ORIGINAL RESEARCH



Intentional action, knowledge, and cognitive extension

J. Adam Carter¹ · Gloria Andrada²

Received: 2 October 2023 / Accepted: 26 June 2024 $\ensuremath{\textcircled{O}}$ The Author(s) 2024

Abstract

Intentional actions exhibit control in a way that mere lucky successes do not. A longstanding tradition in action theory characterizes actional control in terms of the *knowledge* with which one acts when acting intentionally. Given that action theorists, no less than epistemologists, typically take for granted the orthodox thesis that knowledge is in the head (viz., realized exclusively by brainbound cognition), the idea that intentional action is controlled in virtue of knowledge is tantamount to the idea that the knowledge by which intentional actions exhibit control supervenes intracranially. We raise some challenges for this idea, and in doing show, we show how epistemic theories of actional control are naturally aligned moreso than has been appreciated with cognitive extension in the theory of mind.

Keywords Extended cognition · Intentional action · Know-how · Extended mind

1 Introduction

Intentional action has been closely linked with *knowledge* states, and with exercises, in one way or another, of knowledge. Those who act intentionally, as the thought goes, know about what they're doing.

 J. Adam Carter adam.carter@glasgow.ac.uk
Gloria Andrada

gloriandrada@gmail.com

¹ COGITO Epistemology Research Centre, University of Glasgow, Glasgow, Scotland, United Kingdom

² Instituto de Filosofia da NOVA (IFILNOVA), Universidade Nova de Lisboa, Lisbon, Portugal

One way this idea has been captured is by Anscombe's *Practical Knowledge Principle*, which specifies a sense in which *propositional knowledge* is necessary for acting intentionally.

Practical knowledge Principle (PKP): (Necessarily) If an agent is Φ ing (intentionally and under that description), she knows that she is Φ ing (intentionally and under that description).¹

Second, regardless of whether one accepts PKP (or some related propositional knowledge constraint on intentional action), there is a good deal of support for a related way to connect intentional action and knowledge, by linking it to *knowledge-how*:

(*Knowledge-how/Intentionality (KHI*): If S intentionally Φ s, S knows how to Φ .²

A largely unchallenged assumption by those who embrace PKP is that the kind of knowledge (i.e., non-observational propositional knowledge³) that features in PKP supervenes entirely on intracranial states of the agent, and is realized – as traditional cognitive internalist thinking about the mind holds⁴ – exclusively by processes that play out within the boundaries of the brain.

In a similar vein, those who embrace KHI – a thesis accepted widely by both intellectualists and anti-intellectualist camps – take for granted that the kind of know-how that is pertinent to Φ -ing intentionally supervenes exclusively on the biological agent. This will be *either* because (as intellectualists have it) know-how requires propositional knowledge under a practical mode of presentation⁵, where that knowledge (under that mode) is intracranially generated and stored, *or* because (ii) as anti-intellectualists maintain, the kind of disposition (e.g., an ability) that grounds know-how itself supervenes exclusively on body/brain-bound features of the agent.⁶

In sum, because intentional action has been so often linked to intelligence states via knowledge principles such as PKP and KHI which themselves are viewed as a kind of default through a cognitive internalist lens, it is unsurprising that what is thought to separate action that is intentional from action that is not has itself been viewed through this kind of internalist lens.

¹ See, e.g., Anscombe (1957). Cf., Davidson (1970). However, for some recent support of the view, see Velleman (1989), Setiya (2009), and Campbell (2018). For recent defense of a modified version of PKP, which requires just that one intentionally Φ s only if one knows that they are Φ -ing (and even if they don't know that they are Φ -ing intentionally), see Pavese (2022). Cf., Shepherd and Carter (Forthcoming).

² Example defenders include Beddor and Pavese (2021); Ryle (1949); Williamson and Stanley (2001); Hawley (2003); Hornsby (2004); Stanley (2011); Setiya (2012); Carter and Shepherd (2023).

 $^{^3}$ See, e.g., Paul (2009) and Schwenkler (2012) for discussion on the sense in which practical knowledge is taken to be non-observational.

⁴ We discuss this assumption in more detail in § 2.1.

⁵ See Pavese (2015).

⁶ Though for some more permissive thinking about know-how on anti-intellectualist lines, see Carter and Czarnecki (2016) and Andrada (2022).

Our aim here is to show how a more expanded view of the mind – viz., that which is captured by proponents of extended cognition – turns out to fit well with the idea that intentional actions are knowledge implicating. And, in fact, an openness to extended cognition offers a number of advantages for theories of intentional action, at least for those who take actional control to be secured by knowledge.

Here is the plan. In § 2 we raise two problems for the prospects of making sense of the control implicated by acting intentionally in terms of exclusively biologically generated and maintained knowledge states, in a way that an explicit or tacit commitment to a cognitive internalist picture of the mind would have it; in (§ 2.1) we put pressure on the kind of intentional action-knowledge connection captured by PKP – and which shows that PKP combined with cognitive internalism is too restrictive; we show further that the problem here isn't resolved by weakening the epistemic requirement in PKP from knowledge to even belief, so long as cognitive internalism remains in the background. (§ 2.2) then raises a problem for the KHI+cognitive internalism pairing, which is shown to be too restrictive *regardless* of whether one understands know-how along intellectualist or anti-intellectualist lines. (§ 3) contextualizes these results in wider discussions about the epistemic constraints on intentional action in both epistemology and action theory.

2 Two challenges

2.1 PKP + cognitive internalism

One important motivation for some kind of propositional knowledge condition on intentional action (such as Anscombe's PKP principle) has to do with the *control* with which we act when we act intentionally. One thing that seems to separate an intentional action from unintentional movements (e.g., the clown's tumble versus the klutz's tumble) is that the clown's tumble is controlled, and the klutz's is not. If acting intentionally requires control, where control itself implicates that one's attaining the intended outcome isn't too lucky, then we can make sense of this idea nicely by supposing intentional action requires knowing what you're doing (when doing it) – something the tumbling clown has but the tumbling klutz lacks.

Here is not the place to argue positively for a knowledge condition on conditional action – an important causalist tradition in action theory due to Davidson challenges this idea.⁷ For our purposes, it's worth noting that various philosophers have (in some way) followed Anscombe in thinking that whatever control acting intentionally demands on us is secured through the requirement that acting intentionally is in some way acting *knowledgeably* – knowing (as per PKP) what you're doing while you're doing it.

That said, it is no wonder that PKP has traditionally been interpreted against a backdrop of cognitive internalist thinking – viz., which holds that cognitive processes play out entirely within the boundaries of the skull and skin. For one thing, such thinking was certainly the dominant paradigm when PKP was conceived by Ans-

⁷ See, e.g., Davidson (1978).

combe in the mid-20th century, and it has largely only been since the rise of 4E approaches to cognition in the past decades that cognitive internalism more generally has been called into doubt.⁸

And even so, the move towards rejecting cognitive internalism in the philosophy of cognitive science has yet to receive serious uptake in various other areas of philosophy (e.g., ethics and moral responsibility, action theory⁹, epistemology¹⁰) that have proceeded mostly on the assumption (even if not argued explicitly) that the received cognitive internalist picture best captures the boundaries of the human mind – viz., that commonsense is right, and that cognition takes place in the head (e.g., Adams and Aizawa 2008).¹¹

It is also not surprising that PKP has been taken thus far to fit snugly with cognitive internalist thinking for an entirely different reason, one that – crucially – would be available even for those who are open to allowing (in some circumstances) our cognition to extend (contra cognitive internalism) beyond the bounds of the skull and skin.

In order to get this point into view, it will be useful to recall the familiar contrast case of "Inga and Otto" (Clark & Chalmers, 1998) used to motivate *extended cognition* – viz., the view that, put roughly for now, cognitive processes do not 'play out' exclusively in the head. The comparison between Inga and Otto and how they (respectively) remember what they need to know, is as follows:

Inga: Inga has a normally functioning biological memory. When she learns new information, she stores it in her biological memory (as is typical and familiar) and it is her biological memory which she consults when she needs old information.

Otto: Otto suffers from Alzheimer's disease, and like many Alzheimer's patients, he relies on information in the environment to help structure his life. Otto carries a notebook around with him everywhere he goes. When he learns new information, he writes it down. When he needs some old information, he looks it up. For Otto, his notebook plays the role usually played by a biological memory.

⁸ For an overview on 4E approaches to cognition see Newen et al.(2018). The more general assumption here – that knowledge supervenes exclusively on biologically grounded cognition (viz., that knowledge is 'in the head') – was only really first explicitly challenged within mainstream epistemology as recently as 2010 (e.g., Pritchard, 2010), decades after the idea had already begun holding some sway in the philosophy of mind and cognitive science. See Carter et al. (2014) for an overview of the development of cognitive extension.

⁹ An exception here is Vierkant (2022).

¹⁰ An exception here is the 'extended epistemology' subfield that has emerged in the past decade, see, e.g.Pritchard (2010); Carter et al. (2018a).

¹¹ Adams and Aizawa have, in various works, pressed back against challenges to the orthodox picture of the mind, challenging arguments for extended cognition on various fronts. As they put it "For all that the radical philosophers have said, the mind is still in the head" (2009, 78).

Clark and Chalmers' key insight is that Otto's notebook is playing the very same functional role (vis-a-vis information storage and retrieval) as Inga's biomemory is playing. It would be, as the thought goes, an objectionable 'bioprejudice' if we were to discount the process of consulting the notebook as a genuine memory process simply on account of the notebook's physical constitution or location, given that the notebook is performing the same (information storage-and-retrieval) function for Otto as biomemory is for Inga in the memory process she employs. Clark and Chalmers accordingly suggest our thinking about what kinds of things are parts of genuine cognitive processes ought to be guided by a more liberal principle – the *parity principle* – than by considerations of physical constitution or location:

Parity Principle: If, as we confront some task, a part of the world functions as a process which, were it done in the head, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world is (so we claim) part of the cognitive process. (Clark & Chalmers, 1998, p. 8).

Notice the tension between the parity principle and cognitive internalism. If cognitive internalism is true, it is simply impossible for any cognitive process to 'criss-cross' the boundaries of the brain and body and world; cognitive processes, including memory processes, are, as such, *intracranial* processes. However, if the parity principle is correct, then it is going to be entirely an open question, in any given case, whether a given part of the world is part of any given token cognitive process. And if, in some cases, a part of the world would get positively 'ruled in' as part of a cognitive process (via the parity principle), then it follows that *at least some* cognitive processes are *transcranial* and not merely *intracranial*. And that is exactly what the proponent of extended cognition maintains: not the implausible view that *all* cognitive processes are transcranial processes, but rather just that that (*contra* cognitive internalism) at least some are.

Now, against that background, let's consider how a proponent of extended cognition (suppose, on the basis of something like the parity principle) will approach the question of whether the kind of 'knowledge' that features in PKP would always be generated and stored intracranially.

To this end, let's consider the case of two tennis serves, one served simply by accidental movements (one is trying to use the tennis racket to swat a fly, and so happens to hit the ball instead), and another serve executed with purpose by a skilled player. The latter serve is an intentional action, while the former is not. According to PKP, what makes the difference here is that the latter case is one where the player *knows* that she's serving (intentionally) while doing so, but this knowledge is lacked in the former case. But notice that this knowledge of what one is doing *while one is doing it* seems to be (for our tennis player in the second case) 'immediate' propositional knowledge, knowledge that one has in virtue of having an occurrent belief *throughout* the process of serving. But then, as the thought goes, the only kind of machinery that can host *occurrent* beliefs are biological brains – *even if* (as per extended cognition) information can be stored externally, and indeed even if when stored externally (as in the case of Otto) we grant that Otto has *dispositional beliefs* stored externally even when these are not occurrent. The following thinking is suggestive of the thought that the 'knowledge' (of what ones is doing while they're doing it) that features in PKP would always be generated and stored intracranially; and, if this is right, this it looks like – if a knowledge condition like PKP is correct – the *mentality* of intentional action is exclusively *intracranial mentality*.

We maintain that the above thinking is too quick, and that this becomes apparent once we consider more complex, *diachronic* intentional actions – viz., intentional actions that require the performing of multiple steps toward a distal intended end. Compare now a case of Inga and Otto performing such an action – holding fixed the difference in memory storage from the original case pair:

INGA*: Inga* is driving from New York to Los Angeles, via Interstate 70 West. In normal traffic, the trip will take around 50 h, however, Inga* knows she will need to stop at least twice along the way to sleep, and at least five times to refuel, and get food and drinks. As Inga drives through Missouri, she stays on I-70 just as she knows she is meant to, thanks to her carefully memorizing the map.

OTTO*: Otto* is driving from New York to Los Angeles, via Interstate 70 West. In normal traffic, the trip will take around 50 h, however, Otto* knows he will need to stop at least twice along the way to sleep, and at least five times to refuel, and get food and drinks. As Otto* drives through Missouri, he stays on I-70 just as he knows he is meant to, thanks to his notebook.

Before getting into the details about OTTO*, let's think first about INGA*. Inga is not simply driving through Missouri (although she *is* doing that). Inga is doing something more complicated, and intentionally so: she's driving *from New York to Los Angeles*. On that assumption, from PKP, she must know that she's doing *that* (and intentionally) while doing it.

Assume, for *reductio*, that the only way that Inga* can know *this* (viz. that she's driving from New York to Los Angeles) while intentionally doing this by the lights of PKP, is to continually host an occurrent belief *throughout* the trip, with that content (viz., that she is driving from NY to LA). On this assumption, Inga* can't (by PKP) intentionally drive from New York to Los Angeles unless we make over-intellectualising and demanding assumptions. A proponent of PKP might suggest that the right way to read the principle isn't such that it requires occurrent belief (at each stage of the diachronic intentional action) with the content that one is doing the complex action, only that one know she is doing the *part* of the complex action that she is doing at that time, and occurrently believes (at each part of the complex action) only *that*. On such an assumption, we read "she knows that she is Φing ", in PKP, as something more akin to: "she knows that she is *M*-ing, where *M* is a means to Φing ".

However, even this way of thinking of PKP requiring occurrent belief runs into challenges. Consider Sarah Paul's (2009) case of a distracted driver:

DISTRACTED DRIVER: Consider a driver driving home on autopilot; the driver takes a certain route because it is the best way home, uses his turn signal in the appropriate places for reasons of safety, raises the garage door in order to park his car, and so forth. To express the fact that his turning and signaling are purposive, have a means-end structure, and are performed in response to his reasons, it is natural to say that he turns and signals intentionally. If his driving is habitual enough, however, he may have no belief that he is doing these things.¹²

It's hard to deny that Paul's distracted driver is driving home intentionally; yet, it's hard to see how the driver is, while driving¹³, hosting occurrent beliefs, either of the form "I am driving home", or – to the above point – even of a form whose content specifies some particular means she is now taking toward that objective.¹⁴

What this all suggests, then, is that if PKP is plausible at all as a (propositional) knowledge condition on intentional action, the knowledge that the principle maintains is needed for acting intentionally must itself be captured in a way that doesn't implicate occurrent beliefs throughout the duration of the action (regardless of the content of those beliefs). A natural move at this point is to suppose that the right way to explain Inga* as driving from NY to LA intentionally, for one who would like to hold on to PKP, is via allowing that (at some points at least) the beliefs she has about what she is doing will be *dispositional* rather than *occurrent beliefs* – viz., where dispositional beliefs are, in virtue of being stored in memory, available to the mind for endorsement even when the content of a dispositional belief is not being consciously endorsed. To be clear, this interpretation is perfectly compatible with the thought that some actions (going back to our tennis example) will be ones where performing them intentionally will paradigmatically involve occurrent beliefs about what one is doing while they are doing it. Our point is just that a plausible articulation of PKP shouldn't be read as holding that all intentional actions (e.g., including complex intentional actions like Inga*'s) involve knowledge (either that one is Φ ing or that one is taking some means M towards Φ ing) that must be *occurrent* throughout the action, or even *mostly* throughout the action.

The foregoing offers us now a good vantage point to think squarely about the case of Otto*. Let's remember that, in the original Inga/Otto contrast, Otto was really just like Inga when it comes to memory processing, with the exception being just that a key part of the process Otto uses to store and retrieve memories (i.e., his notebook) was located outside of his head. If, as we've been suggesting, the most plausible

¹² This case is slightly modified from Paul's (2009, 4–5) case.

¹³ Granted, in some cases like that of Paul's, we can expect that one will have antecedently – prior to driving, in this case – formed an occurrent belief whose content is a plan (see, e.g., Bratman, 2009). It might be that some cases of driving are exactly like this. However, for our purposes, we take it that not all such 'autopilot' cases will be ones where any occurrent beliefs feature even antecedently in the causal chain that leads to one's intentional performance. We maintain that in order to make sense of such cases of intentional action while also embracing PKP, the proponent of PKP's best move is to relax not insist on 'occurrent' belief (in a specification of the principle) but to allow dispositional beliefs as well. Thanks to a referee at Synthese for requesting clarification on this point.

¹⁴ For instance, if we take seriously the suggestion that some of our intentional actions are (like in DIS-TRACTED DRIVER) done more or less on autopilot (particularly when they are things we've done many times before), we should not expect that the distracted driver will (e.g., at around the midway point of the journey) be hosting an occurrent belief with the content that "I am now taking such and such means, M" where these are (partial) means (e.g., staying on I-70 West) toward getting to L.A.

way of unpacking PKP is as one that allows the relevant propositional knowledge one must possess when acting intentionally (by PKP) to be, at least in some circumstances, knowledge one has in virtue of *dispositional* beliefs, beliefs stored in memory, then we've now got all we need to derive a new kind of parity principle (one pertinent to the mentality of intentional action, specifically), which will follow from the combination of the original parity principle in conjunction with PKP understood as a principle that allows the kind of knowledge required by intentional action to be (at least in some circumstances) stored as dispositional beliefs. Call this new principle *Parity Principle-Extended Intentional Action* (Parity-EIA):

Parity Principle-Extended Intentional Action (Parity-EIA): If, as we confront some task, a part of the world functions as a process which, were it done in the head, we would have no hesitation in recognizing as part of the cognitive process that generates or preserves propositional knowledge required for acting intentionally, then that part of the world is part of a cognitive process capable of generating or preserving propositional knowledge required for acting intentionally.

Recall that if we follow Clark and Chalmers's original parity principle when theorizing about the bounds of cognitive processes generally, it is a short road to countenancing *extended cognition*. An analogous point applies here if we follow *Parity Principle-Extended Intentional Action* (Parity-EIA) when theorizing about the boundaries of the mentality of intentional action. From (Parity-EIA), it is a short road to *extended intentional action* – viz., action that is controlled in the way that acting intentionally requires on account of the fact that an agent knows certain facts about what she's doing when acting, but where this knowledge itself can be in principle stored outside the head.

By way of reminder, our objective here has *not* been to either argue positively for PKP as an epistemic principle on actional control (we remain agnostic on this point), *nor* to argue positively for extended cognition generally. More modestly – though we take this to be in interesting result nonetheless – we've seen that cases like Otto* put serious theoretical pressure on a would-be proponent of a PKP-style propositional knowledge condition on intentional action, in so far as one would hope to defend PKP alongside cognitive internist conception of the kind of knowledge PKP requires. Whereas the conjunction of PKP+cognitive internalism can't make sense of how Otto intentionally drives from New York to LA, the conjunction of PKP and the parity principle – what we called (Parity-EIA) – *can* make sense of this. But it makes sense of this only by, at the same time, allowing the *bounds of intentional action* to criss-cross the brain, body and world – viz., by allowing extended intentional action.

Granted, one might simply reject PKP – and attempt to explain the sense in which intentional action is controlled, by its connection to knowledge, in some other way. Any such story, we submit, should make sense of how Otto* is intentionally driving from New York to L.A. no less than Inga is doing so. Our contention to this point is simply that *if* one wants to capture actional control in terms of propositional knowledge along the lines of PKP, cases like Otto* suggest that the proponent of PKP needs to leave cognitive internalism behind.

We turn in the next section to an entirely different way that actional control has been captured in terms of knowledge – by linking intentional action not with propositional knowledge, per se, but with *knowledge-how*.

2.2 KHI + cognitive internalism

A widely held idea in action theory takes know-how to be *necessary* for intentional action, in the sense captured by KHI:¹⁵

(*Knowledge-how/Intentionality (KHI*): If S intentionally Φ s, S knows how to Φ .¹⁶

Here is a simple kind of case, due to Hawley (2003) that is useful for capturing the intuitive plausibility of KHI:

ANNOYANCE: Susie likes to annoy Joe, and Susie believes that she annoys him by smoking. Joe is annoyed not by Susie's smoking, but instead by Susie's tapping her cigarette box, which she happens to do whenever she smokes.

Susie, despite having the *ability* to annoy Joe, doesn't know-how to annoy him. As Pavese puts it, "a natural explanation of this is that she cannot annoy him *intentionally*" (2021, § 5). At most, she can reliably annoy him accidentally.

KHI also fits well with the way we attribute (and withdraw attributions of) intentional action in connection with know-how. Suppose a lawyer is attempting to prove the defendant *intentionally* hacked into the government's servers to steal secret information. If the defense lawyer replied, "but my client doesn't even *know how* to hack, or to write code", this assertion, if true, would seem to straightforwardly count against the prosecution's claim. But that would fit squarely with KHI, which holds that acting intentionally requires know-how.

KHI, we should note, is compatible with PKP; these are not mutually exclusive knowledge constraints. But let's set PKP aside now and consider the following: if KHI is correct, is the most plausible version of the thesis captured via a *cognitive internist* gloss of the relevant knowledge (know-how) required by intentional action, or is it not?

We have, potentially, a *partial* answer to this question already. Consider that one of the central approaches to the nature of knowledge-how, *intellectualism*, takes knowhow to be a species of propositional knowledge: put roughly, a subject S knows how to Φ iff for some way, w, where w is a way for S to Φ , S knows that w is a way for S to Φ .¹⁷ On the assumption that what KHI requires (under the description of knowhow) for acting intentionally is *intellectualism* know-how, KHI can be reinterpreted as holding that if S intentionally Φ s, then for some way, w, where w is a way for S

¹⁵ See Pavese (2021) for discussion.

¹⁶ Example defenders include Beddor and Pavese 2021; Ryle, 1949; Williamson and Stanley 2001; Hawley, 2003; Hornsby, 2004; Stanley, 2011; Setiya, 2012; Carter and Shepherd 2023.

¹⁷ See, e.g., Stanley (2011); Williamson and Stanley (2001); Pavese (2017); Brogaard (2011).

to Φ , S knows that w is a way for S to Φ . With this formulation in mind, consider again Otto*, who is intentionally driving from New York to Los Angeles, while storing information in extended memory in a way that is functionally on a par with how Inga* is storing her information in biomemory. If an intellectualist construal of KHI is wedded to cognitive internalism, then it looks like the intellectualist construal is not going to make sense of how Otto* counts as driving intentionally, given that (in short) the relevant action-guiding information (i.e., that such and such is a way for him to get to LA) is simply nowhere in his head. Only by embracing a more liberal kind of position (along the lines of Parity-EIA) can a specifically intellectualist reading of KHI make room for how Otto* is driving to LA no less intentionally than Inga is.

The KHI-*anti-intellectualist* combination is a bit more complex. Anti-intellectualists are unified in rejecting the idea that intelligence states (such as know-how) are grounded in propositional attitudes. As such, anti-intellectualism is a negative thesis. But anti-intellectualists also typically embrace, in positively characterizing knowhow, some version of the Rylean insight that know-how to Φ is to be identified with an agent's *dispositions* to Φ rather than with known facts.¹⁸ As Bengson and Moffett (2011) characterize this positive view: know-how is, for the anti-intellectualist, *grounded in* Φ -oriented dispositions of the agent, as opposed to being grounded in the agent's propositional attitudes.¹⁹

The matter of how to *characterize* the relevant dispositions of an agent that ground know-how states is a matter of internal dispute amongst anti-intellectualists.²⁰ But one tacit background assumption of these proposals is that the supervenience base of these dispositions is intrinsic to the biological agent. To use a simple example here, suppose you know how to play the piano; for anti-intellectualists, this will be in virtue of your having a kind of piano-related disposition. And – to make explicit the shared background assumption here – you'd have this disposition on account of how *you* are constituted – e.g., given your muscles, mental states, brain, etc.²¹

As it turns out, the KHI-*anti-intellectualist* combination runs into trouble when paired with the above kind of biological disposition supervenience assumption much as KHI-*intellectualism* ran into trouble with cases like Otto* when paired with cognitive internalism.

To bring this point into sharp relief, let's consider a trio of comparison cases:

DECORATOR-1: Amy is a skilled interior decorator who has been hired to decorate the interior of a new museum. Color-coordination is critical to doing the job well. A master at color coordination, Amy chooses tasteful colors that, when implemented in the new interior design, are very well received.

¹⁸ Though see Kremer (2017).

¹⁹ Note that this characterization is compatible with some propositional attitudes being necessary for know-how. The anti-intellectualist's position, framed as a grounding claim, is simply going to be that the know-how is in virtue of the dispositions, not in virtue of the propositional attitudes, and indeed even if such propositional attitudes are necessary *for* possessing the relevant dispositions.

²⁰ For some discussion, see Carter and Pritchard (2015); Carter and Poston (2018); Noë (2005).

²¹ Being in a *state* of know-how is accordingly being in a state where you are constituted (biologically) such that a success oriented disposition supervenes on that biological constitution.

DECORATOR-2: Neil is naturally colorblind. However, his "Eyeborg" device allows him to perceive color through sound waves. He is also a master at color coordination. He is hired to decorate the interior of a nearby museum; he chooses tasteful colors that, when implemented in the new interior design, are very well received.

DECORATOR-3: Charles is naturally colorblind. He is nonetheless hired to decorate the interior of a nearby museum; through total luck (he picks colors randomly) he chooses tasteful colors that, when implemented in the new interior design, are very well received.

Here are two uncontroversial observations. First, Amy knows how to color-coordinate. Second, Charles does *not* know-how to color-coordinate. In connection with KHI, this means Charles is not in a position to *intentionally* color-coordinate the museum, but Amy is.

But what about Neil? We base the case of Neil on the British cyborg Neil Harbisson, whose Eyeborg, osseointegrated into his skull, enables him to perceive colors under a different mode than the rest of us, but to distinguish colors from each other nonetheless.²² Neil is biologically colorblind, but he is not colorblind with his Eyeborg, which has for decades been his default way of engaging perceptually with color in his environment.



Neil Harbisson with Eyeborg: (Photo: Pedro Henrique Ferreira).²³

When it comes to color-coordinating the museum's interior as he has, can Neil do this *intentionally*? Put another way: is Neil more like Charles in DECORATOR-3 or more like Amy in DECORATOR-1?

A proponent of KHI-anti-intellectualism will need to be very careful here. Intuitively, Neil attains his excellent color-coordination results no less intentionally than

²² For discussion, see Yasenchak (2013); Pearlman (2015); and Carter et al., (2018).

²³ This explanation of the eyeborg appears originally in https://www.straitstimes.com/lifestyle/the-worlds-first-official-cyborg-10-things-to-know-about-him.

Amy does. If that's right, then KHI, for it to be plausible, must allow that Neil has the relevant know-how in DECORATOR-2, which would set him apart from Charlie (who lacks such know-how) in DECORATOR-3. *However*, if KHI-anti-intellectualism is paired with the background thesis that the dispositions in virtue of which an agent is in a state of knowing how to do something are dispositions seated (in that they have their supervenience base in) the biological agent *exclusively*²⁴, then Neil both lacks the relevant color-coordination know-how, and on account of this, is effectively in the same boat as Charles: *neither* would count as intentionally colorcoordinating the museum because neither possesses the relevant know-how that by KHI-anti-intellectualism is required for acting intentionally.

The foregoing suggests, in short, that for KHI-anti-intellectualism to be positioned to rule-in cases like DECORATOR-2 as cases of intentional action, background assumptions that limit the supervenience base of know-how apt dispositions to the

²⁴ One way to think of the kind of dispositions that matter for know-how is in terms of Sosa's (e.g., 2015) tripartite 'SSS' seat/shape/situation triad (for one version of know-how that makes use of this model, see Carter, 2022, Ch. 4; see also Andrada and Carter (Forthcoming). A referee has pointed out that a cognitive internalist might (in assessing Neil Harbisson's visual-perceptual skill), be inclined to reason as follows: the Eyeborg should be understood as part of the situation in which Harbisson's biologically supervenient skill is exercised, and where the disposition associated with that skill is itself best understood as entirely biologically supervenient. In this way, as the thought might go, we can make sense of Harbisson's visual-perceptual skill without (as we go on to suggest in the paper) recourse to transorganismic skill supervenience. We agree with the referee that Sosa's SSS structure is helpful for thinking about the kinds of dispositions that matter for know-how. Our reason for thinking that the above-described option is not a promising one for the intracranialist is as follows: it is a mark against any account of skill if it cannot explain how skills are retained by performers who are (in Sosa's terms) in poor shape and poorly situated and so not in position to exercise them. For example, Carlos Alcaraz's tennis skill is something Alcaraz retains even when he is drunk and under water, without a tennis racket; it is not shared equally by a novice in similar shape and similarly situated. If the intracranialist maintains that the supervenience base of Harbisson's visual-perceptual skill is not transorganismically supervenient, but that it includes only what is biologically endowed to Neil, it is unclear how we can explain what Neil does retain (when in poor shape and poorly situated, without the eyeborg), but which is not retained by someone (call him Neil*) who is like Harbisson in all respects except that they would be unreliable at colour identification even when situated so as to have Eyeborg access. It seems that to explain why Neil but not Neil* retains their skill when in poor shape and poorly situated, the view that Neil's skill (in the sense of the Sosa-style seat of the disposition) supervenes partly on Neil's brain, etc., and partly on the eyeborg offers an explanatory advantage here. To be clear, the intracranialist can give a coherent story here for the difference between Neil and Neil*; our view is just that viewing the supervenience base of Harbisson's visual-perceptual skill as transorganismically supervenient offers a better explanation, and by extension, a better explanation for why Neil can do certain things intentionally. Consider now some further contrast points between the two positions: if (per the intracranialist) the Eyeborg is merely a medium for inputs (and not part of the supervenience base), it seems analogous to sensory transducers like the eyes or ears. But it's not clear that an intracranialist would want to deny that such transducers can be part of the supervenience base for perceptual skills. After all, differences in sensory organs (e.g., an acute vs. impaired ear) can make a difference to auditory dispositions and skills. Second, in the case of Neil, the Eyeborg doesn't just provide inputs, but actively transforms and supplements them to make color information available. It's integral to his color-identifying abilities in a way that mere input media are not. By analogy, if a person had a cochlear implant from birth, it would be odd to deny that it was part of their auditory system and the basis of their auditory skills. That said, we grant that an intracranialist could argue for a distinction between the Eyeborg and biological sensory organs, and hold that only the latter are part of the supervenience base for skills. They would then need to maintain that Neil* is like someone whose sensory inputs aren't being accurately or reliably transmitted to an intracranially-seated recognition capacity. Whereas we want to allow that skills can supervene on an extended system, the intracranialist must insist on an intracranial locus, with performance differences explained in other ways.

biological subject should be left behind. What kind of principled thinking should replace it? There is, of course, the risk that KHI-anti-intellectualism paired to *too permissive* a conception of the supervenience base for (know-how-apt) dispositions will carry its own problem, which is that it will allow KHI-anti-intellectualism to be too easily satisfied.

A way forward here, we think, will require first a distinction between two kinds of dispositions pertinent to know-how on an anti-intellectualist programme. First, it's worth making explicit that any plausible version of anti-intellectualism will be compatible with the idea that we are not always in a position to *exercise* our know-how. A pilot, for instance, will plausibly retain their know-how (e.g., to land a commercial aircraft) when they are at home, miles from any cockpit. What *kind* of disposition, then, does the pilot retain when *at home*, as opposed to, while in a cockpit?

Following Bengson and Moffett (2011), let's distinguish being in a state of knowing how to do something, Φ , from exercising that know-how state in Φ -ing. The kind of disposition that (for a plausible anti-intellectualism) matters for *being in a state of knowing how* is best understood as a *general disposition to* Φ (successfully) across a suitably broad range of circumstances, C_{Φ} , when one tries, and regardless of whether one is (in the here and now) in C_{Φ} . For reference, call this the *skill* to Φ in C_{Φ} . In the case where Φ -ing is *landing a commercial aircraft*, C_{Φ} will include, e.g., conditions where one tries while in a working cockpit, with suitable ambient air pressure, etc. Having the *specific disposition* to Φ requires more than merely possessing the relevant skill to do so: it requires that every prerequisite is met, in the sense that, you have the skill to Φ in C_{Φ} , and moreover, you are in C_{Φ} .

If we restrict our focus to the kind of general skill disposition that KHI-anti-intellectualism will want to identify with *being a state of knowing how to do something*, we can reframe the theoretical challenge raised in this section in the following way: if the skill disposition associated (for a proponent of KHI anti-intellectualism) with being in a state of knowing how is taken to supervene exclusively on properties of a biological subject, then (problematically), KHI-anti-intellecutalism won't be in a position to treat DECORATOR-2 as a case of *intentional action*, as no general skill to color-coordinate would supervene on Neil's biological properties.

A better approach for KHI-anti-intellectualism, we submit, meets two conditions. It (i) allows for *transorganismic skill supervenience*, and does so while placing plausible restrictions on what parts of the world are candidates for being within a skill's supervenience base.

Fortunately, at this point, connecting skill supervenience with the parity principle offers a way forward on this score. Consider now the following 'parity principle' for transorganismic skill supervenience:

Parity Principle (transorganismic skill supervenience) (Parity-TSS): If a given skill X is such that it is characteristically exercised via a process that includes an external part, E, which would, by the Parity Principle, be part of a cognitive

process if E were internal, then that skill, X, is *transorganismic* – viz., it supervenes at least in part on the external part E.²⁵

Notice that, from Parity-TSS, we can make sense of how Neil's color-coordinating skill supervenes in part on the Eyeborg. Neil's (general) color discerning skill is characteristically exercised by a process that includes an external part (the Eyeborg) that by the parity principle gets ruled in as a kind of extended visual-perception faculty; and on account of this, Neil's skill supervenes in part on the Eyeborg, and not only on, e.g., his impaired biological color-vision.

To be clear, Parity-TSS is *not* a thesis about intentional action. Nor is it a thesis about know-how per se. It is a supervenience thesis about *general* dispositions or skills – of the sort that agents may retain (as they may retain *being in a state of know-how*) even when not in a position to exercise them. As such, it is an *alternative* to a different, received, or 'default' way of thinking about skill supervenience that is often taken for granted – viz., that skills supervene exclusively on biological agents.

Our suggestion is that KHI-anti-intellectualism is better off replacing one kind of background view about skill supervenience with a different and more flexible thesis (i.e., Parity-TSS); this does better by way of extensional adequacy (e.g., it accommodates DECORATOR-2). And, zooming out: the reasons for a proponent of KHI-anti-intellectualism to pair their position with Parity-TSS (rather than with a default biological skill supervenience thesis) are broadly analogous to the reasons a proponent of PKP has to pair itself not with cognitive internalism, but with (Parity-EIA). In both cases, the result is the flexibility to countenance actional control that owes not just to occurrent, biologically realized knowledge states, but to extended knowledge and skill.

3 Actional control, extended

In this section, we want to contextualize the results from § 2 in some of the wider issues that concern control conditions on intentional action, and on cognitive extension more generally. First, we want to emphasize that *knowledge conditions* (such as PKP and KHI) on intentional action are not the only way to secure the result that intentional actions must exhibit control. A different tradition in action theory (e.g., Davidson, 1970) characterizes intentional action as action that is (non-deviantly) caused by our intentions. We have argued that in so far as intentional action implicates knowledge states (of the two principal varieties we investigated in § 2) we have good cause to countenance extended intentional action.

It is worth registering at this point that we think it is far from clear that intentional action would 'extend' *only in so far* as knowledge states implicated by acting intentionally would so extend (e.g., via principles such as Parity-EIA and Parity-TSS). Consider, for instance, the position embraced by strong cognitivist views of intention (e.g., Setiya, 2004; Velleman, 1989) which assimilate intending to believing. In a bit more detail, the idea is that an intention to Φ is to be identified with a belief that one

²⁵ See Carter (forthcoming, Ch. 3) for related discussion.

is or will Φ . Note that, regardless of whether acting intentionally implicates knowledge states (either know-how or know-that), if it involves intention-cum-belief, then the question arises whether such beliefs must be occurrent when one acts intentionally, and if not (viz., if intention-cum-belief may be dispositional) whether they must be stored biologically when one acts intentionally. We mention the above because it speaks to a way, beyond what we've attempted to explore here, in which intentional action (via the mentality by which actions are intentional) might be viewed as extended.

Finally, we want to note a further sense in which the results from § 2 connect with existing debates, specifically, about extended cognition *generally*. In the first wave of literature on extended cognition (and associated discussions of the 'extended mind'), the question of whether cognitive processes can play out only in the head (or not) – taking memory as a paradigmatic example – was framed largely as a way of thinking about the bounds of *theoretical cognition* or intelligence. Memory processes, after all, when working well, deliver inputs to practical reasoning: beliefs or knowledge. Subsequent literature on *extended knowledge* reinforced the identification of cognitive extension with the extension of theoretical intelligence. *Action*, by contrast, is an *output* of practical reasoning. Actions, unlike beliefs, have a world-to-mind direction of fit.²⁶ By seeing ways in which – by the lights of some of the standard knowledge conditions on intentional action – we can envision intentional action as extended, we at the same time open a way to see both *theoretical* and *practical intelligence* as extended – and thus, extension applied to both mind-to-world and world-to-mind dimensions of intelligence.

Acknowledgements The authors would like to thank to anonymous referees at *Synthese* for helpful comments.

Funding Gloria Andrada's research is supported by the Foundation for Science and Technology (Fundação para a Ciência e para a Tecnologia) of the Portuguese Ministry of Education and Science under the project UIDB/00183/2020 and the individual contract CEECIND/00591/2021; and by the Spanish Ministry of Science and Innovation under the project PID2021-123938NB-100. J Adam 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.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/ licenses/by/4.0/.

²⁶ See Anscombe (1957); cf., Humbertson (1992), Williamson (2000, 2017), and Miracchi and Carter (2022).

References

- Andrada, G. (2022). Extending knowledge-how. Philosophical Explorations, 1-17.
- Andrada, G. & Carter, J. A. Forthcoming. Mind-Technology Problems for Know-how Anti-Intellectualism. Social Epistemology.
- Anscombe, G. (1957). Intention. Harvard University Press.
- Beddor, B. & Pavese, C. (2021). Practical Knowledge without Luminosity. *Mind*. October, https://doi. org/10.1093/mind/fzab041.
- Bengson, J., & Moffett, M. A. (2011). Two conceptions of mind and action. Knowledge How and the Philosophical Theory of Intelligence.
- Bratman, M. E. (2009). Intention, practical rationality, and self-governance. Ethics, 119(3), 411-443.
- Brogaard, B. (2011). Knowledge-How: A Unified Account. In J. Bengson & M. Moffett (Eds.), Knowing How: Essays on Knowledge, Mind, and Action (pp. 136–60). Oxford University Press.
- Campbell, L. (2018). An epistemology for practical knowledge. *Canadian Journal of Philosophy*, 48(2), 159–177. https://doi.org/10.1080/00455091.2017.1341073.
- Carter, J. Adam. Forthcoming. Digital Knowledge. Routledge.
- Carter, J. A., & Shepherd, J. (2023). Intentional action and knowledge-centred theories of control. *Philosophical Studies*, 180(3), 957–977.
- Carter, J. A., Kallestrup, J., Palermos, S. O., & Pritchard, D. (2014). Varieties of externalism. *Philosophical Issues*, 24, 63–109.
- Carter, J. A., & Pritchard, D. (2015). Knowledge-how and cognitive achievement. *Philosophy and Phenomenological Research*, 91(1), 181–199.
- Carter, J. A., & Czarnecki, B. (2016). Extended Knowledge-How. *Erkenntnis*, 81(2), 259–273. https://doi. org/10.1007/s10670-015-9738-x.
- Carter, J. A., & Poston, T. (2018). A critical introduction to Knowledge-How. Bloomsbury Publishing.
- Carter, J., Adam, A., Clark, O., & Palermos, S. (2018a). New Humans? *Ethics, Trust, and the Extended Mind*, 331–51.
- Carter, J., Adam, A., Clark, J., Kallestrup, S., Orestis, P., & Duncan, P. 2018. *Extended epistemology*. Oxford University Press.
- Clark, A., & Chalmers, D. (1998). The extended mind. Analysis, 58(1), 7–19. https://www.jstor.org/ stable/3328150.
- Davidson, D. (1970). Essays on actions and events: Philosophical essays volume 1. Clarendon.
- Davidson, D. (1978). Intending. Philosophy of history and action: Papers presented at the First Jerusalem Philosophical Encounter December 1974 (pp. 41–60). Springer.
- Hawley, K. (2003). Success and Knowledge-How. American Philosophical Quarterly, 40(1), 19–31. https://www.jstor.org/stable/20010094.
- Hornsby, J. (2004). Agency and actions. Royal Institute of Philosophy Supplements, 55, 1-23.
- Humberstone, I. L. (1992). Direction of fit. Mind, 101(401), 59-83.
- Kremer, M. (2017). A capacity to get things right: Gilbert Ryle on Knowledge. European Journal of Philosophy, 25(1), 25–46. https://doi.org/10.1111/ejop.12150.
- Miracchi, L., & Carter, J. A. (2022). Refitting the mirrors: On structural analogies in Epistemology and Action Theory. Synthese, 200, 9.
- Newen, A., De Bruin, L., & Gallagher, S. (2018). The Oxford handbook of 4E cognition. Oxford University Press.
- Noë, A. (2005). Against Intellectualism. Analysis, 65(4), 278-290.
- Paul, S. K. (2009). How we know what we're doing. Philosophers' Imprint, 9, 1-24.
- Pavese, C. (2015). Practical senses. Philosophers' Imprint 15.
- Pavese, C. (2017). Know-How and Gradability. The Philosophical Review, 126(3), 345-383.
- Pavese, C. (2022). Practical knowledge first. *Synthese*, 200(5), 375. https://doi.org/10.1007/ s11229-022-03848-y.
- Pavese, C. n.d. Knowledge-How. In Edward N. Zalta (Ed.), Stanford Encyclopedia of Philosophy (Summer 2021 Edition).
- Pearlman, E. (2015). I, Cyborg. PAJ: A Journal of Performance and Art, 37(2), 84-90.
- Pritchard, D. (2010). Cognitive Ability and the Extended Cognition Thesis. Synthese, 175(1), 133–51. https://doi.org/10.1007/s11229-010-9738-y.
- Ryle, G. (1949). The Concept of mind. Hutchinson & Co.

- Schwenkler, J. (2012). Non-observational knowledge of action. *Philosophy Compass*, 7(10), 731–740. https://doi.org/10.1111/j.1747-9991.2012.00513.x.
- Setiya, K. (2004). Explaining action. The Philosophical Review, 112(3), 339-393.
- Setiya, K. (2009). Practical knowledge revisited. Ethics, 120(1), 128-137. https://doi.org/10.1086/606000.
- Setiya, K. (2012). XIV—Knowing How. In Proceedings of the Aristotelian Society (Hardback) (Vol. 112, pp. 285–307). Wiley Online Library.
- Shepherd, J. and J. Carter, A. Forthcoming. Knowledge, practical knowledge, and intentional action. *Ergo:* An Open Access Journal of Philosophy.
- Stanley, J. (2011). Know How. OUP Oxford.
- Velleman, D. (1989). Practical Reflection. https://philpapers.org/rec/VELPR.
- Vierkant, T. (2022). The tinkering mind: Agency, Cognition, and the extended mind. Oxford University Press.
- Williamson, T. (2000). Knowledge and its limits. Oxford University Press.
- Williamson, T., & Stanley, J. (2001). Knowing how. The Journal of Philosophy, 98(8). https://ora.ox.ac. uk/objects/uuid:1fcf732b-731b-47e6-943d-aaf34905c02a.
- Williamson, T. (2017). Acting on knowledge. In J. A. Carter, E. C. Gordon, & B. Jarvis (Eds.), Knowledge First: Approaches in epistemology and mind (pp. 163–181). Oxford University Press.
- Yasenchak, D. (2013). Filling the Colorless Void: The Cybernetic Synaesthesia of Neil Harbisson. *Retrieved Oct*, 25, 2015.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.