem to negatively answer the main question of this chapter.

Vekony et al. (2021) focus on two *distinct* cases—one concerning skilled (and non-automatic, controlled action) and the other concerning habitual (automatic) action. In the first experiment, they assigned vignettes such as:

[Basketball] Andy is a 92% free throw shooter. One evening, he is at the gym practicing his free throws. He lines up and takes the shot. But just as the ball leaves his hands, lightning strikes the building. The power goes out and it is pitch black. There is also a loud clap of thunder. Due to this, Andy could not see or even hear whether he made the shot. He is completely unaware of whether he sank the shot. But he did in fact sink the shot.

Beneath the scenario, on the same page, participants responded to three test statements (presented in random order):

When Andy was sinking the shot, he knew that he was sinking it. (Knowledge)

While Andy was sinking the shot, he was aware of sinking the shot. (Awareness)

Andy intentionally sank the shot. (Intentionality)

Vekony et al. (2021) found that in cases involving skilled action, participants were overwhelmingly inclined to say that Andy intentionally sank the shot but they were considerably less inclined to say that Andy knew, or was aware, that he was sinking the shot.<sup>5</sup>

This experiment nicely details ordinary intuitions about whether intentional action  $\Phi$  requires knowledge, or awareness, that one is  $\Phi$ -ing. Does this experiment show that accurately recognizing intentional action does not require people to recognize factive mental states in others? There are several reasons that the answer to this question is ‘no’. Let me mention two. For one thing, the participants might be recognizing that the subjects are in a factive mental state of some sort, and so be inclined to ascribe intentional action, and yet be reluctant to call that mental state ‘knowledge’, since the vignette makes it clear that Andy is not aware of what they are doing, and the participants might be uncertain of whether knowledge requires awareness. So the fact that the

---

<sup>5</sup> As a referee observes, it would be interesting to see whether these results would replicate if [Basketball] was adjusted slightly so that the lights went out *just before* the ball was released, rather than just after. In this circumstance it is much less clear the shot is intentional.

participants are not inclined to ascribe awareness nor knowledge is not sufficient to show that they are not factive mindreading. Here is yet another explanation of what is going on, also compatible with the factive mindreading hypothesis. It might be that, although the basketball player is judged not to know *that they are sinking the basket*, they are judged to know *some other propositions* about their performance. Vekony et al.'s (2021) experiment was not devised to check whether the basketball player has been judged to act intentionally *because* they are judged to know that they are performing some action as means for sinking the basket.

While both alternative explanations are viable, let me develop the latter in particular. Practical knowledge is *knowledge of what one is doing when  $\Phi$ -ing*; but what one is doing need not be *that one is succeeding at  $\Phi$ -ing*. Philosophers thinking about practical knowledge have long recognized that the most plausible formulation of the practical knowledge condition is not one on which intentionally  $\Phi$ -ing requires knowing that *one is  $\Phi$ -ing*. Rather, the practical knowledge requirement is best thought as the requirement that intentionally  $\Phi$ -ing requires knowing *that one is taking certain means in order to  $\Phi$* . This weakening of the practical knowledge condition is required to deal with classical counterexamples, such as Davidson's (1971) carbon copier case, on which the subject is not sure that they will succeed at producing ten carbon copies, and yet seems to do so intentionally (Pavese, 2021; 2022). Though the carbon copier does not know that he is succeeding at producing the ten carbon copies, he *does* know what means he is taking to that effect.<sup>6</sup>

For all we know from Vekony et al.'s (2021) first experiment, it might very well be that subjects judge the basketball player to be sinking the basket intentionally only because they judge him to know that, e.g., he is taking the means available to him in order to sink the basket—i.e., is taking position, positioning his arms and hands in the right way, directing his attention at the basket, and taking a shot. Thus, it is compatible

---

<sup>6</sup> As a referee reminds me, even my preferred weaker formulation according to which intentionally  $\Phi$ -ing requires knowing that one is taking certain means in order to  $\Phi$  might need to be weakened further to deal with Paul's (2009) 'distracted driver' cases. In Pavese (2022), I suggest all one needs for intentionally  $\Phi$ -ing is *that one be in position to know that one is taking the means for  $\Phi$ -ing*.

with Vekony et al. (2021)'s first experiment that the basketball player is represented as being aware of these means for taking the shot, despite the blackout. So, since Vekony et al. (2021)'s first experiment only targets a *particular* construal of the practical knowledge condition, it should not be taken as evidence that factive mindreading does not play a role in the folk psychology of action.

Similar considerations hold for Vekony et al.'s (2021) second experiment, which looks at habitual actions. They used the following vignette:

[Door] Suzy locks her door every morning as she leaves for work. On her way out to work one morning, she locks the door. But, because she is preoccupied with thoughts about her day, she is completely unaware of doing so. She gets to her car, pauses, and wonders if she locked the door. Because she was unaware of locking it, she didn't even remember locking it. So she walks back from her car to check if she locked the door.

And they asked three questions (presented in random order):

When Suzy was locking her door, she knew that she was locking it. (Knowledge)

While Suzy was locking her door, she was aware of locking it. (Awareness)

Suzy intentionally locked her door. (Intentionality)

They found that a majority of participants say that Suzy intentionally locked her door, despite the fact that Suzy failed to satisfy the knowledge and awareness conditions regarding that action. They conclude that, on the ordinary concept of intentional action, neither knowledge nor awareness is required for intentional action.

Like the first, this second experiment only looks at whether knowledge, or awareness, *that one is  $\Phi$ -ing* is judged to be necessary for ascriptions of intentionally  $\Phi$ -ing. So this result is compatible with the hypothesis that subjects would be inclined to ascribe some *other* piece of knowledge to Susy. For example, it is rather plausible that they would be inclined to credit Susy with, e.g., knowledge of where to find her keys.

Thus, the considerations developed above about the first experiment apply equally to this second experiment.

In conclusion, Vekony et al. (2021)'s experimental findings are an important contribution on how to understand the knowledge requirement on the ordinary concept of intentional action. But while they might also be taken to suggest that factive mindreading does not play a role in the folk psychology of action—intentional, skilled, or habitual—this conclusion would be too quick. These findings are overall compatible with factive mindreading entering in the folk psychology of action.

#### 4. Kraemer's puzzle

This section and the next section overview some positive evidence of factive mindreading in the folk psychology of action. This section discusses in particular a puzzle about ascriptions of intentional action and recent experimental work I have conducted, jointly with Paul Henne, suggesting that the puzzle arises because people's ascriptions of intentional action are affected by their ability to recognize and ascribe know-how (Pavese and Henne, 2023). In the closing of the section, I speculate on how our results might also bear on the question of the role of factive mental states in our ability to recognize and ascribe intentional action.

Suppose Brown will win a game if he throws a six with an ordinary, fair dice. Brown rolls a six, so he wins the game. We are more prone to agree with 2 than with 1:

1. Brown intentionally rolled a six.
2. Brown intentionally won the game.

Rolling a six, however, is as likely as winning the game, so this pattern of judgments is puzzling. Specifically, it seems that agents bring about ends (e.g., winning the game) intentionally but also that they do not bring about the means that brought about the ends (e.g., rolling a six) intentionally, even though bringing about the ends and means is just as likely. This contrast in judgments was first raised by Butler (1978) and sharpened by Kraemer (1978). Call it the *Kraemer effect* and the puzzle it raises as *Kraemer's puzzle*.

The Kraemer effect is mysterious: why is it that people judge that agents bring about ends intentionally but they do not bring about the means that brought about those ends intentionally, even though bringing about the ends and means is just as likely? While some philosophers had discussed this puzzle before (e.g., Nadelhoffer, 2004, 2006; Cova et al., 2012), extant accounts have fallen short from a systematic discussion. Pavese and Henne (2023) have explored Kraemer’s effect more systematically, in order to see whether it could be explained by a related concomitant effect concerning judgments about know-how. In other words, we conjectured that what gave rise to Kraemer’s effect was a perceived difference as to the degree at which the agent knew how to perform the means versus the ends.

To get the general motivating idea, consider the game example at the outset. In this scenario, we judge that there *is* a reliable way of performing the ends such that the agent knows: throwing a six is a perfectly reliable way to win the game, and Brown knows that. By contrast, we do not judge that there is a reliable way of performing the means in these cases that the agent knows—in the fair dice game, there is just no reliable way such that Brown knows that way to throw a six. Thus, it is plausible that we judge that Brown knows how to win to a higher degree than he knows how to throw a six. As such, there is an apparent difference in the degree at which the agents know how to perform the ends and the means. We conjecture that this observation could explain the Kraemer effect; people’s intentionality judgments vary between ends and means because of a more fundamental difference in the extent to which they represent the agents as knowing how to perform the ends and the means.

We call this the *Know-How Hypothesis*. The Know-How Hypothesis uniquely predicts that in the scenarios where there is the Kraemer effect there should also be a *corresponding* effect for know-how ascriptions. For instance, it predicts that people should agree more with statement 4 than with statement 3:

3. Brown knows how to throw a six.
4. Brown knows how to win the dice game.

In several experiments, we found evidence for the Know-How Hypothesis (All materials, data, and analysis code for all experiments in this manuscript available at <https://osf.io/bj4np/>.) Let me describe two experiments in particular. In Experiment 1, we asked the participants to read vignettes such as the following:

### **Vignette 1**

Jane is a contestant on a game show. In the game, Jane is given the opportunity to push a button that will randomly open exactly one of the ten doors in front of her.

A brand-new car is behind one of the ten doors. If Jane pushes the button and the door with the brand-new car behind it opens, then she will win the car.

Jane has no idea which door will open if she pushes the button. But she does know that the brand-new car is behind door three. And Jane really wants to win that car. Hoping to win the car, Jane pushes the button. To her great satisfaction, door three opens, and Jane wins the brand-new car.

Before each statement, we asked participants “To what extent do you agree with the following statement about the passage you just read?” Participants who received the game show vignette responded to the following:

Jane knows how to win the brand-new car.

Jane knows how to open door three.

Jane intentionally won the brand-new car.

Jane intentionally opened door three.

For each question, we asked participants for their level of agreement with each statement on a -50-50 scale [-50 = strongly disagree, 0 = neutral, 50 = strongly agree].

We found that just as people are more inclined to judge that agents brought about the ends intentionally than the means, people also are more inclined to judge that agents know how to bring about the ends more than the means. That is, we conceptually replicated the Kraemer effect in a variety of non-moral scenarios (Nadelhoffer, 2004; 2005), and we found a new effect predicted by our new hypothesis: the Know-How



effect. We also found that people's intentionality judgments that give rise to the Kraemer effect are fully mediated by people's know-how judgments.<sup>7</sup>

Encouraged by these results, we planned more experiments to further test the Know-How Hypothesis. Here is a particularly significant one. The Know-How Hypothesis holds that people's know-how judgments explain the Kraemer effect. One might think that perhaps there is another explanation of the effect that is just as plausible. Perhaps, people judge that there is a reliable way for the agent to bring about the ends but no reliable way to bring about the means and that this difference in judgments independently explains both the know-how effect and the Kraemer effect. This alternative hypothesis predicts that, if there is a reliable way to perform both the ends and the means, then there should be neither a know-how effect nor a Kraemer effect. The Know-How Hypothesis instead predicts that these effects should persist independently of the presence of a reliable way to perform the ends and the means when there is a difference in know-how.

To test these predictions, we used a modified version of the game show vignette from Experiment 1. In the modified vignette, there is a reliable way to perform the end and there is a reliable way to perform the means: the agent can win the brand-new car by opening door three, and they can open door three by pushing the button, which is fixed. However, the agent has no idea how to open door three, since they do not know that the button is fixed. In this setting, the Know-How Hypothesis predicts both the know-how effect and the Kraemer effect. By contrast, the Common Cause Hypothesis predicts that both the Kraemer effect and the know-how effect would disappear when there is no difference in reliable ways to execute means and ends.

Again, we found a know-how effect and a Kraemer effect, replicating our findings from Experiment 1. The Know-How Hypothesis predicted these results because, in this view, it is not the presence or absence of reliable ways to perform the means and the end

---

<sup>7</sup> Participants agreed to a greater degree that the agent knew how to perform the ends ( $M = 12.91$ ,  $SD = 34.88$ ,  $n = 843$ ) than that the agent knew how to perform the means ( $M = -28.38$ ,  $SD = 29.39$ ,  $n = 843$ ) ( $b = -41.29$ ,  $SE = 1.55$ ,  $t = -26.49$ ,  $p < .001$ ,  $CI_b [-44.35, -38.24]$ ,  $d = -1.29$ ). Participants also agreed to a greater degree that the agent intentionally performed the ends ( $M = -10.97$ ,  $SD = 35.73$ ,  $n = 843$ ) than that the agent intentionally performed the means ( $M = -28.01$ ,  $SD = 30.31$ ,  $n = 843$ ) ( $b = -17.04$ ,  $SE = 1.57$ ,  $t = -10.83$ ,  $p < .001$ ,  $CI_b [-20.13, -13.96]$ ,  $d = -.52$ )

that is responsible for the know-how effect or the Kraemer effect. Rather, according to this view, it is the difference in know-how for the end and the means (i.e., a know-how effect) that is responsible for a Kraemer effect. Critically, these results are incompatible with the Alternative Hypothesis, which predicts no know-how effect and no Kraemer effect when there is a reliable way for the agent to perform both the ends and the means. Pavese and Henne (2023) go on to replicate these results also for morally loaded versions of Kraemer’s puzzle (Experiment 5 and Experiment 6).

The relevance of these findings for the topic of this chapter should be clear, given the discussion in §2, but it is nonetheless worth elaborating. To begin with, if the ordinary concept of know-how requires that of propositional knowledge, then these findings suggest that knowledge representation plays a role in folk ascribing intentional action, for they suggest that know-how representation is central to the ordinary practice of ascribing intentional action. Second, though Pavese and Henne (2023: 10-11) only briefly discuss this further complexity, our results are actually hard to explain on any view on which our concept of know-how is independent of propositional knowledge. Consider, for example, a view on which know-how is conceived as a mere ability of sort—understood as a reliable disposition to success, not further grounded in any factive mental state. In Vignette 1, subjects are presented as having the same ability to perform both means and ends, since means and ends are presented as just as likely. So, this view does not seem equipped to explain the know-how effect that we have found, since in this case, the subjects are as disposed to succeed at performing the ends as they are at performing the means. Moreover, notice that in Experiment 3, manipulating the know-how *effectively* amounted to manipulating the subject’s factive mental state—here Jane is presented as having no idea of how to open door three, so lacking awareness of how to open door three.

This observation goes in the direction of supporting the thought that the ordinary concept of know-how is at least partly intellectualistic—something that others have argued in recent literature as well (Gonnerman et al., 2018; 2021). So, these findings are

not only compatible with the hypothesis that factive mental states play a role in explaining Kraemer's puzzle—to a closer look, they are *suggestive* of such a role.

### 5. Epistemic luck and the ordinary concepts of Know-How and Intentional Action.

Though our findings on Kraemer's puzzle are suggestive of a role for factive mental states in our capacity to track intentional action, they cannot be regarded as conclusively showing that *knowledge* specifically plays such a role. After all, in these experiments, we have not controlled for the possible difference between knowledge and other factive mental states, such as awareness. This section discusses experimental findings on the ordinary concept of know-how and of intentional action that provide evidence specifically for the role for knowledge representation in the folk psychology of action.

One respect on which propositional knowledge might be thought to differ from other factive mental states—such as e.g., awareness—is that the former but not the latter would be defeated in Gettier scenarios, and more generally in cases involving 'veritic epistemic luck' (Pritchard, 2005). Indeed, some recent experimental work shows that the folk concept of knowledge is sensitive to Gettierization (Nagel et al., 2013; Turri, 2013; Machery et al., 2017). These findings open up a novel avenue for the research on factive mindreading. If the ordinary concept of knowledge is sensitive to Gettierization and to veritic epistemic luck, one way to ascertain whether knowledge representation plays a role in the folk psychology of action is to test whether epistemic luck also undermines the ordinary concept of intentional action.

This question can be investigated directly by looking at whether epistemic luck undermines people's inclination to ascribe intentional action. It can also be investigated *indirectly* by looking at whether epistemic luck undermines people's inclination to ascribe some *related* mental state concept that is plausibly related to that of intentional action. As we have seen in §2, the concept of know-how seems relevant to the concept of intentional action. Many philosophers and cognitive scientists have argued independently that the concept of intentional action is fundamentally related to the concept of *skills* (e.g., Ryle, 1949; Mele & Moser, 1994; Malle & Knobe, 1997; Malle, 2003; Guglielmo et al, 2009;

Guglielmo & Malle, 2010; Pavese 2013; Setiya, 2012; Cath, 2015; Pavese 2016; Pavese & Beddor, 2022).<sup>8</sup> Moreover, it is commonly assumed that a person is skilled at an action only if one knows how to perform it: skill requires know-how (e.g., Ryle, 1949; Setiya, 2012; Cath, 2015; Pavese, 2016). On this account, intentional action is fundamentally related to know-how and skills.

Thus we might ask whether epistemic luck undermines people's inclination to ascribe know-how. If we found evidence for this claim, we would have some *indirect* evidence that epistemic luck undermines people's inclination to ascribe intentional action too. Let us begin to explore this indirect path.

In epistemology, it is often claimed that the folk concept of know-how is compatible with epistemic luck (e.g., Poston, 2009; Cath, 2011; Carter & Pritchard, 2015). One source of evidence for this claim are cases such as this, adapted from Cath (2011; 2015):

**Lucky Lightbulb.** Charlie wants to change a lightbulb. Being unversed in such matters, he pulls down a manual of everyday household tasks, looks up the instructions for lightbulb-changing, and proceeds to follow them. It turns out that the author of the manual was a prankster, who riddled the book with inaccurate instructions. But by a fluke, when the instruction manual went to the printers, a correct set of lightbulb-changing instructions were substituted at the last minute, due to a misprint.

There is reason to question whether this case, and other examples that have been proposed, are genuine Gettier cases (cf Pavese & Beddor 2022). Indeed, in the following, I am going to suggest that this example, and other examples that have been produced in the literature, do not even have the same *structure* as Gettier's original cases.

---

<sup>8</sup> In morally loaded cases, this connection between intentionality and skill or know-how is somewhat weakened (Knobe, 2003b; Nadelhoffer, 2005). In our studies, however, we focus on *morally neutral cases* for which the connection between intentional action and skills and know-how has been widely confirmed empirically (cf. Malle & Knobe, 1997; Nadelhoffer, 2005). For more discussion on the relation between intentionality and skills, see Guglielmo & Malle, 2010).

In prototypical Gettier cases, the agent's belief is *unsafe*—that is, there is a nearby circumstance where the agent forms *the very same belief* on similar grounds, but their belief is false. For example, consider the following variant of one of Gettier's (1963) examples:

**Occupational Hazard.** Suppose that Smith and Jones have applied for a certain job. And suppose that Smith has strong evidence for the following conjunctive proposition:

(d) Jones is the man who will get the job, and Jones has ten coins in his pocket.

Smith's evidence for (d) might be that the president of the company assured him that Jones would in the end be selected, and that he, Smith, had counted the coins in Jones's pocket ten minutes ago. Proposition (d) entails:

(e) The man who will get the job has ten coins in his pocket.

Let us suppose that Smith sees the entailment from (d) to (e), and accepts (e) on the grounds of (d), for which he has strong evidence. In this case, Smith is clearly justified in believing that (e) is true. But imagine, further, that unknown to Smith, he himself, not Jones, will get the job. And, also, unknown to Smith, he himself has ten coins in his pocket. Proposition (e) is then true, though proposition (d), from which Smith inferred (e), is false. It is equally clear that Smith does not know (e).

In this example, there is a nearby world where either Smith did not get the job, or Smith did not have ten coins in his pocket. In that world, Smith would have believed the *same* proposition (e), but his belief would have been false. So Smith's belief is unsafe. Indeed, the lack of safety appears to be the general mark of veritic epistemic luck (Sosa, 1999; Williamson, 2000; Pritchard, 2005; Beddor & Pavese, 2020).

By contrast, in **Lucky Lightbulb**, it is less clear that Charlie's belief is unsafe. In the nearby world where the instruction manual is free from misprints, Charlie would have formed a *very different belief* about how to change a lightbulb, since he would have come

to believe an altogether different set of instructions. So there is no nearby world where he holds the very same belief falsely. Thus **Lucky Lightbulb** does not even have the same structure as typical Gettier cases, since the main protagonist appears to have a safe belief.

Some might protest that safety principles should not hold fixed the content of the relevant belief, on pain of making all beliefs in necessary propositions trivially safe (Pritchard, 2012). Suppose I believe that that is a barn, where ‘that’ picks up a barn within sight. The content of my belief is necessarily true. So I could not have had the same belief falsely. And yet intuitively, I might fall short of knowing that that is a barn. So a fix is needed to the standard formulation of a safety condition on knowledge and any such fix might extend to **Lucky Lightbulb**. On this ground, one might object that on the right formulation of the safety principle, Charlie’s belief in **Lucky Lightbulb** will come out unsafe after all.

However, this observation is simply incorrect. To see this, consider a prominent strategy that has been proposed to deal with the problem of trivializing safety for necessary propositions. Pritchard (2012) proposes that a perceptually based belief *B* qualifies as safe if and only if any belief *B'* which is formed in response to the *same perceptual stimuli* would have been true. This formulation of safety deals with many cases of necessary but empirically formed beliefs and it provides a desirable restriction on the relevant counterpart beliefs. It seems relevant to cases such as **Lucky Lightbulb** too, where the belief is formed on the basis of perception (as well as on the basis of reading skills). Notice, however, that if we apply this corrected definition of safety to **Lucky Lightbulb**, we still get that Charlie’s belief is not unsafe. After all, if Charlie had received the manual with the misprint, he would have been in a different perceptual state, since he would have been reading a very different set of instructions. So, even by the lights of a formulation of safety for perceptual belief that does not run afoul of the problem of necessary propositions, Charlie’s belief turns out to be safe.

Many putative examples of Gettiered know-how that the literature on know-how has produced are, just like **Lucky Lightbulb** and for the same reasons, actually cases of

safe beliefs. Consider, for example, this example put forward by Stanley & Williamson (2001:435) as an example of Gettier case for know-how:

[T]here are indeed Gettier cases for know-how. Bob wants to learn how to fly in a flight simulator. He is instructed by Henry. Unknown to Bob, Henry is a malicious imposter who has inserted a randomizing device in the simulator's controls and intends to give all kinds of incorrect advice. Fortunately, by sheer chance the randomizing device causes exactly the same results in the simulator as would have occurred without it, and by incompetence Henry gives exactly the same advice as a proper instructor would have done. Bob passes the course with flying colors. He has still not flown a real plane. Bob has a justified true belief about how to fly.

Just like in **Lucky Lightbulb**, Bob's justified true belief about how to fly *is* safe—had the randomizing device returned Henry a different advice to give Bob, Bob would indeed have a false belief about how to fly but it would not have been the same belief as the one that Bob turned out to have. That belief that Bob does have would have been true in that counterfactual circumstance too.

So several widely discussed counterexamples to the hypothesis that know-how can be Gettiered do not resemble closely enough paradigmatic Gettier cases, in that they are not cases of unsafe beliefs. The same holds for Carter & Pritchard (2015)'s variant of **Lucky Lightbulb**—originally designed to more closely resemble the fake barns case. On their variant, Charlie reads a manual that contains no misprints, but he could have easily received a phony manual with incorrect instructions. Here too, the earlier diagnosis applies: had he received the phony manual, he would have held a false belief but the content of the belief would have been importantly different.

Indeed, in this particular case, this diagnosis has received (unintended) empirical confirmation. Carter et al. (2019) have conducted an experimental study on the Know-how Question, with the goal of assessing Intellectualism about Know-How—the view that knowing how to perform some task  $\phi$  requires or consists in propositional

knowledge about how to  $\varphi$  (Stanley and Williamson, 2001; Pavese, 2017; 2018). In particular, Carter and colleagues used 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 et al. (2019) found that people tend to agree with the claim that Charlie knows how to change a lightbulb and take these results to be evidence that people attribute know-how in cases of epistemic luck.<sup>9</sup> However, there is reason to question Carter and colleagues’s diagnosis. They meant their vignette as a case of *environmental luck*—the sort of luck present in fake barn cases (Pritchard, 2005). However, 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.<sup>10</sup> 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

---

<sup>9</sup>This line of argument has been pushed by Cath (2015).

<sup>10</sup>Nagel et al. (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, & 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.



way of changing a lightbulb—suggesting that people are willing to attribute *both* know-how and knowledge-that in this scenario.<sup>11</sup> So, all in all, it is far from clear that **Lucky Manual** is an authentic case of epistemic luck.

So, in conclusion: Cath’s original counterexample to the claim that know-how is not Gettierizable (**Lucky Lightbulb**) is questionable, since it does not have the same structure as Gettier’s original cases—indeed, it does not even have the same modal profile as Gettier cases. The same is true for other putative counterexamples as well as for Carter and Pritchard’s variant. Indeed, Carter et al.’s own experimental findings encourage this diagnosis.

A second recent experimental study on this issue is worth mentioning. Gonnerman et al. (2021:11333-5) found that participants were more inclined to attribute know-how when protagonists were described as having a justified true belief than when they were described as having an *unjustified but luckily true* belief. They use this vignette:

MEDICAL DIAGNOSIS. Amber, Brenda, Carrie, and Diana are four undergraduate students. They have recently applied to a selective medical program. Multiple factors go into any application decision for this program. One of these is whether the applicant knows how to diagnose an autoimmune disorder called Schnitzler syndrome. The way to diagnose Schnitzler syndrome is to look for signs of urticarial eruptions, recurrent fevers, and relapsing arthralgia. We’ll call this way of diagnosing the syndrome “w”. Coming into their admissions interview, Amber, Brenda, Carrie, and Diana had each heard of Schnitzler syndrome. But none of them had ever tried diagnosing the syndrome before. Still they were excited to give it a try, even when they were informed that they would have to diagnose twenty people in a row who may or may not have the syndrome while getting no feedback as to whether they are getting the diagnosis correct. What follows are the admission officer’s notes. Assuming that these notes are

---

<sup>11</sup> We also note a further puzzling aspect of Carter et al.’s (2019) result: they found that participants were more inclined to attribute know-how in the epistemic luck conditions than in the control conditions. Though very puzzling, as they themselves note, this result might have a non-epistemic explanation (cf. pp. 717-8). Accordingly, we set this finding aside for the purpose of our discussion.

accurate, do you agree or disagree that the applicant knows how to diagnose Schnitzler syndrome?

This vignette was followed by descriptions such as:

[Applicant] – Believes that *w* is the way to diagnose Schnitzler syndrome. Has very good [poor] reason for believing this, having learned this fact from someone very [in]experienced in autoimmune disorders. Correctly diagnosed 9 [1] out of 10 people with Schnitzler syndrome.

They found (p. 11335) that participants were more inclined to attribute know-how when protagonists were described in ways suggesting that they had a justified true belief about the way to execute the task ( $M=4.20$ ,  $SD=1.54$ ) than when they were described to have an unjustified true belief ( $M=3.45$ ,  $SD=1.64$ ). They note that this is evidence that *veritic epistemic luck* affects folk ascriptions of know-how, so the folk concept of know-how is sensitive to the presence of propositional knowledge.

While this diagnosis is plausible, it is worth noting that their study does not involve a genuine Gettier case either—not at least in the standard understanding of Gettier cases as involving *justified* true belief that falls short of knowledge—nor the structure of the case involved matched those of Gettier’s original cases. After all, in the lucky condition the prompt did present the applicants as “having poor reason for believing the diagnosis”, so as having an unjustified true belief. So, Gonnerman et al. (2021) recent experimental work does not address the question whether the ordinary concept of know-how is sensitive to Gettierization.

It is now time to look at some positive evidence that the ordinary concept of know-how is sensitive to epistemic luck. Pavese et al. (forthcoming) have conducted experimental work on the role of epistemic luck on know-how that differs from these recent works in that we focused on a variety of uncontroversial cases of epistemic luck which are also exact variants of Gettier’s original cases, as well as in Dharmottara’s desert traveler case, Russell’s stopped clock case, in addition to lottery cases. In all of

these cases, one of Pavese et al.'s (forthcoming) goals was to test whether this sort of epistemic luck undermines people's tendency to ascribe know-how. All of the vignettes we used feature someone acting 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, they asked participants the extent to which they agreed that the agent knows-that and knows-how.

For example, in the first study, they used two vignettes for condition (lucky/control). In the lucky condition, the vignette was:

Lucky:

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.

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.

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.

Alvin's belief that sending an email to the email address barns678@yahoo.com is a means to send it to the person who got the promotion is Gettiered—it is true and justified, but falls short of knowledge. Indeed, this vignette structurally matches Gettier's original coins in the pocket case. Interestingly, Pavese et al. (2023) found that 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 ( $M = 3.78$ ,  $SD = 35.36$ ,  $n = 241$ ) ( $b = 36.14$ ,  $SE = 2.47$ ,  $t = 14.58$ ,  $p < .001$ ,  $CI_b [31.28, 41.01]$ ,  $d = 1.33$ ).

As another example, in a third study, we tested the following vignette, which matches structurally Dharmottara's (770 CE: D:4429, 9a2-3) desert mirage case. 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. In the lucky condition, the vignette was:

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.

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.

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.

Here, they replicated the results of the first experiment: 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 ( $M = -19.43, SD = 25.79, n = 83$ ) ( $t(156.58) = 9.34, p < .001, d = 1.49, CI [1.14, 1.85]$ ).

In each case, Pavese et al. (forthcoming) found that epistemic luck undermines people's tendency to ascribe know-how, suggesting that the folk concept of know-how is Gettierizable. These results fit nicely with the idea, voiced by others too (cf. Bengson et al. 2009; Gonnerman et al. 2018) that the folk concept of know-how is at least partly 'intellectualistic'—it is at least related to the concept of propositional knowledge, since this position *predicts* an effect of epistemic luck on ordinary ascriptions of know-how. These findings speak against those who have argued that the folk are purely anti-intellectualists with respect to know-how (cf. Harmon and Horne, 2016). It also is strongly suggesting that our ability to recognize and ascribe know-how requires knowledge-representation. Why else would a distinctively epistemic kind of luck undermine people's ability to ascribe know-how?

Now suppose the ordinary concept of know-how is sensitive to epistemic luck. Since the ordinary concept of intentional action is closely related to that of know-how, we should expect the ordinary concept of intentional action to be sensitive to know-how too. However, this conclusion has been often rejected without much argument. For example, Cath (2015) argues that the ordinary concept of know-how cannot require that of knowledge, since the ordinary concept of know-how goes together with the ordinary concept of intentional action and the latter is independent of the concept of knowledge.

Is it true, though, that the presence of epistemic luck is compatible with people's willingness to ascribe intentional action? In order to address this issue, in the same studies, Pavese et al. (forthcoming) simultaneously checked for whether epistemic luck affected both people's willingness to agree with a know-how ascription and their willingness to agree with an intentional action ascription. For example, in Study 1 above, they asked whether Alvin intentionally sent the email to the person who got the promotion. Here is what we found: Participants who read about the agent in the control

condition agreed to a greater extent that the agent intentionally performed the action ( $M = 43.06$ ,  $SD = 15.14$ ,  $n = 237$ ) than participants in the lucky condition ( $M = -9.82$ ,  $SD = 37.72$ ,  $n = 241$ ) ( $b = 52.77$ ,  $SE = 2.60$ ,  $t = 20.24$ ,  $p < .001$ ,  $CI_b [47.66, 57.89]$ ,  $d = 1.85$ ). Similarly, for Study 2, they asked whether Maria intentionally brought the group to water. Here is what they found: 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 ( $M = -10.78$ ,  $SD = 35.46$ ,  $n = 83$ ) ( $t(148.96) = 8.02$ ,  $p < .001$ ,  $d = 1.31$ ,  $CI [0.96, 1.67]$ ). This result strongly suggests that the epistemic luck defeated people's inclination to ascribe intentional action.

Pavese et al. (forthcoming) found very similar results in a lottery case and in a variant of Russell's clock case. Thus, far from epistemic luck sparing people's inclination to ascribe intentional action, the folk concept of intentional action, just like that of propositional knowledge, *is* sensitive to the presence of epistemic luck. This argument is reinforced by adding know-how to the mix: since the ordinary concept of know-how is sensitive to epistemic luck in those scenarios too, if intentional action requires know-how, we should expect that the ordinary concept of intentional action is sensitive to epistemic luck too. Though more investigation is needed on this question, these findings are already strongly suggestive of a role for knowledge representation in the folk psychology of action.

## 6. Open ends.

This chapter has concerned the role of factive mindreading in the folk psychology of action. Granting that factive mindreading plays a role in our ability to predict and explain behavior, how does it enter (if at all) in our ability to track intentional and skilled action? I sharpened this question and motivated it by looking at several views which assign factive mental states a central role to play in action theory. After discussing and rebutting some *prima facie* reasons to be skeptical, I have overviewed some recent studies that suggest a role for factive mindreading in the folk psychology of intentional action. While these results are preliminary, they are strongly suggestive of factive mindreading in the

folk psychology of action. Further experimental investigation might look more closely at people's judgments about skilled action, and about the relation between skilled action and know-how; at the relation between knowledge representation and our ordinary concepts of culpability, responsibility, and intent; at the relation between knowledge representation and the ordinary concept of creditability, as well as at that between knowledge representation and legal and moral concepts.

## References

- Anscombe, Gertrude Elizabeth Margaret (1958). *Intention* (London: Basil Blackwell)
- Beddor, B. & Pavese C. (2022). "Practical Knowledge without Luminosity", in *Mind*, 131: 523, 917–934.
- Bengson, J., Moffett, M. A., & Wright, J. C. (2009). "The Folk on Knowing How." *Philosophical Studies*, 142, 387-401.
- Butler, R. J. (1978). "Report on Analysis "Problem"" No. 16. *Analysis*, 38(3), 113-114.
- Cath Y. (2011). "Knowing how without knowing that," In Bengson and Moffett (eds.), *Knowing How: Essays on Knowledge, Mind and Action*. Oxford University Press: 113–135.
- Cath Y. (2015). "Revisionary Intellectualism and Gettier." *Philosophical Studies* 172(1): 7–27.
- Cath, Y. (2020). "Know how and Skill: The Puzzles of Priority and Equivalence." In *The Routledge Handbook of Philosophy of Skill And Expertise* (pp. 157–167). Routledge.
- Colaco, D., Buckwalter, W., Stich, S., & Machery, E. (2014). "Epistemic intuitions in fake-barn thought experiments." *Episteme* 11(2): 199–212.
- Cova, F., Dupoux, E., & Jacob, P. (2012). "On doing things intentionally." *Mind & Language*, 27(4), 378–409.
- Dharmottara (770 CE). *Explanation of (Dharmakirti's) Ascertainment of Valid Cognition*.

- Davidson, D. (1971). "Agency", in R. Binkley, R. Branaugh, and A. Marras (eds.), *Agent, Action, and Reason*, Toronto: University of Toronto Press.
- Dreyfus, H. L. (2006). "Overcoming the myth of the mental." *Topoi*, 25, 43-49
- Guglielmo, S., Monroe, A. E., & Malle, B. F. (2009a). "At the Heart of Morality Lies Folk psychology." *Inquiry*, 52(5), 449–466.
- Guglielmo, S., Monroe, A. E., & Malle, B. F. (2010). "Enough Skill to Kill: Intentionality Judgments and the Moral Valence of Action." *Cognition*, 117(2), 139–150.
- Kneer, M. (2021). "Success and knowledge in action: Saving Anscombe's account of intentionality." *Context dependence in language, action, and cognition*, 131–154.
- Kraemer, E. R. (1978). "Intentional Action, Chance, and Control." *Analysis*, 38(3), 116-117.
- Knobe, J. (2003a). "Intentional Action and Side Effects in Ordinary Language." *Analysis*, 63(3), 190–194.
- Knobe, J. (2003b). "Intentional Action in Folk Psychology: An Experimental Investigation." *Philosophical Psychology*, 16(2), 309–325.
- Knobe, J. (in press). "Difference and Robustness in the Patterns of Philosophical Intuition Across Demographic Groups." *Review of Philosophy and Psychology*.
- Knobe, J. and Burra, A. (2006). "Intention and Intentional Action: A Cross-Cultural Study." *Journal of Culture and Cognition*, 6, 113-132.
- Kraemer, E. R. (1978). "Intentional action, chance and control." *Analysis*, 38(3), 116–117.
- Gettier, E. (1963). "Is Justified True Belief Knowledge?" *Analysis* 23 (6): 121–123.
- Gibbons, J. (2001). "Knowledge in Action," *Philosophy and Phenomenological Research* 62(3): 579–601.



Gonnerman, C., Mortensen, K., & Robbins, J. (2018). “The ordinary concept of knowledge how.” *Oxford Studies in Experimental Philosophy*, 2, 104–115.

Gonnerman, C., Mortensen, K., & Robbins, J. (2021). “*Knowing how* as a philosophical hybrid”, *Synthese*, 199(3), 11323–11354.

Hampshire, S. (1959). *Thought and Action* (University of Notre Dame Press).

Harmon, I., & Horne, Z. (2016). “Evidence for anti-intellectualism about know-how from a sentence recognition task.” *Synthese*, 193(9), 2929–2947.

Heider, F. (1958). *The Psychology of Interpersonal Relations*. New York: Wiley.

Horschler, D., Santos, L., and MacLean, E. (2019): “Do non-human primates really represent others’ ignorance? A test of the awareness relations hypothesis,” *Cognition* 190: 72–80.

Malle, B. F., & Knobe, J. (1997). “The Folk Concept of Intentionality.” *Journal of experimental social psychology*, 33(2), 101–121.

Malle, B. (2003). “The Social Cognition of Intentional Action.” *Malingering and illness deception*, 83–92.

Machery, E. et al., (2017). “Gettier Across Cultures.” *Nous* 51(3): 645–664.

Mele, A. & Moser, Paul K. (1994). “Intentional Action.” *Nous*, 28(1), 39–68.

Nadelhoffer, T. (2004). “The Butler Problem Revisited.” *Analysis*, 64(3), 277–284.

Nadelhoffer, T. (2005). “Skill, Luck, Control, and Intentional Action.” *Philosophical Psychology*, 18(3), 341–352.

Nagel, J. (2013). “Knowledge as a Mental State.” *Oxford Studies in Epistemology* 4: 273.

Nagel, J. (2017). “Factive and non-Factive Mental State Attribution.” *Mind & Language*, 32(5), 525–544.

- Nagel, J., San Juan, V., and Mar, R.A. (2013). “Lay Denial of Knowledge for Justified True Beliefs.” *Cognition* 129(3): 652–661.
- Paul, S. K. (2009). “How We Know What We’re Doing.” *Philosopher's Imprint*, 9(11).
- Pavese, C. (2020). “Probabilistic Knowledge in Action.” *Analysis*, 80(2), 342–356.
- Pavese, C. (2016). “Skill in Epistemology, I and II” *Philosophy Compass*, 641–660.
- Pavese, C. (2018). “Know-how, Action, and Luck.” *Synthese*, 198(7), 1595–1617.
- Pavese, C. (2021a). “Knowledge, Action, and Defeasibility.” In Brown, J., & Simion, M. *Reasons, Justification, and Defeaters*, Oxford University Press, pp. 177-200.
- Pavese, C. (2021b). “Knowledge and Mentality.” *Philosophical Perspectives*, 35(1), 359–382.
- Pavese, C. (2022). “Practical Knowledge First”, *Synthese*.
- Pavese, C. (forthcoming). *The Practical Mind*, Cambridge University Press.
- Pavese, C. & Beddor, B. (2022). “Skills as Knowledge.” *Australasian Journal of Philosophy*, 1-16.
- Pavese, C. & P. Henne (2023). “The Know-How Solution to Kraemer’s puzzle”. <https://osf.io/bj4np>, *Cognition*, 238.
- Pavese, C., P. Henne, & B. Beddor (forthcoming). “Epistemic Luck, Knowledge-How, and Intentional Action” *OSF*. <https://osf.io/s4hyd>, *Ergo*.
- Perner, J. (1993). *Understanding the Representational Mind* (Cambridge, Mass.: MIT Press).
- Phillips, J., & Norby, A. (2021). “Factive theory of mind.” *Mind & Language*, 36(1), 3–26.

- Phillips, J., Buckwalter, W., Cushman, F., Friedman, O., Martin, A., Turri, J., & Knobe, J. (2021). "Knowledge before belief." *Behavioral and Brain Sciences*, 44.
- Poston, T. (2009). "Know how to be Gettiered?" *Philosophy and Phenomenological Research*, 79(3), 743-747.
- Povinelli, D. J., Perilloux, H. K., Reaux, J. E., & Bierschwale, D. T. (1998). "Young and juvenile chimpanzees' (Pan troglodytes) reactions to intentional versus accidental and inadvertent actions." *Behavioural Processes*, 42(2-3), 205-218.
- Pritchard, D. (2005). *Epistemic Luck*. Oxford University Press.
- Pritchard, D. (2012). Anti-Luck Virtue Epistemology. *The Journal of Philosophy*, 109(3), 247–279.
- Ryle, G. (1949). *The Concept of Mind*. Hutchinson, London.
- Shepherd, J. (2021). *The Shape of Agency: Control, Action, Skill, Knowledge* (p. 208). Oxford University Press.
- Schwenkler, J. (2019). *Anscombe's Intention: A Guide*. Oxford University Press.
- Sosa, E. (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge, volume I* (Vol. 1). OUP Oxford.
- Stanley, J., & Williamson, T. (2001). "Knowing how." *The Journal of Philosophy*, 98(8), 411-444.
- Tomasello, M. (2000). "Primate cognition: introduction to the issue." *Cognitive science*, 24(3), 351-361.
- Turri, J. (2013). "A conspicuous art: putting Gettier to the test." *Philosophers' Imprint* 13(10): 1–16.
- Turri, J. (2017). "Knowledge attributions in iterated fake barn cases." *Analysis* 77(1): 104–115.

Vekony, R., Mele, A., & Rose, D. (2020). “Intentional action without knowledge.” *Synthese* 1–13.

Westra, E., & Nagel, J. (2021). “Mindreading in conversation.” *Cognition*, 210.

Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press, Oxford.

Wood, J. N., Glynn, D. D., Phillips, B. C., & Hauser, M. D. (2007). “The perception of rational, goal-directed action in nonhuman primates.” *Science*, 317(5843), 1402-1405.

Wright, J. (2010). “On Intuitional Stability: The Clear, the Strong, and the Paradigmatic.” *Cognition* 115(3): 491–503.