

EASY PRACTICAL KNOWLEDGE*

In central or well-ordered cases of intentional action, an agent *knows what she's doing as she does it*; her intentional actions manifest “practical knowledge.” Some authors, notably Carlotta Pavese and Bob Beddor,¹ have recently defended the broadly Anscombian thesis that intentional actions are *necessarily* accompanied by practical knowledge;² manifesting such knowledge is an essential part of *actional control*, the control that distinguishes intentional actions from unintentional ones and mere (non-actional) behavior.

In this article, we draw new connections between the epistemologies of mental rehearsal and suppositional reasoning to offer a novel perspective on skilled behavior and its relationship to practical knowledge. In particular, we argue that practical knowledge is “easy” in the sense that, by manifesting one’s skills, one has *a priori* propositional justification for certain beliefs about what one is doing as one does it. One important point of contrast between our “skill-first” theory of actional control and other “practical knowledge-first” theories is that, because agents sometimes act intentionally in epistemically hazardous environments, (skill-)justified beliefs do not necessarily rise to the level of practical knowledge. This view charts a middle way through the Anscombian tradition of defending a necessary connection between intentional action and practical knowledge, on the one hand, and the more recent tradition of explaining away any substantive epistemic condition on intentional action, on the other.³

Our criticism, as our appeal to “easy knowledge” might suggest, relies on motivating parallels between the knowledge requirement for

* Carter’s research is supported by the Arts and Humanities Research Council’s *Expanding Autonomy* (AH/W005077/1) and *Digital Knowledge* (AH/W008424/1) projects as well as by the European Research Council grant number 101199255 (KNOW-HOW). Kearl’s research is supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (948356, *Knowledge-Lab*, PI: Mona Simion).

¹ Carlotta Pavese and Bob Beddor, “Skills as Knowledge,” *Australasian Journal of Philosophy*, CI (2023): 609–24; and Carlotta Pavese, “Practical Knowledge First,” *Synthese*, CC, 5 (2022): 1–18.

² G. E. M. Anscombe, *Intention* (Blackwell, 1963).

³ Sarah Paul, “How We Know What We’re Doing,” *Philosophers’ Imprint*, IX (2009): 1–24; and Juan Piñeros-Glasscock, “Practical Knowledge and Luminosity,” *Mind*, CXXIX, 516 (2019): 1237–67.

intentional action and the so-called “easy knowledge” problem.⁴ If we are correct, debates surrounding the connection between intentional action, knowledge, and actional control are more intimately related to traditional epistemological questions about basic knowledge than has so far been appreciated. It is our hope that by exploring these novel questions about the epistemology of skill, we might enrich our collective grasp of the relationship between knowledge and action.

I. THE EPISTEMOLOGY OF MENTAL REHEARSAL

It is common for elite athletes to mentally rehearse⁵ their competitions, visualizing, sensing, or otherwise simulating their performance before they compete. Sometimes this is a process of addressing contingencies; Michael Phelps, one of the greatest competitive swimmers of all time, is reported to have imagined himself swimming with broken goggles and torn trunks, internalizing and automating how he would respond to those conditions in the pool. Other times, it is a process of reducing performance anxiety by anticipating the accomplishment of a desired end. A gymnast might mentally rehearse her bars event, focusing on how to stick a tricky dismount; a kicker might spend hours visualizing his field goals before the playoffs: where to place the tee, how many steps back to take when lining up, how quickly to call for the ball to snap, and the precise force and angle with which to hit it, depending on field position and wind conditions, and so on. When all goes well, the kicker’s mental rehearsal can provide them with knowl-

⁴ See, for instance, Richard Fumerton, *Metaepistemology and Skepticism* (Rowman and Littlefield, 1995); Jonathan Vogel, “Reliabilism Leveled,” this JOURNAL, xcvi, 11 (2000): 602–23; Stewart Cohen, “Basic Knowledge and the Problem of Easy Knowledge,” *Philosophy and Phenomenological Research*, lxxv, 2 (2002): 309–29; Stewart Cohen, “Why Basic Knowledge Is Easy Knowledge,” *Philosophy and Phenomenological Research*, lxxv, 2 (2005): 417–30; Peter Markie, “Easy Knowledge,” *Philosophy and Phenomenological Research*, lxxv, 2 (2005): 406–16; Michael Titelbaum, “Tell Me You Love Me: Bootstrapping, Externalism, and No-Lose Epistemology,” *Philosophical Studies*, cxlix, 1 (2010): 119–34; Jonathan Weisberg, “Bootstrapping in General,” *Philosophy and Phenomenological Research*, lxxxi, 3 (2010): 525–48; and Jonathan Weisberg, “The Bootstrapping Problem,” *Philosophy Compass*, vii, 9 (2012): 597–610.

⁵ Mental Rehearsal, as we use the term, is a form of Mental Imagery Practice. For discussion, see Martha Epstein, “The Relationship of Mental Imagery and Mental Rehearsal to Performance of a Motor Task,” *Journal of Sport Psychology*, ii, 3 (1980): 211–20; Jean Decety et al., “Vegetative Response During Imagined Movement Is Proportional to Mental Effort,” *Behavioral Brain Research*, xlii (1991): 1–5; Jennifer Cumming and Craig Hall, “Deliberate Imagery Practice: The Development of Imagery Skills in Competitive Athletes,” *Journal of Sports Science*, xx, 2 (2002): 137–45; Jörn Munzert, Britta Lorey, and Karen Zentgraf, “The Role of Motor Imagery in the Study of Motor Representations,” *Brain Research Reviews*, lx, 2 (2009): 306–26; and Franck Di Renzo et al., “Online and Offline Performance Gains Following Motor Imagery Practice: A Comprehensive Review of Behavioral and Neuroimaging Studies,” *Frontiers of Human Neuroscience*, x (2016): 315.

edge of how to accommodate various conditions on the field so as to bring about their desired end.⁶

Sports are an obvious place to look for examples of mental rehearsal, but that should not give readers the impression that this phenomenon is tied uniquely to competitive pursuits. The night before a long trip, one might imaginatively simulate possible routes to the airport, paying special attention to potential road closures and traffic. Let's focus on this everyday example.

TRAFFIC: Tom has to wake up early to catch a flight from Los Angeles to New York. Traffic in Los Angeles is stressfully unpredictable and predictably stressful, and Tom, sensitive to this fact, has heightened anxiety before the trip. To make it to LAX from Tom's West Hollywood bungalow, he would typically take La Cienega, but recent construction has rendered that typical route unreliable. To ease his anxiety, and to better prepare for the journey ahead, Tom mentally rehearses the trip as follows: "If La Cienega is jammed from construction, I can take the 10 to the 405 and only lose a few minutes."

It is, of course, inessential that Tom engage in inner monologue; he might simply imaginatively work through the sensations of navigating city traffic. One plausible upshot of Tom's mental rehearsal in TRAFFIC is that he comes to know (or is at least in a position to know) that he can make it to the airport on time by employing one of two means; the first is to take La Cienega, and the second is to take a detour down the 10 and 405. The next morning, when construction renders the first route too slow, he takes the second, and he makes it to the airport on time. This is obviously something Tom does intentionally, and in so doing manifests *both* knowledge of how to get from West Hollywood to LAX, *and* knowledge that the means he employs (taking the detour) are very likely sufficient to bring about his desired end, a timely airport arrival.

In Tom's case, to make things especially clear, the following three claims are true:

⁶While the above sports examples are meant to illustrate how mental rehearsal/imaginative simulation can furnish one with the relevant propositional knowledge of means, it's worth registering that—perhaps more controversially—some sports psychologists take knowledge of the relevant means (at least in some domains of athletic performance) to be inextricably dependent upon mental rehearsal. This emphasis can be found in the sports psychology of putting in golf (for example, Bob Rotella, *Putting Out of Your Mind* (Simon and Schuster, 2001)); according to Rotella, for example, imaginative simulation of the line of the putt is necessary for knowing the correct line of the putt (where to aim and with what speed), and an inclination to undertake such imaginative rehearsal is common to most elite putters.

- T1. *Tom makes it to the airport on time intentionally,*
- T2. *Tom knows how to get from West Hollywood to LAX, and*
- T3. *Tom knows that in order to get from West Hollywood to LAX, one either takes La Cienega, or one takes a detour down the 10 and 405.*

Cases like TRAFFIC suggest that, in virtue of knowing how to φ , by mentally rehearsing φ -ing one can acquire knowledge that, for some means ψ , it is sufficiently likely for one to succeed at φ -ing by ψ -ing. In virtue of Tom knowing how to navigate Los Angeles, and in particular, in virtue of Tom knowing how to get from West Hollywood to LAX, he can—to a first approximation—come to know via engaging in mental rehearsal a proposition concerning the means by which one gets to LAX in various counterfactual circumstances. In this way of viewing Tom's case, T2 is explanatorily prior to T3, since T2 is what puts Tom in a position to know that *in order to get to West Hollywood from LAX, one either takes La Cienega, or one takes a detour down the 10 and 405*, the proposition featuring in T3.

Why think that what Tom gains through mentally rehearsing his know-how is in the market for *know-that*? Well, he gains a non-accidentally true belief, which is the mark of knowledge. The truth of the belief is stipulated; its non-accidentality is, of course, the interesting and distinctive part of the story. Tom's belief that the means he employs (taking the detour) are very likely to bring about his timely airport arrival *inherits* its non-accidentality from the know-how underwriting it.

To expand on this last claim a bit, consider that, insofar as Tom's know-how regarding navigating Los Angeles, and in particular getting from West Hollywood to LAX, exhibits a certain kind of counterfactual success. He would succeed in arriving at LAX (on time) across a range of counterfactual scenarios in which he tried or intended to do so—scenarios in which the weather is a bit worse, in which the flight's departure time is a bit earlier or later, in which he is fighting a cold, in which the make and model of his car varies, in which he decides to make morning tea rather than coffee, and so on.⁷ Part of what makes it the case that Tom possesses this know-how is that, when he tries or intends to navigate to LAX, the means he employs reliably ensure his (likely) on-time arrival across these variations in circumstance.⁸ His attempts are, so to speak, safely connected to his success. And given that Tom's know-how is grounded in this safe connection

⁷ Compare here with the discussion of Phelps above.

⁸ Katherine Hawley, "Success and Knowledge-How," *American Philosophical Quarterly*, XL, 1 (2003): 19–31; and Romy Jaster, *Agents' Abilities* (De Gruyter, 2020).

to success, Tom's mentally rehearsing his know-how affords the beliefs he acquires during mental rehearsal a safe connection to the truth about how he would likely succeed.⁹ The know-that that results from mental rehearsal merely propositionally encodes the skill rehearsed. One might say that the justificatory heavy-lifting is done by the skill being mentally rehearsed; the justification for Tom's know-that is "easy" precisely in that it is parasitic on the know-how it encodes, as opposed to, for example, being inferred from any other of his justified beliefs or knowledge.

Thus far, we have argued only that Tom's *actually* mentally rehearsing his knowledge of how to navigate LA traffic yields knowledge that in order to get from West Hollywood to LAX, one either takes La Cienega, or one takes a detour down the 10 and 405. But suppose Tom never bothered to rehearse, as many of us might refrain from doing, especially when the stakes are low. He would, presumably, nevertheless *be in a position to know* that proposition concerning the means likely to bring about his desired end; after all, he could come to know it by mentally rehearsing the relevant skill.

In sum, mental rehearsal enables practitioners to rehearse a skill, or simulate a movement, *without actually executing it, much as, for example, imagination allows us to simulate possibilities without actualizing those pos-*

⁹There is a structural worry here: know-how is widely (albeit not unanimously) thought to be *less resilient* against epistemic luck than propositional knowledge is. This is a point that has been made variously by Yuri Cath, "Knowing How Without Knowing That," in John Bengson and Marc A. Moffett, eds., *Knowing How: Essays on Knowledge, Mind, and Action* (Oxford University Press, 2011), pp. 113–35; Ted Poston, "Know How to Be Gettiered?," *Philosophy and Phenomenological Research*, LXXIX, 3 (2009): 743–47; J. Adam Carter and Duncan Pritchard, "Knowledge-How and Epistemic Luck," *Noûs*, XLIX, 3 (2015): 440–53; and others. First, we want to emphasize that it is mistaken to infer from the kind of inheritance claim being made above that know-how and know-that must be *equally* resistant to epistemic luck. *Resistance* to being undermined by epistemic luck (in the debate registered above) is thought to matter for exercises of know-how and states of know-that; one can—as per Cath, "Knowing How Without Knowing That," *op. cit.*—exercise know-how in screwing in a lightbulb even if one's belief about a way to screw in the lightbulb is Gettiered. But, *even if that is right*, it might be that one's reliable ability to screw in lightbulbs (in the very case that the corresponding target belief is Gettiered) is extremely high, *as high* for all that has been said, as the level of reliability that, for example, one's perceptual or memorial capacities would need to be to furnish one with propositional knowledge. Second, and perhaps more importantly, the core position advanced here is best interpreted (as we will go on to develop it in later sections) as a view about *prima facie* propositional justification. The view we embrace is *not* a kind of 'closure' thesis to the effect that if you know how to φ and mentally rehearse that know-how, then you know *that* (for some relevant skill-encoding proposition); the *prima facie* justification furnished to one by mentally rehearsing their know-how *puts one in a position* to know. When all goes well (as in the case of Tom) one then gains knowledge; in other cases, one's *prima facie* justification may not secure *ultima facie* justification, and one thereby falls short of knowing.

sibilities.¹⁰ In rehearsing, whether it concerns navigating traffic on the way to the airport or kicking the game-winning field goal in the play-offs, one comes to know (weaker: puts oneself in a position to know) a proposition that encodes some aspect of the skill being rehearsed, and that specifies the means by which one will achieve a desired end (arriving at the airport on time, kicking the game-winning field goal, and so on).

Let us briefly remark on one potential objection. In particular, one might not find it obvious that mentally rehearsing one's knowing how to φ , in the absence of knowing *that* one knows how to φ , suffices to put one in a position to know the relevant skill-encoding propositions. To see the problem, suppose that an underconfident locksmith has a mismeasure of her own abilities to pick a warded lock. Though she in fact knows how to do so perfectly well, her irrational underconfidence leads her to doubt that she has the relevant know-how when facing a particular warded lock she is called upon to pick.¹¹ Suppose that, despite this underconfidence, she undergoes the relevant mental rehearsal before, to her surprise, picking the lock successfully. The view under consideration implicates, counterintuitively, that our underconfident locksmith is in a position to know the relevant skill-encoding propositions, and simply in virtue of the know-how she (unbeknownst to her) possesses. A more plausible view—as this line of thought goes—will be more restrictive: the locksmith is put in a position to know skill-encoding propositions not simply in virtue of possessing the relevant know-how, but rather in virtue of *knowing that* she possesses the relevant know-how—knowledge our underconfident locksmith lacks, though which she could gain through additional confirmation of her skill.

This objection suggests that possessing a skill does not suffice, all on its own, to put a skilled agent in a position to know a corresponding skill-encoding proposition. Instead, possessing a skill suffices, all on its own, to *prima facie* propositionally justify a skilled agent in believing a corresponding skill-encoding proposition.¹² In unfavorable epis-

¹⁰ Indeed, this general idea—that is, that mental simulation is closely connected to counterfactual knowledge—is central to a prominent strand of thinking in modal epistemology about philosophical knowledge as sourced in epistemically efficacious thought experiments. See, for example, Timothy Williamson, *The Philosophy of Philosophy* (Wiley-Blackwell, 2007), chapters 5–7.

¹¹ Her underconfidence might function as a “psychological defeater” à la Jennifer Lackey, “Testimonial Knowledge and Transmission,” *The Philosophical Quarterly*, XLIX, 197 (1999): 471–90; and Jennifer Lackey, “Socially Extended Knowledge,” *Philosophical Issues*, XXIV, 1 (2014): 282–98.

¹² An anonymous reviewer for this JOURNAL invites us to clarify how this idea applies to, for example, young children and animals. Such agents plausibly manifest skills and

temic conditions, perhaps those involving considerations of underconfidence, defeat, or other knowledge-undermining features, this *prima facie* propositional justification may not yield *ultima facie* justified beliefs. The underconfident locksmith possesses—and can avail herself of—this *prima facie* justification even if, were she to believe the corresponding skill-encoding proposition, she would fail to know it. But since TRAFFIC is not a case of mental rehearsal performed in inhospitable epistemic conditions, T1–T3 still hold.

In the next section we show how our picture of the epistemology of mental rehearsal is part of a much more general perspective on the epistemology of skills and the propositions encoding them.

II. SUPPOSITIONAL REASONING AND EASY KNOWLEDGE

The last section explored what might be considered “practical skills,” concerning how to kick a field goal, how to navigate traffic, and so on. But a parallel point can be made about cognitive skills, which we explore here.

One issue that animated epistemologists in the 1990s and 2000s was how to explain the possibility of *basic knowledge*.

Basic Knowledge: S has basic knowledge of P just in case S knows P prior to knowing that the cognitive source of S’s knowing P is reliable.¹³

know-how but less obviously have the capacity to believe corresponding skill-encoding propositions or, perhaps further, to be *prima facie* propositionally justified in believing (or to know) such propositions.

There is, however, reason for optimism that animals and children are in a position to possess propositional knowledge of the sort implicated by skilled and intentional behavior, even (i) absent the capacity to verbally articulate that knowledge (J. A. Fodor, “Appeal to Tacit Knowledge in Psychological Explanation,” this JOURNAL, LXV (1968): 627–40, p. 34; Jason Stanley, *Know How* (Oxford University Press, 2011), pp. 157–63; and Jason Stanley and Jon W. Krakauer, “Motor Skill Depends on Knowledge of Facts,” *Frontiers of Human Neuroscience*, VII (2013): 1–11, p. 6); and (ii) insofar as they would have knowledge of means, absent the concept of probability or the capacity to conceptualize means as such (Pavese and Beddor, “Skills as Knowledge,” *op. cit.*, p. 615). This is so even if such agents lack the capacities for what might be thought of as *paradigmatic, explicit* belief (Michael Brownstein and Eliot Michaelson, “Doing Without Believing: Intellectualism, Knowledge-How, and Belief-Attribution,” *Synthese*, CXCIII, 9 (2016): 2815–36, p. 2826).

We follow these authors in thinking that non-human animals and young children are not precluded from having knowledge of skill-encoding propositions. Even views stronger than ours are developed so as to be sensitive to (or at least compatible with) the actuality of, for example, Fido catching the rabbit.

¹³ This is the presentation due to Cohen, “Why Basic Knowledge Is Easy Knowledge,” *op. cit.*, p. 417. See also Weisberg, “The Bootstrapping Problem,” *op. cit.*; and Jesper Kallestrup, “Bootstrap and Rollback: Generalizing Epistemic Circularity,” *Synthese*, CLXXXIX (2012): 395–413.

What made the issue particularly urgent was that many popular epistemological views—including versions of reliabilism, virtue epistemology,¹⁴ and evidentialist foundationalism—countenanced basic knowledge, but to countenance basic knowledge was, ostensibly, to license problematic justificatory bootstrapping.¹⁵

To see the problem, start with a case due to Vogel:

[Roxanne] believes implicitly what her gas gauge says, without knowing that the gauge is reliable. But she undertakes the following, admittedly curious, procedure. She looks at the gauge often. Not only does she form a belief about how much gas is in the tank, but she also takes note of the state of the gauge itself. So, for example, when the gauge reads ‘F,’ she believes that, on this occasion, the tank is full. She also believes that, on this occasion, the gauge reads ‘F.’ Moreover, Roxanne combines these beliefs; she believes:

[G] On this occasion, the gauge reads ‘F’ and *F*.

... Now, it is a completely straightforward logical consequence of [G] that:

[A] On this occasion, the gauge is reading accurately.

Assume that Roxanne deduces [A] from [G]. Deduction is certainly a reliable process, so there is no loss of reliability at this step. Consequently, it seems that Roxanne must be credited with knowing [A].¹⁶

By repeating applications of this procedure of checking the gauge, Roxanne could infer *that the gas gauge is generally reliable*. Vogel thinks that Roxanne’s belief in A, along with the inference about the gauge’s general reliability, are illegitimately “bootstrapped” from her merely reading the gas gauge.¹⁷ We could even imagine Roxanne continuing

¹⁴ See here, in particular, Sosa’s attempt to respond to the problem of basic knowledge, by identifying species of ‘virtuous circularity’ in his Ernest Sosa, *Reflective Knowledge: Apt Belief and Reflective Knowledge*, Volume II (Oxford University Press, 2009).

¹⁵ This problem was originally due to Fumerton, *Metaepistemology and Skepticism*, *op. cit.*; and picked up again by Vogel, “Reliabilism Leveled,” *op. cit.* However, see also Cohen, “Basic Knowledge and the Problem of Easy Knowledge,” *op. cit.*; Cohen, “Why Basic Knowledge Is Easy Knowledge,” *op. cit.*; Stewart Cohen, “Bootstrapping, Defeasible Reasoning, and A Priori Justification,” *Philosophical Perspectives*, xxiv (2010): 141–59; Markie, “Easy Knowledge,” *op. cit.*; Titelbaum, “Tell Me You Love Me,” *op. cit.*; Weisberg, “Bootstrapping in General,” *op. cit.*; and Weisberg, “The Bootstrapping Problem,” *op. cit.*

¹⁶ Vogel, “Reliabilism Leveled,” *op. cit.*, p. 613.

¹⁷ Vogel’s aim was to raise a problem for Goldman-style reliabilism, but Cohen showed that the problem was much more general to any view that admits of basic knowledge, including, for instance, foundationalist evidentialism. See Cohen, “Basic Knowledge and the Problem of Easy Knowledge,” *op. cit.*; and Cohen, “Why Basic Knowledge

this process over and over again to become confident in the “super-reliability” of her gas gauge.¹⁸ (Arbitrarily secure) knowledge of the reliability of one’s gas gauge is surely not so easy as *merely* checking the gauge, no matter how many opportunities one has to check. Many thought, more generally, that one cannot bootstrap one’s way into knowledge that one’s belief-formation methods are reliable merely by employing those very methods.¹⁹ There are, of course, many different proposed solutions to the problem of easy knowledge, and we cannot hope to do them all justice here.

Instead, we hope to explore one throughline in the literature surrounding the easy knowledge problem, namely that of the epistemology of suppositional reasoning, and show its intimate connection to mental rehearsal. Consider, to that end, the following example:

SUPPOSITION: In a reflective moment and blindfolded, Stew anticipates various experiences $E_1, E_2, E_3 \dots$ he might undergo upon removing his blindfold. Perhaps the wall will appear red to him, or blue, or $\dots$. He also considers various hypotheses $H_1, H_2, H_3 \dots$ about the color of the wall, that it is red, that it is blue, or $\dots$. As he imagines these various experiences and hypotheses, he considers which hypothesis would be most reasonable to believe given various experiences, perhaps by reasoning thusly: “Supposing that the wall looks red, that the wall is red best explains why this is so. Supposing that the wall looks blue, that the wall is blue best explains why this is so. $\dots$ ” Stew discharges his assumptions and, making an inference to the best explanation, comes to believe various contingent material conditionals of the form *if E , then H* ; for instance, *if the wall looks red, the wall is red*. Blindfold removed, Stew has a (veridical) experience as of a red wall and comes to know the wall is red.²⁰

Is Easy Knowledge,” *op. cit.* And this not only rendered knowledge of the reliability of one’s faculties or methods too easily attained, but licensed a certain objectionably cavalier anti-skepticism. For suppose that I open my eyes and see what looks like a red table, from which, according to theories that admit of basic knowledge, I am in a position to know that the table is red. Moreover, suppose that I know (*a priori*) that if the table is red, it is not white but illuminated by red lights. By an application of *modus ponens*, I am in a position to know that the table is not white but illuminated by red lights. Similar reasoning puts me in a position to know that I am not a brain in a vat being radically deceived about the color of the room’s furniture. But anti-skepticism is surely not so easily attained. The “closure problem” is sometimes given a separate treatment. See Stewart Cohen, “Contextualism and Skepticism,” *Philosophical Issues*, x (2000): 94–107; Cohen, “Basic Knowledge and the Problem of Easy Knowledge,” *op. cit.*; and Markie, “Easy Knowledge,” *op. cit.* However, those details are not important for our presentation here.

¹⁸ Weisberg, “Bootstrapping in General,” *op. cit.*

¹⁹ For helpful overview, see Weisberg, “Bootstrapping in General,” *op. cit.*; and Titelbaum, “Tell Me You Love Me,” *op. cit.*

²⁰ We borrow this general idea from a similar case, “The Explainer,” which appears in John Hawthorne, “Deeply Contingent *A Priori* Knowledge,” *Philosophy and Phenomenological Research*, LXV, 2 (2002): 247–69.

According to Hawthorne and Cohen, the upshot of Stew's suppositional reasoning is that he gains contingent *a priori* knowledge of various material conditionals of the form *if E, then H*, in this case various material conditionals concerning the relationship between how things look and their color.²¹ Cohen, building off Pollock, suggests that the epistemology of suppositional reasoning resembles that of conditional proof, in which I assume *p* to derive *q* and can then discharge my assumption and infer *if p, then q*.²² In SUPPOSITION, of course, Stew is not relying on any natural deduction system with derivation rules such as a rule of conditional proof. Instead, he is relying on what might be called the "perception rule," an informal, defeasible rule of inference according to which:

S looks C

Therefore, S is C.

As we invoke it here, the perception rule is highly schematic; it captures quasi-inferential mental transitions from something's *looking* a certain way to its *being* a certain way, and one's *competence with* the perception rule is a competence in making the relevant mental transitions.

With that said, note that there is nothing special, at least according to Cohen-cum-Pollock, about the perception rule with respect to its potential to confer justification to competent agents; it is a feature of defeasible inference rules more generally (even more "formal" rules like statistical syllogism) that, when applied in suppositional reasoning, afford a defeasible *a priori* justification for believing the corresponding contingent material conditionals encoding particular instantiations of the rule.

How, exactly, can suppositional reasoning play this justificatory role? According to Cohen:

In the case of the suppositional reasoning, I can arrive at the relevant conditional via trivial reasoning *simply in virtue of my competence in using the rule*. It is important to see that the conditional is not a premise in the reasoning. If we view the conditionals corresponding to inference rules as premises in our reasoning, then as Lewis Carroll observed, we

²¹ Hawthorne, "Deeply Contingent *A Priori* Knowledge," *op. cit.*; Cohen, "Basic Knowledge and the Problem of Easy Knowledge," *op. cit.*; Cohen, "Bootstrapping, Defeasible Reasoning, and *A Priori* Justification," *op. cit.*; and Stewart Cohen, "Suppositional Reasoning and Perceptual Justification," *Logos and Episteme*, VII, 2 (2016): 215–19.

²² Cohen, "Bootstrapping, Defeasible Reasoning, and *A Priori* Justification," *op. cit.*; and John Pollock, *Cognitive Carpentry: A Blueprint for How to Build a Person* (MIT Press, 1995).

are committed to a vicious regress of premises. But crucially, the bootstrapping reasoning is not independent of my justification for believing the conditional. I can engage in the bootstrapping reason only if I can competently use the inference rule. *And if I am so competent, I am (propositionally) justified in believing the conditional.* So I cannot do the bootstrapping reasoning unless I am already justified in believing the conditional. Moreover, the bootstrapping reasoning cannot increase my degree of justification for the reliability of my color vision beyond the initial *a priori* degree of justification afforded by the suppositional reasoning. The only source of justification in the suppositional reasoning comes from the application of the perceptual inference rule.²³

It is an agent's *competence with the perception rule* that affords that agent a *a priori* justification for believing a proposition that encodes an instance of the rule itself. In SUPPOSITION, Stew's justified belief that *if the wall looks red, the wall is red* encodes an instance of the rule according to which something's looking thus-and-so is defeasible reason to conclude that it is thus-and-so, a rule with which, by stipulation, he is competent. Moreover, nothing in this explanation requires that Stew *actually* engage in the relevant suppositional reasoning in order to be justified in believing that the wall is red (on the basis of its looking red). Instead, when an agent like Stew manifests his competence with the perception rule in suppositional reasoning, he can avail himself of the *a priori* propositional justification *he already has in virtue of that very competence*; in other words, Stew's competence with the perception rule puts him in a position to know various conditionals encoding it.²⁴

There are two things to note about this rational competence-based view of bootstrapping. The first is that, by engaging in Cohen-cum-Pollock bootstrapping reasoning, one cannot *gain* justification one

²³ Cohen, "Bootstrapping, Defeasible Reasoning, and *A Priori* Justification," *op. cit.*, pp. 154–55, our italics.

²⁴ Granted, facts about an agent's conceptual repertoire may constrain the propositional justification available to her via her competence with the perception rule. An agent lacking the concept CERULEAN but otherwise competent with the perception rule, for instance, may not have propositional justification for believing the conditional *if S looks cerulean, S is cerulean* precisely because, given her limitations, nothing can *look cerulean* to her, nor could she believe or suppose that something looks cerulean, at least not until she acquires CERULEAN.

How does this square with Cohen's claim—one which we take on board—that an agent's competence with the perception rule affords her *a priori* propositional justification for conditionals encoding that rule? We take it that the most plausible answer here is a straightforward one: the perception rule gives one justification only for instances where one possesses the relevant concepts. If one lacks concepts required to grasp the relevant instance of the perception rule, one isn't in a position to apply the rule to the instance. Thanks to an anonymous reviewer for this JOURNAL for inviting us to clarify this point.

did not already have; this avoids the problematic form of bootstrapping discussed by Vogel and Weisberg, which is predicated on agents *gaining* knowledge of the reliability (or super-reliability) of their belief-formation processes by deploying those very processes.²⁵ If one is competent with an inference rule, then one is *already* justified in believing conditionals encoding instances of it.²⁶ In other words, agents competent with the perception rule have a *a priori* propositional justification for believing that perception is reliable, of which agents can avail themselves by engaging in a bit of suppositional reasoning.²⁷

The second thing to note is the striking parallel between the conditions of intentional action and the conditions of rational or justified belief deploying the perception rule. Recall the TRAFFIC case from section 1, in which:

- T1. *Tom makes it to the airport on time intentionally,*
- T2. *Tom knows how to get from West Hollywood to LAX, and*
- T3. *Tom knows that in order to get from West Hollywood to LAX, one either takes La Cienega, or one takes a detour down the 10 and 405.*

In section 1, we argued that T2 explained T3, the latter of which encoded Tom's skill in navigating Los Angeles traffic. The knowledge featuring in T3 was "easy" because it was parasitic on his know-how, and he could avail himself of justification for this knowledge by mentally rehearsing the relevant skill. In unreflective moments, or perhaps if circumstances were a bit different and Tom didn't mentally rehearse, he would nevertheless be in a position to know the relevant skill-encoding proposition. And note that in SUPPOSITION, the following three claims are true:

- S1. *Stew knows the wall is red,*

²⁵ Vogel, "Reliabilism Leveled," *op. cit.*; and Weisberg, "Bootstrapping in General," *op. cit.*

²⁶ For relevantly similar accounts, see John Turri, "Contingent *A Priori* Knowledge," *Philosophy and Phenomenological Research*, LXXXIII, 2 (2011): 327–44; and Ralph Wedgwood, "Primitively Rational Belief-Forming Processes," in Andrew Reisner and Asbjørn Steglich-Petersen, eds., *Reasons for Belief* (Cambridge University Press, 2011), pp. 180–200. One might also look to Timothy Williamson, "The Contingent *A Priori*: Has It Anything to Do With Indexicals?," *Analysis*, XLVI, 3 (1986): 113–17; and Timothy Williamson, "The Contingent *A Priori*: A Reply," *Analysis*, XLVIII, 4 (1988): 218–21; in reply to Gareth Evans, "Reference and Contingency," *The Monist*, LXII, 2 (1979): 161–89; Gareth Evans, *The Varieties of Reference* (Oxford University Press, 1982); and Graham Oppy, "Williamson and the Contingent *A Priori*," *Analysis*, XLVII, 4 (1987): 188–93, for a precursor.

²⁷ If Hawthorne, "Deeply Contingent *A Priori* Knowledge," *op. cit.*, is correct, then if it is true that S can rationally believe P on the basis of E, S has a *a priori* propositional justification for the contingent proposition *if E, then P*. See especially *ibid.*, pp. 250–52.

- S2. *Stew is competent with (knows how to employ) the perception rule, and*
 S3. *Stew knows that if the wall looks red, the wall is red.*

According to the rational competence-based view we've been exploring, S2 explains S3. In particular, Stew is (*a priori*) propositionally justified in believing S3 in virtue of S2, and he can avail himself of this justification by engaging in trivial suppositional reasoning that employs the perception rule. Stew's knowledge that *if the wall looks red, the wall is red* plays a relatively minor role in the drama of his gaining perceptual knowledge. His possessing such knowledge (or being in a position to, in virtue of facts about what he is propositionally justified in believing) is not a part of what explains S1; it is instead a necessary consequence of S2.

This parallel between TRAFFIC and SUPPOSITION is important and non-coincidental. It suggests that practical skills play a role in the epistemology of mental rehearsal analogous to cognitive skills in the epistemology of suppositional reasoning. *In each case, by imaginatively simulating one's (practical or cognitive) skills, one can come to possess easy know-that encoding the skill being simulated.* Moreover, even when one does not imaginatively simulate one's (practical or cognitive) skills, one is in a position to acquire easy know-that encoding the relevant skills. Consequently, in both cases, the resulting know-that is "easy," ultimately grounded in one's skill, competence, or know-how.

The connection between mental rehearsal and suppositional reasoning has been largely ignored in the action theory debate.²⁸ Our proposal explains the connection between intentional action, know-how, and know-that—as we will see in more detail in sections III–IV—in a way that not only renders the connection between the epistemic and actional non-mysterious, but also renders the connection between the practical and the cognitive non-mysterious, via connecting the epistemologies of mental rehearsal and suppositional reasoning.

Of course, in articulating our view, and in developing the parallels between the epistemologies of mental rehearsal and suppositional reasoning, we have relied on the idea that the agents engaged in rehearsal or supposition are in a position to know rather task-specific, plan-encoding propositions. One might worry that this is *too demanding* on the agent, in the sense that they might have only limited powers of imagination, or that their mental rehearsal might be highly general or partial. Moreover, even those with know-how who undergo mental

²⁸ Though see Carlotta Pavese, "Lewis Carroll's Regress and the Presuppositional Structure of Arguments," *Linguistics and Philosophy*, XLV, 1 (2022): 1–38, for recent discussion of presuppositions in connection with Carroll's Regress.

rehearsal that *is* highly specific are not “clairvoyant”—that is to say, the imaginative simulations they undertake will at best be rough approximations of the relevant means that they in fact go on to take—except in the most highly idealized cases.

To make this worry more concrete, one might consider for example the kind of mental rehearsal undertaken by expert athletes such as Tiger Woods, who may have robust powers of imagination in matters of golf, but who describes mentally rehearsing a mere partial visualization of the swing. How is Tiger Woods in a position to know a task-specific, plan-encoding proposition via mental rehearsal when that mental rehearsal is “incomplete”?

We grant that, when mentally rehearsing a skill, one might only imagine exercising that skill in some *general* sense, or one might imagine *approximately* the circumstances one will face in actually employing the skill. And so, through engaging in mental rehearsal, one might *come to know* something more general than or approximating what one in fact does intentionally. This does not undermine the idea that, in virtue of possessing a practical skill, one *is in a position to know* a task-specific, plan-encoding proposition too. Presumably, in virtue of one’s skill, one is in a position to know task-specific, plan-encoding propositions that are *as specific or granular as* the skill itself, even if one is also in a position to know more related but general or partial plan-encoding propositions. What this objection shows, we think, is only that the final and completed epistemology of skill will have to explain how skills can put one in a position to know these other, more general, partial, or approximating propositions too.

In the next section, we explore the implications of our account for theories of control, and in particular, the extent to which epistemic control characterizes intentional action.

III. THE EPISTEMIC CONDITION ON ACTIONAL CONTROL

Here, we apply the epistemological framework developed in sections I–II to offer a novel account of the epistemic condition on actional control.

TRAFFIC and SUPPOSITION motivated the general idea that imaginatively simulating skilled behavior puts an agent in a position to know a proposition encoding the skill(s) being simulated. Tom (TRAFFIC) imagined his way around Los Angeles traffic, coming to know that he could take a detour in pursuit of a timely arrival at LAX, and Stew (SUPPOSITION) reflected on potential experiences and the hypotheses that would best explain them, coming to know that if the wall looks red, it is red. In each case, what explains the agent’s skill-encoding

knowledge is not any other beliefs or knowledge but rather the skill itself; the latter renders the former “easy.”

Insofar as our view is motivated by attending to these sorts of cases, two points of clarification about them are in order. First, although TRAFFIC and SUPPOSITION involved agents imaginatively simulating a skill by consciously and articulably reflecting on that skill, these features of the case were inessential. Both Tom and Stew were *in a position to know* the relevant proposition without engaging in mental rehearsal or suppositional reasoning, and, if we imagine that Tom or Stew were of limited powers of imagination or articulation, it would only follow that *coming to know* what they were in a position to know would be more difficult than otherwise.²⁹

But once we’ve acknowledged that *possessing* skills begets *a priori* justification for skill-encoding propositions, one might naturally wonder whether *manifesting*—and not *merely* possessing—skills bears an important relationship to “practical knowledge,” the kind of knowledge that agents have about what they’re doing when they act intentionally. To that end, we consider how our position interacts with several distinct conceptions of the epistemic condition on actional control.

To start, let’s consider Pavese and Beddor’s recent defense of what they call an “Epistemic Theory of Control,”³⁰ according to which the control characteristic of intentional actions satisfies the *Epistemic Condition* (or “EC”):

EC: Whenever an agent φ -s intentionally, they know that they are φ -ing, and they have this knowledge in virtue of their knowledge of how to φ .³¹

What EC purports to capture is the distinctive sense in which intentional actions are *under the agent’s control as they unfold* in a way that

²⁹ So long as Tom’s imaginative simulation of driving to LAX is guided by his knowledge of how to navigate LA traffic, he is in a position to know a proposition encoding the means to his desired end. If his know-how is inarticulable, Tom will not be in a position to know this proposition under a “semantic mode of presentation,” a way of knowing that enables him to articulate and perhaps defend what he knows. Similarly, so long as Stew’s suppositional reasoning is guided by his knowledge of how to employ the perception rule, he is in a position to know a proposition encoding the relationship between how things look and how they are, perhaps only under some non-semantic mode of presentation if his know-how is inarticulable. For agents with limited powers of articulation, perhaps it is more appropriate to say that they are in a position to know a skill-encoding proposition *de re*. Nothing about our position turns on whether this is so, so we do not address it further.

³⁰ Pavese and Beddor, “Skills as Knowledge,” *op. cit.*

³¹ *Ibid.*, p. 922. Pavese, “Practical Knowledge First,” *op. cit.*, articulates and defends a substantively equivalent thesis, according to which “[o]ne intentionally φ -s just in case when φ -ing, one knows that one is ψ -ing in order to φ for some action ψ , in virtue of knowing how to φ .”

unintentional actions and non-actional behavior are not. Moreover, *EC* gives voice to a broadly Anscombian³² thesis: the mark of intentional control is that whenever one φ -s intentionally, one knows that one is φ -ing.³³

Still, one might worry that *EC* is too strong; it leaves no room for intentional action performed in the absence of *knowledge* of what one is doing. As such, it cannot, in particular, countenance cases of ostensibly intentional actions performed without a belief about what one is doing, nor can it countenance cases of ostensibly intentional actions performed *with* a belief about what one is doing that does not rise to the level of knowledge.³⁴ Consider the first of these:

DISTRACTED: Denise is up early to catch a flight from Los Angeles to New York. To make it to LAX from Denise's West Hollywood bungalow, she would typically take La Cienega, but recent construction has rendered that typical route unreliable, and she knows this. But Denise, like many Los Angeles drivers, is distracted; as she drives, she contemplates whether she'll become a partner at her law firm, the rising price of eggs in her local supermarket, the war in Ukraine, whether she missed trash collection day, and so on. She snaps out of it as she arrives at LAX, having taken the 10 to the 405 because La Cienega was jammed.

Denise makes it to the airport on time intentionally. After all, her arriving at the airport on time has a purposive, means-end structure that manifests a skill. But given her absent-mindedness, it is difficult to attribute to her a belief about what she is doing as she does it; her mind is entirely elsewhere.³⁵ The problem is that, if Denise arrives at the airport on time intentionally *without a belief about what she is doing as she does it*, her case is one of intentional action that violates *EC*, precisely because in lacking the relevant belief, she lacks the relevant knowledge.³⁶

³² Anscombe, *Intention*, *op. cit.*

³³ There are other, stronger characterizations of the contents of practical knowledge. *Nonreductive* views of practical knowledge according to which, when an agent, S, φ -s intentionally, S knows *that she is* φ -ing intentionally. We don't consider these stronger nonreductive views, in part because we are convinced that they face insurmountable objections. See, for instance, Pavese and Beddor, "Skills as Knowledge," *op. cit.*; and Pavese, "Practical Knowledge First," *op. cit.*, for criticism.

³⁴ See, for example, Joshua Shepherd and J. Adam Carter, "Knowledge, Practical Knowledge, and Intentional Action," *Ergo: An Open Access Journal of Philosophy*, ix (2021), for a discussion on this point. However, see Pavese, "Practical Knowledge First," *op. cit.*, section 5, for a reply.

³⁵ See Paul, "How We Know What We're Doing," *op. cit.*, for extended discussion. Pavese, "Practical Knowledge First," *op. cit.*, grants this point.

³⁶ This simple gloss of the case adverts to the epistemological platitude that knowledge entails belief. This platitude has had scattered historical challenges (for exam-

Nevertheless, Denise *is in a position to know* what she is doing as she does it, or to know what means she's employing to secure her desired end, in virtue of the very skill she's employing. Thus, the following three claims are true of DISTRACTED:

- D1. *Denise makes it to the airport on time intentionally,*
- D2. *Denise knows how to get from West Hollywood to LAX, and*
- D3. *As Denise (distractedly) navigates Los Angeles traffic, she is in a position to know that in order to get from West Hollywood to LAX, she's taking a detour down the 10 and 405 because La Cienega was jammed.*

Even if Denise arrives at the airport intentionally *without a belief about what she is doing as she does it*, so that her case is one of intentional action that violates *EC*, one might think that this is a problem with the letter and not the spirit of that proposal. A defender of *EC* might hope to weaken the principle accordingly:

Epistemic Condition - Position to Know (EC-PK): Whenever an agent φ -s intentionally, they are in a position to know that they are φ -ing, and they are in this position in virtue of their knowledge of how to φ .³⁷

EC-PK is strictly weaker than *EC*, since knowing that p trivially entails one is in a position to know that p , but the reverse does not hold. And

ple, as in "unconfident examinee"-style cases (A. D. Woozley, "Knowing and Not Knowing," *Proceedings of the Aristotelian Society*, LIII (1952): 151–72, p. 155; and Colin Radford, "Knowledge—By Examples," *Analysis*, xxvii (1966): 1–11; cf. Keith Lehrer, "Belief and Knowledge," *The Philosophical Review*, Lxxvii (1968): 491–99). A recent line of argument due to David Rose and Jonathan Schaffer, "Knowledge Entails Dispositional Belief," *Philosophical Studies*, CLXVI, 1 (2013): 19–50, reports experimental evidence that folk attributions of 'belief' pattern in a way predicted by the thesis that knowledge entails dispositional belief, even if not as would be predicted by the thesis that knowledge entails occurrent belief. We mention this because, as we see it, the import of DISTRACTED *vis-à-vis EC* does not turn on whether knowledge entails (as Rose and Schaffer think) dispositional rather than occurrent belief. This is because, even though Denise lacks an occurred belief (that is, that in order to get from West Hollywood to LAX, she's taking a detour down the 10 and 405 because La Cienega was jammed) she *also* lacks a dispositional belief to this effect, at least on a plausible construal of the case where she is effectively adapting on the fly without having previously endorsed the relevant content. Note: Denise might well have (as per Robert Audi, "Dispositional Beliefs and Dispositions to Believe," *Noûs*, xxviii (1994): 419–34) a *disposition to believe* a proposition with the relevant content in DISTRACTED, but unless we assume (which we needn't) that she has affirmed that specific content in the past, she *needn't* be understood as hosting the relevant belief dispositionally.

³⁷ Pavese, "Practical Knowledge First," *op. cit.*, suggests this sort of refinement in response to cases like DISTRACTED, although she does not explore it in detail. One might also weaken *EC* so that it only requires intentional actions to manifest *dispositional* or *implicit knowledge* of what one is doing as one does it. According to this alternative, what Denise lacks is *occurrent* or *explicit* knowledge that she is driving to the airport as she is. We think this alternative faces the same questions we raise for *EC-PK*, so we do not consider it at length.

plausibly, *DISTRACTED* satisfies *EC-PK*, for if Denise were to believe, as she was driving to the airport, that she was driving to the airport, her belief would amount to knowledge.

Here is an obvious sense in which *EC-PK* fares better as a general account of epistemic control than *EC*: any case that *EC* would count as an intentional action, *EC-PK* would also count as an intentional action, and intuitively intentional actions (like Denise's arriving at the airport on time) that *EC* leaves out of its explanation, *EC-PK* captures. This advantage is not, strictly speaking, limited to the matters of extensional fit (that is, to capturing cases like *DISTRACTED*). More generally, *EC-PK* insulates itself from difficult questions about belief attribution; to put it colloquially, *whether or not* we can find a way, in any given putative case of intentional action, to plausibly locate in an agent's belief-box propositional contents concerning what the agent is doing as she does it, we can appeal to the fact that she is *in a position* to acquire such a belief, and that that belief would amount to knowledge. Thus, a necessary connection between intentional action and knowledge of what one is doing is preserved, and so an epistemic condition on actional control is preserved, albeit a weaker one than what is posited by *EC*.

Here, we want to note how the view developed in sections I–II can shed light on *EC-PK*, in particular on the 'in virtue of' relation featuring therein. According to our view, Denise is in a position to know what she's doing as she does it *in virtue of her know-how* because knowing what one is doing as one does it is easy knowledge; it's as easy as the knowledge that Tom gained (weaker: was in a position to gain) via mental rehearsal. Denise's skills afford her skill-encoding beliefs a certain privileged epistemic status; this is why Denise's knowledge of how to navigate Los Angeles traffic, as she's taking a detour down the 10 and 405, puts her in a position to know that she's taking the detour because La Cienega was jammed. In other words, D2 explains D3.

This much fits nicely with the view developed by Pavese & Beddor and Pavese, which are explicitly "practical knowledge-first" views of actional control. But we want to note that *EC-PK* has a different sort of problem, one which may not move die-hard practical knowledge-firsters, but which might animate a neutral reader. In particular, there are some cases of ostensibly intentional actions in which an agent is not in a position to know what she is doing, precisely because, were she to form a belief about what she is doing, her belief would not amount to knowledge. Here we note two that have some purchase in the literature:

PARALYSIS: Paolo has recently undergone a medical procedure that required temporarily paralyzing his right arm. The doctors informed him

that the paralysis wears off in four to six hours, and as he sits at home, post-procedure, he notes that only two hours have passed. On a whim, he thinks to himself, “What harm is there in trying?” He raises his arm above his head.³⁸

MARGIN: Marge has finished a double shift at work and is really tired, but she’s signed up for a study in which experimenters have asked her to lie in a bed and to move her finger for as long as she can. Marge starts out moving her finger, but she gradually gets more and more lethargic. As the minutes go by, her finger moves less and less energetically. Eventually, she falls asleep. At some point in the process, she can be correctly characterized as trying to move her finger even when she doesn’t succeed; but as she dozes off, her intentions turn to whims. Even if she forms a true belief that she is trying to move her finger, the belief will be too epistemically hazardous to constitute knowledge. This suggests that an agent may be trying to do something and not know that she is so trying.³⁹

In PARALYSIS, Paolo raises his arm above his head intentionally. His arm-raising is purposive, is non-deviantly initiated and sustained by his intention to raise his arm, and so on. But he is not in a position to know that he is raising his arm above his head, since that proposition is defeated by the doctor’s testimony. In MARGIN, Marge wiggles (or tries to wiggle) her finger intentionally, but she is not in a position to know that she is wiggling (trying to wiggle) her finger, given how close she is to possibilities in which she no longer intends to do so. Given those nearby possibilities, her belief that she is wiggling (trying to wiggle) her finger is unsafe. If PARALYSIS and/or MARGIN are cases of intentional action, then *EC-PK* is too strong a principle to characterize the epistemic condition on actional control; Paolo and/or Marge act intentionally but are *not even in a position to know* what they’re doing as they do it.

One option, of course, is to bite the bullet on cases like PARALYSIS and MARGIN;⁴⁰ if these are penumbral cases within action theory (they are admittedly not paradigmatic ones), one might think, *to the victor go the spoils*. But we think that there is a plausible weakening of *EC* and *EC-PK*—one that is implied by the view developed so far in

³⁸ Kieran Setiya, “Practical Knowledge,” *Ethics*, CXVIII, 3 (2008): 388–409; Paul, “How We Know What We’re Doing,” *op. cit.*; and Carlotta Pavese, “Probabilistic Knowledge in Action,” *Analysis*, LXXX, 2 (2020): 342–56.

³⁹ This case is borrowed—only very slightly modified—from Piñeros-Glasscock, “Practical Knowledge and Luminosity,” *op. cit.* (See especially pp. 1256–59 for discussion.)

⁴⁰ Pavese, “Probabilistic Knowledge in Action,” *op. cit.*; and Pavese and Beddor, “Skills as Knowledge,” *op. cit.*, footnote 3.

sections I–II—that accommodates all of DISTRACTED, PARALYSIS, and MARGIN:

Epistemic Condition - Propositional Justification (EC-PJ): Whenever an agent φ -s intentionally, they are *a priori prima facie* propositionally justified in believing that they are φ -ing, and they have this *a priori prima facie* justification in virtue of their knowledge of how to φ .

Note that *EC-PJ* is strictly weaker than *EC-PK*, at least on the standard assumption that, if one is in a position to know that p , one is propositionally justified in believing that p , but not the other way around. (*EC-PJ* is thus strictly weaker than *EC* too.) Notice how *EC-PJ* explains each of the problem cases above. In DISTRACTED, Denise is *prima facie* propositionally justified in believing that she is driving to the airport because she is manifesting her skill in driving to the airport. This is precisely—in the absence of any defeaters, and given that safety concerns are irrelevant—what puts her in a position to know that she is. In PARALYSIS, Paolo’s belief (if he has one) that he is raising his arm above his head cannot amount to knowledge in light of his doctor’s testimony. Nevertheless, that Paolo is manifesting knowledge of how to raise his arm above his head affords him *a priori prima facie* propositional justification for believing that he is. Although Paolo is not in a position to know that he is raising his arm above his head, given that his *a priori* justification is *ultima facie* defeated, *EC-PJ* plausibly captures a sense in which his know-how grounds a form of epistemic control over his behavior as it unfolds that renders it actional. And in MARGIN, Marge’s belief that she is wiggling (trying to wiggle) her finger is too close to the margins for it to be safe, and so too close to the margins for it to count as knowledge. Even though, like Paolo, Marge is not in a position to know that she is wiggling (trying to wiggle) her finger, she nonetheless has *a priori prima facie* propositional justification for believing that she is. And this, as in Paolo’s case, plausibly affords her belief a special epistemic status about what she is doing as she does it that grounds a form of epistemic control.

In short, one advantage of *EC-PJ* over *EC* and *EC-PK* is how the former handles non-paradigmatic, marginal, or defective cases. Instead of painting us into a dialectical corner, as the latter, stronger theses might do, *EC-PJ* allows us to see each case as occupying a more general position on a spectrum of epistemic control, each of which is ultimately grounded in the relationship between one’s know-how and the *a priori* propositional justification such know-how affords. One virtue of *EC-PJ* is that it can agree with the practical knowledge-firsters about core cases of intentional action while accommodating these more difficult ones too. But if, in these more difficult cases, the agents act

intentionally *without* practical knowledge, it is nevertheless true that they act intentionally with a certain kind of *justification*—justification that would, in more hospitable epistemic environments, yield practical knowledge. One might view our proposal as a “skill-first” picture of actional control, according to which one’s skills yield propositional justification for believing that one is φ -ing as one φ -s. According to this skill-first picture, the epistemic condition on intentional action is captured by *EC-PJ*, not by the stronger, practical knowledge-first principles like *EC* or *EC-PK*.

It’s worth pausing to reflect on what, precisely, one might want out of one’s theory of actional control, and whether, for instance, a principle like *EC-PJ* can really secure those *desiderata*. In the post-Anscombe era of action theory, one gets the sense that any plausible theory of actional control must explain the connection between control and practical knowledge—again, the knowledge one has *of* what one is doing as one does it intentionally. Informally and roughly, this explanation must be able to characterize the sense in which agents acting intentionally are in a position to answer “Why?” questions (Question: “Why are you whistling?” Answer: “To hail a cab.”). On a fairly narrow view of the ambitions of a theory of actional control, to give a theory of actional control *just is* to give a theory of practical knowledge, or is *at least* to give a theory of practical knowledge. *EC* and *EC-PK* are versions of this practical knowledge-first approach to theorizing about actional control, while *EC-PJ* is not.

A related but more modest way to view what we want out of a theory of actional control is that it is (at least) a theory of *whatever puts agents in a position to have practical knowledge in central or well-ordered cases*. Even if one has practical knowledge in well-ordered, central, and paradigm cases of intentional action, it does not follow that whenever one acts with the control characteristic of intentional action, one manifests practical knowledge.⁴¹ As we’ve argued, in epistemically hazardous circumstances like those involving considerations of safety or defeat, practical knowledge might be entirely out of the question. Still, our view preserves something like the *kernel* of practical knowledge, namely the justification that, in hospitable epistemic conditions, puts agents in a position to know what they are doing as they do it. A bit informally, Paolo (from *PARALYSIS*) may lack practical knowledge in light of the relevant defeaters he possesses, but he is surely in a position to answer “Why are you doing that?” with “Out of curiosity,” or

⁴¹ Compare Paul’s “Inferentialist” account of the (contingent inferential) knowledge that agents have of what they are doing as they do it in paradigm cases of intentional action. See Paul, “How We Know What We’re Doing,” *op. cit.*

"To see whether I need to find a new doctor." We can make sense of his behavior as actional *even if* there is a legitimate sense in which it is "actionally defective," to the extent that it falls short of some epistemic ideal of actional control (that of manifesting practical knowledge).

In our view, jettisoning the necessary connection between intentional action and practical knowledge is not, ultimately, a significant cost, especially insofar as we have offered a more general principle that explains why and when agents acting intentionally possess such knowledge. Here, we are broadly sympathetic with Piñeros-Glasscock's remark that,

... [T]here is an important link between intentional action and practical knowledge, one that we must respect if we are to understand their nature; but there may be an important link between the two without the need to posit a necessary connection. . . . One way is to say that. . . the function of the will (understood as the capacity to act intentionally) is to yield practical knowledge. Another (compatible) way is to say that. . . exercises of the will can on their own give us practical knowledge. If either of these ideas is correct, we would expect that cases where an agent acts intentionally with practical knowledge would form a central core of cases of intentional action. However, since powers can sometimes fail to achieve their function, and since even non-inferential sources can fail to yield knowledge in epistemically inhospitable circumstances, both ideas are compatible with the rejection of [a necessary connection between intentional action and practical knowledge].⁴²

EC-PJ is a novel way to understand the epistemic control that is distinctive of intentional action, one which gives voice to *both* the idea that the function of the will is to produce practical knowledge, *and* that the will can give us practical knowledge all on its own. Our position is able to secure these verdicts because it is built out of independently plausible considerations from within the epistemology of skill, and, in particular, the connection between what one knows how to do and what one would be justified in believing. In particular, if *EC-PJ* is correct, then practical knowledge is the mark of central and well-ordered cases of intentional action, but as we move away from these core theoretical data, intentional actions exhibit epistemic control along a spectrum; in epistemically hazardous environments where considerations of safety or defeat are especially salient, one's intentional actions might fail to exhibit the epistemic control characteristic of paradigm cases while nevertheless enjoying some lesser but substantive epistemic status characterized by *EC-PJ*.

⁴² Piñeros-Glasscock, "Practical Knowledge and Luminosity," *op. cit.*, p. 1262.

IV. THE KNOWLEDGE-ACTION PRINCIPLE

We have laid out a position in the epistemology of skill (in sections I–II) that supports a “skill-first” conception of the epistemic condition on actional control (section III). Here, we discuss how our view fits into the larger contours of the debate around the connection between intentional action, know-that, and know-how.

Much of contemporary action theory is oriented around two theses, which we call, following Pavese,⁴³ *Know-How/Intentionality* and *Intentionality/Belief*:

Know-How/Intentionality: If S φ -s intentionally, then S knows how to φ .⁴⁴

Intentionality/Belief: If S φ -s intentionally, then there are some means $m_1, \dots, m_n$ to φ such that S truly believes $m_1, \dots, m_n$ are means for oneself to φ .⁴⁵

The plausibility of *Know-How/Intentionality* derives in no small part from its vindicating the pretheoretic datum that intentional actions are *manifestations of skill* and *under the agent’s control* in a way that distinguishes them from unintentional ones or non-actional behavior. As Ryle famously noted, the clumsy person who trips and falls is importantly different from the clown who “trips and tumbles just as clumsy people do, except that he trips and tumbles on purpose and after

⁴³ Carlotta Pavese, “Know-How, Action, and Luck,” *Synthese*, cxcviii, 7 (2021): 1595–617.

⁴⁴ A partial list includes Gilbert Ryle, *The Concept of Mind* (Hutchinson, 1949); Jason Stanley and Timothy Williamson, “Knowing How,” this JOURNAL, xcvi, 8 (2001): 411–44; Hawley, “Success and Knowledge-How,” *op. cit.*; Jennifer Hornsby, “Agency and Actions,” *Royal Institute of Philosophy Supplement*, lv (2004): 1–23; Jennifer Hornsby, “Ryle’s *Knowing-How*, and Knowing How to Act,” in Bengson and Moffett, eds., *Knowing How*, *op. cit.*, pp. 80–98; Stanley, *Know How*, *op. cit.*; Kieran Setiya, “Knowing How,” *Proceedings of the Aristotelian Society*, cxii (2012): 285–307; J. Adam Carter and Jesús Navarro, “The Defeasibility of Knowledge-How,” *Philosophy and Phenomenological Research*, xcvi (2017): 662–85; Pavese, “Probabilistic Knowledge in Action,” *op. cit.*; Carlotta Pavese, “Knowledge, Action, and Defeasibility,” in J. Brown and M. Simion, eds., *Reasons, Justification, and Defeaters* (Oxford University Press, 2021), pp. 177–200, chapter 8; and Carlotta Pavese and Bob Beddor, “Practical Knowledge Without Luminosity,” *Mind*, cxxxi, 523 (2021): 917–34.

⁴⁵ See Alvin Goldman, *A Theory of Human Action* (Princeton University Press, 1970); Gilbert Harman, “Practical Reasoning,” *The Review of Metaphysics*, xxix, 3 (1976): 431–63; Robert Audi, “Intending,” this JOURNAL, lxx (1973): 387–403; Robert Audi, “Intending, Intentional Action, and Desire,” in Joel Marks, ed., *The Ways of Desire: New Essays in Philosophical Psychology on the Concept of Wanting* (Precedent, 1986), pp. 17–38; Michael Bratman, *Intention, Plans, and Practical Reason* (Harvard University Press, 1987); David Velleman, *Practical Reflection* (Princeton University Press, 1989); Carl Ginet, *On Action* (Cambridge University Press, 1990); Alfred Mele and Paul Moser, “Intentional Action,” *Noûs*, xxviii, 1 (1994): 39–68; and John Gibbons, “Knowledge in Action,” *Philosophy and Phenomenological Research*, lxii, 3 (2001): 579–600.

much rehearsal and at the golden moment and where the children can see him and so as not to hurt himself.”⁴⁶ Knowledge-how, being the mark of skill and a critical component of control, seems especially suited to mark this central action-theoretic difference.

And *Intentionality/Belief* captures the idea that intentional actions issue from plans concerning *how to act*, which are beliefs that encode the means by which one accomplishes a desired end. Intentional actions, in some important sense, *manifest* or *are guided by* one’s plans concerning how to act, and this connection between plan and success is absent in unintentional actions and mere (non-actional) behavior. But familiar problems of epistemic luck—particularly cases where beliefs about plans are only accidentally true—have prompted some fans of *Intentionality/Belief* to reach for a stronger thesis bearing an obvious family resemblance, *Intentionality/Knowledge*:

Intentionality/Knowledge: If S φ -s intentionally, then there are some means $m_1, \dots, m_n$ to φ such that S knows that $m_1, \dots, m_n$ are means for oneself to φ .⁴⁷

Perhaps more cautiously, and being sensitive to cases like DISTRACTED, we might do better to consider a weaker principle:

Intentionality/Position-to-Know: If S φ -s intentionally, then there are some means $m_1, \dots, m_n$ to φ such that S is in a position to know that $m_1, \dots, m_n$ are means for oneself to φ .

Together, *Knowledge-How/Intentionality* and *Intentionality/Position-to-Know* yield the following “*The Knowledge-Action Principle*,” or *KAP*:

KAP: If S φ -s intentionally, then S knows how to φ , and there are some means $m_1, \dots, m_n$ to φ such that S is in a position to know that $m_1, \dots, m_n$ are means for oneself to φ .

If *KAP* is correct, then intentional action bears an important relationship to both know-how and know-that.⁴⁸ Moreover, if *KAP* is correct,

⁴⁶ Ryle, *The Concept of Mind*, *op. cit.*, p. 33.

⁴⁷ See Gibbons, “Knowledge in Action,” *op. cit.*; and Carlotta Pavese, “Knowledge How,” in Edward N. Zalta and Uri Nodelman, eds., *The Stanford Encyclopedia of Philosophy* (Fall 2022 Edition), <https://plato.stanford.edu/archives/fall2022/entries/knowledge-how/>, section 5.

⁴⁸ There are other types of knowledge-that that are often thought to line up with intentional action; compare, for instance, knowledge *that one is φ -ing when one is*, and which Anscombe took to be necessary for φ -ing intentionally. We do not assume that the “plan-based” propositional knowledge captured in *Intentionality/Knowledge*, and by extension in *KAP*, implicates Anscombian practical knowledge-that; however, we will discuss the latter kind of intentional action-relevant propositional knowledge in section IV.

it is highly plausible to think that the relationship it captures is *non-coincidental*. In other words, know-how and know-that do not align in cases of intentional action merely by happy accident.

Pavese has recently argued that *Intellectualism about Know-How* (“*Intellectualism*” for short) offers an elegant account of the non-coincidental alignment of know-how and know-that in cases of intentional action, and so offers an elegant explanation of KAP.⁴⁹ According to *Intellectualism*:

Intellectualism: S knows how to φ only if, for some means ψ for S to φ , S knows that it is sufficiently likely for them to succeed at φ -ing by ψ -ing.

To be clear, proponents of *Intellectualism* are not merely in the business of securing an extensionally adequate principle according to which know-how is *necessarily accompanied by* know-that. They are, instead, defending an explanatory principle according to which one’s know-how is or is *reducible to*⁵⁰ (weaker: is *grounded in*⁵¹) one’s know-that. Of course, the weaker, extensional claim is secured by the truth of the robust explanatory one.

And if *Intellectualism* is correct, we can readily account for why know-how and know-that accompany intentional action in the way specified by KAP. According to this way of thinking, since know-how *just is* a kind of know-that, and, in particular, a kind of know-that encoding the means by which one accomplishes a desired end, it is no coincidence that cases of intentional action will involve know-how and know-that. To wit, such cases will involve know-that encoding the means by which one accomplishes a desired end because the required know-how *just is* such a form of know-that!⁵² Moreover, one might think that *Intellectualism* gains abductive support by its ability to unify and explain the relationships KAP captures.

If our arguments are correct, however, KAP is, strictly speaking, false. Its *approximate truth* consists in the fact that intentional action,

⁴⁹ Pavese, “Knowledge How,” *op. cit.*

⁵⁰ Stanley and Williamson, “Knowing How,” *op. cit.*; Pavese, “Probabilistic Knowledge in Action,” *op. cit.*; and Pavese, “Knowledge, Action, and Defeasibility,” *op. cit.*

⁵¹ John Bengson and Marc A. Moffett, “Two Conceptions of Mind and Action: Knowledge How and the Philosophical Theory of Intelligence,” in Bengson and Moffett, eds., *Knowing How*, *op. cit.*, pp. 3–55; and J. Adam Carter and Ted Poston, *A Critical Introduction to Knowledge-How* (Bloomsbury, 2018).

⁵² Pavese says, “[KAP’s] truth is grounded on the very nature of knowledge-how: one knows how to Φ *in virtue of* knowing, for some means m to Φ , that m is a means for oneself to Φ ” (Pavese, “Knowledge How,” *op. cit.*, section 5). And Pavese says, “The intellectualist picture provides the best explanation for why [KAP] should hold. According to this explanation, [KAP] is true not just out of a coincidental aligning of propositional knowledge and know-how in intentional action. Rather, its truth is grounded in the very nature of know-how.” Pavese, “Knowledge, Action, and Defeasibility,” *op. cit.*, p. 191.

know-how, and know-that align in the way it describes in central or well-ordered cases of intentional action—those performed, notably, in epistemically hospitable environments. Just as the agents in cases like PARALYSIS and MARGIN are precluded from having (or being in a position to have) practical *knowledge* because of salient considerations of defeat and safety, we should expect that agents can act intentionally without being in a position to have knowledge connecting means to a desired end. *KAP* is not true, but it nevertheless captures a kernel of truth around which any plausible theory of action should orient itself: intentional action, know-how, and know-that do not align by coincidence.

The non-coincidence of intentional action, know-how, and know-that is explained by the fact that intentional actions involve the manifestation of skill, the possession of which gives agents *a priori* justification for believing propositions encoding the means by which one accomplishes a desired end, which, in epistemically hospitable environments, yields knowledge.

The more fundamental principle connecting intentional action, know-how, and know-that looks like this:

Knowledge-Action Principle (KAP*)*: If S φ -s intentionally, then S knows how to φ , and there are some means $m_1, \dots, m_n$ to φ such that S is *a priori* (and *prima facie*) propositionally justified in believing that $m_1, \dots, m_n$ are means for oneself to φ .

*KAP** captures what we take to be at the heart of *KAP*. And insofar as one always acts from belief-plans specifying the means by which one can achieve a desired end, the second conjunct of *KAP** will, modulo considerations of defeat and safety, give rise to propositional knowledge, thereby vindicating *KAP* in central cases of intentional action.

What of the claim that *Intellectualism* gains abductive support from its ability to explain *KAP*? We have shown that one can vindicate the kernel of truth in the vicinity of *KAP*, namely *KAP**, without taking on any commitments about the ultimate metaphysics of skill. This is not an argument against *Intellectualism*, at least not a straightforward one. Instead, it is an argument against the idea that *Intellectualism* is *especially poised* to explain the connection between intentional action, know-how, and know-that.⁵³

⁵³ There are, to be sure, intellectualist accounts of skill on the market—most notably, due Stanley and Williamson—which seem perfectly compatible with the crux of our easy knowledge narrative. According to Stanley and Williamson, “to be skilled at the action type of φ -ing is to be disposed to form knowledge appropriate for guiding

One virtue of our position is that we can vindicate the intimate and non-coincidental connection between intentional action, know-how, and know-that that *KAP** purports to capture without thereby necessarily committing to the controversial metaphysics that underwrites *Intellectualism*.⁵⁴ After all, our position is motivated by considerations in the *epistemology*, not the *metaphysics*, of skill; its theoretical neutrality is among its chief virtues.

V. CONCLUSION

By attending to the epistemologies of mental rehearsal and suppositional reasoning, we have offered a “skill-first” conception of the epistemic condition on actional control, according to which practical knowledge is “easy.” Even if it is, the connection between intentional action and practical knowledge is still intimate and important; since skilled behavior necessarily *a priori* propositionally justifies the skilled agent in believing a corresponding skill-encoding proposition, in epistemically hospitable environments, agents acting intentionally will (be in a position to) have practical knowledge. Not only does our account bring to bear upon central topics in the theory of action traditional epistemological questions about basic knowledge, but it charts a middle way through the Scylla of Anscombian theories of action, on the one hand, and the Charybdis of denying any substantive, necessary epistemic condition on intentional action, on the other.

TIMOTHY R. KEARL

University of Glasgow

J. ADAM CARTER

University of Glasgow

tokens of φ -ing.” See Jason Stanley and Timothy Williamson, “Skill,” *Noûs*, LI, 4 (2017): 713–26, p. 715. Correlatively, for the skill-intellectualist, skill *manifests* in knowledge-wh. With this point in mind, we can see how the skill-intellectualist might purport to simply “subsume” the view defended here (at least in well-ordered cases) as follows: according to one version of skill intellectualism, skill manifests in know-wh states, and mental rehearsal is nothing more than one of potentially multiple “triggers” for skill manifestation. According to another, skill is propositional knowledge under a special, practical mode of presentation, and mental rehearsal is a means of knowing the same proposition in a new way. Without assessing the merits of either form of skill intellectualism here, we note simply that one could hold those views and take on board our easy practical knowledge proposal.

⁵⁴ See, for instance, Alva Noë, “Against Intellectualism,” *Analysis*, LXV, 4 (2005): 278–90; Imogen Dickie, “Skill Before Knowledge,” *Philosophy and Phenomenological Research*, LXXXV, 3 (2012): 737–45; and Arie Schwartz and Zoe Drayson, “Intellectualism and the Argument from Cognitive Science,” *Philosophical Psychology*, xxxii, 5 (2019): 662–92, for criticism.