

Intentional action and knowledge-centred theories of control

J Adam Carter and Joshua Shepherd

Abstract: Intentional action is, in some sense, non-accidental, and one common way action theorists have attempted to explain this is with reference to control. The idea, in short, is that intentional action implicates control, and control precludes accidentality. But in virtue of what, exactly, would exercising control over an action suffice to make it non-accidental in whatever sense is required for the action to be intentional? One interesting and *prima facie* plausible idea that we wish to explore in this paper is that control is non-accidental in virtue of requiring *knowledge* – either knowledge-that or knowledge-how (e.g., Beddor and Pavese 2021; cf., Setiya 2008; 2012 and Habgood-Coote 2018). We review in detail some key recent work defending such knowledge-centric theories of control, and we show that none of these accounts holds water. We conclude with some discussion about how control opposes the sort of luck intentional action excludes without doing so by requiring knowledge (that- or how).

1. Introduction

A widely-shared platitude about intentional action is that it must be, in some sense, non-accidental. With a golf club in hand, even good players don't intentionally make a hole in one, and this is so even when they desire to do so, hit a ball, and it happens to go in.

Spelling out the sense in which intentional action must be non-accidental is famously contentious. A common criticism of causalist views is that they have no account of the sense in which the causation of action is non-deviant (di Nucci 2013, Williamson 2017). But a similar issue confronts Anscombian views of action, which stress the role of practical knowledge in rendering an action intentional (Paul 2011, Small 2012). One reason is that, as Frost notes, 'It's plausible that knowledge of any kind involves non-accidental agreement of the mind of the knower with the object of knowledge (2019, 317)' (see also Small 2012, 173).

A plausible and commonly taken move at this point is for action theorists to seek to explain the non-accidentality of intentional action in terms of control, or in terms of nearby notions such as skill (Mele and Moser 1994, Aguilar 2012, Dickie 2012, Beddor and Pavese 2021, Shepherd 2021). Control, the thought goes, is opposed to accidentality. The big question that this move raises, then, is this: just what is control, such that it could explain non-accidentality?

One interesting thought, that we wish to explore in this paper, seeks to explain control in terms of knowledge. Perhaps we can explain control, and thus the non-accidentality of intentional action, by appeal to some sort of involvement of knowledge in action.

In this paper we argue that this cannot be done. Along the way, we argue that recent knowledge-centred accounts of control (or skill) fail. Satisfactorily countenancing the tight relationship between intentional action and control doesn't require any appeal to knowledge.

Here is the plan. We begin (§2) by challenging a notable recent attempt to defend a knowledge-centred theory of a control condition on intentional action due to Bob Beddor and Carlotta Pavese (2021) – and in doing so, we focus on problems for thinking that such control requires propositional knowledge – specifically, knowledge (at t) that one is ϕ -ing at t . We then, in §3 consider – and drawing on some recent ideas about intention and know-how due to Josh Habgood-Coote (2018) and Kieran Setiya (2008; 2012) – consider a variation of the knowledge account which places exclusively a know-how rather than a standard know-that condition on control necessary for intentional action. After rejecting both proposals, we conclude with some general discussion about how control opposes the sort of luck intentional action excludes without doing so by requiring knowledge (that- or how).

2. Beddor and Pavese on intentional action, control, and knowledge

In a recent article, Bob Beddor and Carlotta Pavese (2021) have attempted to salvage an epistemic condition on intentional action in the face of Juan S. Piñeros Glasscock's (2020) anti-luminosity argument. Their core strategy has been to give up the thesis that intentional action requires knowing that one is *acting intentionally* and to retreat to a more moderate epistemic

condition on intentional action, according to which intentional action requires merely knowing what one is *doing*.

Beddor and Pavese motivate their preferred epistemic condition on intentional action – a knowledge condition – by considering the connections between intentional action and control. Before challenging Beddor and Pavese, we want to make clear where we think they are right.

Their critical focus is Piñeros Glasscock's strategy of rejecting the following Practical Knowledge Principle:

Practical Knowledge Principle (PKP): Necessarily, if an agent is ϕ -ing (intentionally and under that description), they know that they are ϕ -ing (intentionally and under that description).

Piñeros Glasscock's tack is to argue that an anti-luminosity argument, of the sort developed in Williamson (2002, Ch. 4; see also Srinivasan and Hawthorne 2013), can be readily adapted to provide a reductio of an epistemic condition on intentional action (with PKP being an especially popular way to capture such a condition).¹ If, as the thought goes, Williamson is right there can be no non-trivial luminous conditions whatsoever, then an epistemic condition on intentional action such as PKP must be rejected because it implies otherwise.

Assuming for the sake of argument that Williamson's anti-luminosity argument is correct (we won't dispute this here²), Piñeros Glasscock is surely right that PKP must go. However, the importance of Beddor and Pavese's intervention lies in the fact that they've shown, rightly we think, that from the fact that PKP violates anti-luminosity we should not on this basis (as Piñeros Glasscock encourages us to) jettison an epistemic condition on intentional action. And, indeed, Beddor and Pavese formulate and defend one such condition – framed in terms of knowledge – which they think we should accept, as follows:

¹ Notable defences of PKP include Anscombe (1957), Hampshire (1983), Newstead (2006), and Rödl (2011), among others.

² Though, for criticism, see, e.g., Brueckner and Fiocco (2002) and Ramachandran (2009).

Epistemic Condition (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 .³

As Beddor and Pavese see it, we should embrace EC because we should accept two theses that entail it: what they call **Control Constraint**, which is meant to capture the intuitively tight connection between ϕ -ing intentionally and being in control of one's ϕ -ing, and the **Epistemic Theory of Control**, which they offer as a plausible characterisation of how one who is in control of their ϕ -ing must be related, epistemically, to their ϕ -ing.⁴

Control Constraint (CC): Whenever an agent ϕ s intentionally, they are in control of their ϕ -ing.

Epistemic Theory of Control (ETC): Someone is in control of their ϕ -ing at t if and only if they know they are ϕ -ing at t , and they know this in virtue of exercising their knowledge of how to ϕ .⁵

While we think that Beddor and Pavese are right to accept the control constraint, we have two worries about ETC, which is essential to getting them to their preferred EC. The first is that, on a natural interpretation of the theory, ETC is circular. The second is that the clause requiring knowledge of what one is doing is undermined by a series of cases.

2.1 *Circularity*

To see the circularity worry, consider the place of knowledge-how in ETC. The agent must be exercising their knowledge how. But what is it to exercise knowledge how? Different theorists

³ (Beddor and Pavese 2021, 6).

⁴ The final key step in Beddor and Pavese's wider argument — one that will not be our main focus here — is to sketch how the threat of a version of Piñeros Glasscock's anti-luminosity reductio resurfaces (despite initial appearances to the contrary) for even their preferred weaker epistemic condition (EC), at least, if one accepts the widely shared assumption that some actions are essentially intentional (see, e.g., Anscombe (1957)). However, as Beddor and Pavese see it, when forced between the choice of accepting (i) EC (which they take to be motivated by the ex ante plausible CC and the ETC); and (ii) countenancing the existence of essentially intentional actions, we should go with (i), especially, as they suggest, because we have independent reasons to reject essentially intentional actions.

⁵ (Beddor and Pavese 2021, 4-6).

of knowledge-how may offer different answers.⁶ But an initial and natural thought is that the exercise of knowledge-how at least involves the exercise of a well-crafted, or well-enough-crafted, *intention*, where the contents of one's intention can be thought to bear relationships to the things that one knows how to do. In a separate forthcoming paper, Pavese offers an interpretation in line with this thought. She writes,

It seems that in order for agentic control to obtain, in addition to knowing what one is doing, one ought to know this in virtue of exercising their intention to perform the action and in virtue of being guided by one's knowledge-how (Pavese 2021b).

So exercising one's intention is necessary for agentic control. What is it to exercise one's intention? It is not enough that one's intention cause intended movements (or outcomes), since one's intention might do so deviantly.⁷ So it seems the best way to think of the exercise of knowledge-how involves reference to the non-deviant exercise of an intention. A difficulty for the ETC proponent here is that by far the most natural way to elucidate the non-deviant execution of intention is in terms of control.

Consider, for example, Bishop's account of control – Bishop holds that the agent exercises control iff 'the causal link from basic intention to matching behavior is sensitive,' where sensitivity specifies that 'over a sufficiently wide range of differences, had the agent's intention differed in content, the resulting behavior would have differed correspondingly' (1989, 150). Control, for Bishop, is the execution of an intention when the right condition – namely, a relation of counterfactual sensitivity between intention and behavior – is in place.

Or consider Shepherd's account of control, on which (roughly) the exercise of control involves the production of behavior that sufficiently matches a plan-state (such as an intention), in circumstances for which the causal pathway the behavior follows is a reliable route to success (2021, 46). Control, for Shepherd, is the non-deviant execution of a plan-state (such as an

⁶ For discussion, see, e.g., Carter and Poston (2018) and Bengson and Moffett (2011).

⁷ For an example case here, see the case Beddor and Pavese discuss, Inadvertent Crash, due originally to Adams and Mele (1989, 519).

intention) when the right conditions – in particular, reliable causal pathways spelled out in terms of normalcy across ‘well-selected’ sets of circumstances (2021, 46) – are in place.

Finally, consider Valaris’s account of control, on which ‘if you are phi-ing with control, then you are phi-ing not just in the actual case, but also in all relevant counterfactual cases in which you attempt to phi’ (2021, 7). Valaris does not speak of intentions here – only of phi-ing as a process of acting that involves ‘the exercise of the agent’s relevant capacities’ (9). But it is natural to think that talk of the exercise of the agent’s relevant capacities is consistent with thinking that, when those capacities are being exercised in acting, the agent is thereby exercising some relevant intention. If that is right, then for Valaris, control is the exercise of intention when the right condition – here, a modal specification of success across all relevant counterfactual cases – obtains.

We need not endorse any of these accounts of control to make the present point. The point is that the most natural way to elucidate the non-deviant execution of intention is in terms of control. But this way is not open to Beddor and Pavese, who need the execution of intention to amount to less than control, on pain of rendering ETC circular.

What other options are available to the ETC proponent? Here is one. Suppose one holds that the non-deviant execution of intention is to be understood in terms of the execution of a knowledge-involving ability – an ability to know. Frost, for example, argues that we should think of practical knowledge as a capacity that is manifested in action. Practical knowledge is ‘at once a *capacity* to do something rational (i.e. something that exhibits a particular order of practical reasoning) and also an *exercise* of a capacity to know particular actions’ (Frost 2019, 326).

Frost’s proposal is interesting, and might help someone who wants to understand the execution of *action* in terms of knowledge. But action and control are separate (even if closely linked) subjects – especially for theorists who wish to explain action in terms of control. It would take extra work to say how the exercise of a capacity to know particular actions relates to an exercise of control, to explore the implications for a theory of control, and to explore implications for the control-action relationship. It is not obvious to us that the ETC proponent should want to take on these tasks. But if they do, it is worth noting that this kind of account

runs into problem cases. Shepherd and Carter (forthcoming) discuss the cases we have in mind in a discussion of knowledge and intentional action. At the end of section 2.2, we discuss these kinds of cases in the context of control.

A second option might be found in Sosa's (e.g., 2015, Ch. 1) theory of performance normativity, which uses the notion of 'manifestation' to block causal deviance.⁸ For Sosa, an agent acts intentionally – and an intention is exercised non-deviantly – when the agent's success stems, in the right way, from her competence, where her competence is fleshed out in terms of her having the right dispositions (Skill), her being in the right physical and mental condition (Shape), and her being in circumstances that present no obvious impediments to normal action execution (Situation).⁹ What explains the difference between success stemming from competence in the right way from success stemming from competence in the wrong way is that, in the good case, the agent's success manifests her competence, and manifestation is a primitive relation between success and competence.

Now, the place at which Sosa appeals to manifestation is close to the place at which others appeal to control. So one idea is to take Sosa's manifestation of competence as a virtue-theoretic explication of control: the exercise of control and the manifestation of competence come to the same thing. Of course, this reading will not help the ETC proponent. So let us assume that the manifestation of competence is not an explication of control. Well, now we run into a question about the Control Constraint. If all we need to explain intentional action is competence manifestation, what explains the Control Constraint? The plausible option, here, seems to be this. Competence manifestation explains the presence of control: we might say that someone is in control of their ϕ -ing in virtue of the fact that their ϕ -ing manifests competence to ϕ .

⁸ Thanks to an anonymous referee for encouraging discussion of this line of thinking.

⁹ For a fuller discussion here, see Sosa (2015, 96), which distinguishes between innermost, inner, and complete competence.

We then need to say that the manifestation of competence is closely linked to knowledge – viz., that someone manifests competence in ϕ -ing at t if and only if they know they are ϕ -ing at t , and they know this in virtue of exercising their knowledge of how to ϕ . This expresses a knowledge-centred account of manifestation, and affords a manifestation-based account of control.

One worry is that this package of claims seems to make control explanatorily otiose. But it fits the technical requirements of ETC. So, although it seems more plausible to us to understand Sosa's account as an explication of control, rather than an explanation of control as some further property of action, the reading we have offered does provide a non-circular and ETC-friendly way of understanding the exercise of intention.

However, just as with the Frost-inspired proposal, this proposal runs into trouble with certain kinds of cases, considered below.

2.2. *Problem cases*

Let us now set aside the matter of whether Beddor and Pavese's ETC is subject to a circularity charge. For there is a separate problem with ETC, which is that it is too demanding; we turn now to casting doubt on whether one who is control of their ϕ -ing at t must, as per ETC, know *that* they are ϕ -ing at t .

First, though, a concession. Beddor and Pavese are careful to motivate ETC, and their strategy for doing so is at least initially plausible (even though, as we will soon suggest, misguided). Consider their case DEER HUNTING.

DEER HUNTING. Artemis, the goddess of the hunt, spies what appears to be a deer. She expertly notches her bow, aims, and releases her arrow. It turns out that she was aiming at a lifesize papier mâché statue of a deer — one that had been constructed so artfully that it would fool even the most discerning eye. But, as luck would have it, a

deer was standing directly behind the cervine statue. Artemis' arrow passes directly through the papier mâché sculpture and into the hapless deer (2021, 7).

Here are two observations about DEER HUNTING. First, Artemis' shooting the deer is not under her control (even if many of her subsidiary actions were). Second, though we may assume Artemis has a justified, true, belief that she is shooting a deer, she does not know she is shooting a deer, as her justified, true, belief is Gettiered.¹⁰ What Beddor and Pavese conclude from these two observations – both of which are correct – is that Artemis' action is not under her control because she does not know she is shooting a deer – which is precisely the diagnosis that would be implied by ETC. But the conclusion here is too quick.¹¹

Consider now another case, due to Ernest Sosa (2010).¹²

SIMONE: Simone, a pilot in training who could easily be, not in a real cockpit, but in a simulation, with no tell-tale signs ... trainees are strapped down asleep in their cockpits, and only then awakened. Let us suppose Simone to be in a real cockpit, flying a real plane, and shooting targets accurately (2010, 468).

¹⁰ The subject's belief in this case is Gettiered in a way that is structurally analogous, as they point out, to the way that (would-be) knowledge is Gettiered in Roderick Chisholm's well-known (1977) 'sheep in the field' case.

¹¹ It is worth noting that Pavese and Beddor (forthcoming), in more recent work, offer additional cases beyond DEER HUNTER, such as their case OCCUPATIONAL HAZARD (forthcoming, p. 3) in support of ETC. We take it that OCCUPATIONAL HAZARD is broadly analogous to their original DEER HUNTER case, in that it is a case that features a target belief undermined by intervening epistemic luck (and thus falling short of knowledge), and where control is intuitively lacking. The point we raise in what follows, however, is that even though DEER HUNTER (and cases with an intervening epistemic luck structure, such as OCCUPATIONAL HAZARD) are *compatible* with what ETC implies, other epistemic luck cases, which feature environmental rather than intervening epistemic luck (see Pritchard 2005), are more problematic and, as we suggest, ultimately force a dilemma.

¹² We want to emphasise that the structure of this case is importantly different from the structure of the fist-clenching case given by Setiya (2008, 2012), and which Beddor and Pavese discuss in fn. 3. Setiya's case is not a case of environmental luck.

When Simone deftly hits targets successfully, flying in a real cockpit, her shots surely are well controlled¹³ – much more so, we may suppose, than novices who were in the same position as Simone but who, upon being awakened, shoot at the targets haphazardly, missing most. This is not to say that Simone intentionally hits the real targets. We can remain neutral regarding intentional action at this point, focusing only on control. Now, just as there is an epistemic luck-style argument that Artemis fails to know that she is shooting a deer in DEER HUNTER, there is likewise an epistemic luck-style argument to the conclusion that Simone doesn't know that she is shooting targets, and despite the impressive control she exercises. The reason is that her beliefs – viz., her intellectual shots – while in the simulator are, as Ernest Sosa has pointed out, 'correct only through a kind of accidental luck seemingly inimical to true knowledge'.¹⁴ This particular kind of epistemic luck is described by Duncan Pritchard as 'environmental epistemic luck' – viz., luck of the sort one encounters in barn-façade cases (e.g., Goldman 1979; Ginet 1975), where the subject's belief is unsafe in that very easily it could be mistaken given how it was formed, but where the fact that one's belief is unsafe is simply due to one's being situated in a risky environment.¹⁵

¹³ After all, consider that Simone is manifesting the same inner skill (e.g., which supervenes on her brain and body) to hit targets that she carefully hones via the simulator. One might object to this suggestion and hold that Simone's skill (by which she controls her shots) is indexed specifically to the simulator environment. This suggestion, however, does not line up well with recent demonstrated success in using virtual reality simulators to train fighter pilots. For just one example, Mass Virtual, which supplies VR pilot training to the US Navy, reports reliable success in transitioning from virtual to non-virtual environments. For discussion, see Brooks (2008). See also Carter and McDonnell (2018).

¹⁴ Note that Sosa's own view of knowledge is idiosyncratic, in that he distinguishes between levels of knowledge, animal and reflective. The kind of knowledge Sosa takes to be lacking here is reflective knowledge (as well as, more recently, what he calls 'knowledge full well' (2010) and later 'judgmental knowledge' (2015). Judgmental knowledge and knowledge full well, like reflective knowledge, require that one assess knowledgeably that one's (first-order) belief would not too easily have been such that its correctness would not manifest competence. This point, however, won't concern us here.

¹⁵ It is worth clarifying further how we are characterizing the action we take to be under Simone's control. Define a "potential target" as a target that is either genuine or simulated. Define a "real target" as a genuine (non-simulated) target. One might question whether the action in question we take to be under Simone's is *shooting potential targets* or *shooting real targets*. Importantly, while the former is uncontentionally under her control, we maintain the latter plausibly is. A critic might worry that this assessment is less clear when we draw attention to the shots' modal fragility (viz., that she could have been shooting at simulations were she to have woken up in the simulator as she easily could have been). The 'macarena' case and our fist-clenching

To sharpen the above point, just consider how we could very easily run an environmental-luck twist on Beddor and Pavese's DEER HUNTING case – and in doing so get the result that (as in SIMONE) one is in control of one's shots while still failing to know – given the presence of environmental epistemic luck – that one's shots are successful. To run this twist, just imagine that Artemis sees 10 deer in the field, and chooses one of them. Drawing the bow back skilfully (thanks to, we suppose, years of extensive training plus raw talent) Artemis aims, releases, and hits, as she reliably does from this distance, with low winds, and perfect lighting conditions. The twist in the tale, though, is that every one of the other deer in the field are surrounded by invisible forcefields that would have deflected her arrow, had she happened to aim at any of those. Artemis, as the thought goes, does not know that she is shooting a deer as she releases the arrow, any more than Simone knows she is firing at a real target – or for that matter, any more than one knows one is looking at a real barn when that barn is surrounded by undetectable facades. Nonetheless, we want to suggest, Artemis's shot does exhibit control in this revised case, and this is so even if control is not manifested in the original DEER HUNTING CASE. The lesson, it seems, is that regardless of whether exercising control of one's shot is incompatible with the kind of knowledge-undermining epistemic luck that features in Beddor and Pavese's DEER HUNTING case, it is very plausibly compatible with the kind of purely environmental knowledge-undermining epistemic luck we find in both SIMONE and in our revised environmental-luck twist on DEER HUNTING.¹⁶

case (a tweak on Setiya's case), which we discuss further in this section, offer the reader additional support for thinking that mere modal fragility is not enough to undermine one's control even when we characterise the relevant action as involving 'real' (and not merely potential) success. Thanks to an anonymous reviewer for suggesting clarification on this point.

¹⁶ Note that the idea relied on here is just that some cases of environmental epistemic luck look to be compatible with the exercise of control over a successful action, even if we grant that intervening luck (of the sort present in the original DEER HUNTER case) is incompatible with exercising such control. Although we are sympathetic to the thought that environmental epistemic luck as such isn't enough to wreck one's exercising control over a successful action, we don't intend to rely on this stronger type-level idea and so for present purposes remain neutral on the point.

What this all means for Beddor and Pavese is that it becomes much less plausible to think that ETC gains support from DEER HUNTING. After all, we've seen that even if we grant that (in the original case) Artemis' shot is not under her control, and that she does not know that she is shooting a deer, it would be too quick to reason to the kind of diagnosis of the case that is implied by ETC, namely, that Artemis' action is not under her control because she does not know she is shooting a deer.

In response to the above point, Beddor and Pavese seem to face a dilemma, either horn of which has theoretical costs. On the one hand, they could insist (contrary to what we've suggested) that Artemis (in our environmental luck variation of the case) doesn't exercise control, and by extension, neither does Simone (in the real cockpit). Or, they could deny that environmental epistemic luck – of the sort that features in traditional barn-façade cases (as well as in our twist on DEER HUNTING and in SIMONE) – is incompatible with knowledge.

Neither route looks very promising. If Beddor and Pavese opt for the first horn, this will involve putting a strong safety requirement on control, one with reference to which we could explain how the modal proximity of Simone's waking up in the simulator prevents her from controlling her shots when she does not in fact wake up in the simulator, but in the real plane. Such a strong safety requirement on control would generate arguably even less plausible predictions elsewhere. For example, let us suppose – to take a mundane case – a dancer is beginning to dance the macarena, having done so previously hundreds of times on nights out. Unbeknownst to the dancer, on the other side of the world, a maniac is deciding by a flip of a coin whether to destroy the world with a bomb.¹⁷ Heads = destroy. Tails = don't destroy. The coin lands tails, suppose, right while the dancer is mid-macarena, moving her body according to the macarena steps she's so accustomed to, and which are now almost second-nature. Here is the point: if Simone is not in control of her shots simply due to the modal proximity that she would be firing at fake targets in the simulator scenario, then neither does our macarena dancer exercise control in her dance, simply because of the modal proximity of the bomb.

¹⁷ For a variation on this case put to a different purpose – regarding background conditions in epistemology – see Sosa (2021).

One might complain that the modal proximity of the bomb is not relevant to control, since the bomb would not inhibit control so much as destroy the agent. But other cases are available. Consider a variant of Setiya's (2008) case in which a patient's hand has been paralyzed and deafferented. By noon, the paralytic has worn off, but the patient is in an unsafe condition in that in most cases, the paralytic effects last much longer. Suppose that the patient has unjustified but true beliefs that the paralysation will wear off at noon, that they will be able to clench their fist at noon, and that this is what they will do. The beliefs are unjustified because the patient's information is bad. Nonetheless, at noon the patient, believing that they can clench and that this is what they are doing, clenches their fist. A strong safety condition on control predicts that the patient does not exercise control in doing so. But it seems clear that they do.¹⁸

Is taking the second horn of the dilemma a better option for Beddor and Pavese? Problems lurk unavoidably in this direction as well. First, and foremost, going the second route involves denying that propositional knowledge must be safe, where a belief is safe in the sense knowledge demands just in case in most nearby possible worlds in which the subject continues to form her belief about the target proposition in the same way as in the actual world, the belief continues to be true (Pritchard 2007).¹⁹ While the claim that propositional knowledge requires safety is widely endorsed in epistemology – Pritchard (2012) refers to it as a 'platitude' in most epistemologists thinking about knowledge – it has nonetheless been found to be in tension with patterns of folk knowledge attributions, as reported in empirical studies by (among others) Colaco, Buckwalter, Stich & Machery (2014) and Turri, Buckwalter & Blouw

¹⁸ Notice that the patient also lacks knowledge of what they are doing. Their belief about what they are doing is unjustified. Since they are deafferented, they get no proprioceptive feedback about whether they are doing so. And we can assume they are not watching their hand. Of this sort of case, Beddor and Pavese suggest that given the lack of proprioceptive feedback due to deafferentation, 'it becomes less clear that the agent's clenching is intentional' (2021, fn. 3) (fn. 3, forthcoming). This is tough to accept. Deafferentation seems to do no real work here regarding control or intentionality: agents who have been permanently deafferented are able to clench their fists, and there is no great reason to deny that they can do so with control, or intentionally. Similarly, this patient's clenching seems clearly to be an intentional action.

¹⁹ For further discussion of safety conditions in epistemology, of which Pritchard's formulation is a very standard one, see Rabinowitz (2011).

(2015).²⁰ Further, some epistemologists (e.g., Hetherington 2013; Sosa 2007) have denied the incompatibility of knowledge with safety on theoretical grounds.²¹

Let's assume for the sake of argument that Beddor and Pavese opt to take the second horn of the dilemma and maintain that even if (as we grant them) knowledge is not present in DEER HUNTING, it is present in cases where the unsafety of the target belief is down to purely environmental luck, as in SIMONE. Even if this much is granted – we want to now suggest – there remain safety-based problems for ETC which don't turn at all on the matter of whether specifically environmental epistemic luck is compatible or not with knowledge.

All we need to get this second kind of safety-based argument against ETC up and running are cases where an elite performer's success rate grades out somewhere in the middle percentiles.²² Take, for example, Megan Rapinoe's shot-on-target percentage (excluding penalty kicks), which is (at 48%) much higher than most players. Or take golfer Justin Rose, who is leading the PGA tour this season in putting from 10-15 feet – making 43% of all attempts from this range. Rapinoe and Rose are not merely skilled at shots on goal and putting, respectively, but they by esteem and acclaim of their peers are among the most skilled in the

²⁰ In these studies, knowledge is attributed more than not in barn façade cases, by participants without self-reported expertise in epistemology. Results from Horvath & Wiegmann (2016) and Carter, Pritchard and Shepherd (2019), however, indicate that participants with self-reported training in epistemology are inclined to deny knowledge in fake barn cases much more so than participants without self-reported training in epistemology. For discussion of the significance of such results, see Carter and Sosa (2022).

²¹ In Hetherington's case, the compatibility of knowledge and luck (of the sort a safety condition purports to block) is defended on two kinds of bases. First, Hetherington thinks safety is too strong, for reasons discussed in his (2013). But, perhaps more interestingly, Hetherington (e.g., 2011) has long defended a heterodox view about the gradability of propositional knowledge, according to which propositional knowledge is gradable in such a way that *some* (low-grades) of propositional knowledge are compatible to a greater extent with epistemic luck than higher grades. It's worth registering two points about this gradability view. First, in so far as the view implicates that knowledge is present in intervening epistemic luck cases, the view seems to run counter to a point that both we as well as Beddor and Pavese accept, which is that intervening epistemic luck is incompatible with exercising control. Secondly, note that – as we make explicit later in this section – our wider argument against ETC's propositional knowledge commitment (*viz.*, that exercising control requires propositional knowledge in the way described) does not ultimately rely on our contention that (some) cases of environmental epistemic luck are compatible with exercising control. See, in particular, our discussion of Rapinoe/Rose-style cases later in this section.

²² Shepherd and Carter (forthcoming) offer a discussion of these kinds of cases as they apply to the relationship between knowledge and intentional action.

world. We think it is accordingly very plausible to attribute a significant level of control to Rapinoe and Rose, on pain of denying that these activities are activities that admit of any control. That said, on effectively any mainstream theory of propositional knowledge – and crucially, not just those that demand a safety condition that would rule-out cases of environmental epistemic luck – it is false that Rapinoe and Rose know that they are making a successful shot on goal or sinking a putt (in the 10-to-15-foot-range, respectively). After all, they miss more than they make.²³ The lesson from Rapinoe and Rose-style cases seems to be that exercising control and knowing that one is succeeding at what one exercises one's control over have different levels of permissiveness regarding failure in similar circumstances. There's no clear way to accommodate this point in a view that holds (as ETC does) that someone is in control of their ϕ -ing at t if and only if they know they are ϕ -ing at t .²⁴

One might respond that a weaker principle is available. Following a suggestion Beddor and Pavese (2021, footnote 2) make, one might propose the following.²⁵

ETC (Weakened): If S is in control of their ϕ -ing at t , then there is some action ψ that they are performing at t in order to ϕ such that i) they know they are ψ -ing at t , ii) they know this in virtue of exercising their knowledge of how to ϕ .

Set aside the (significant) worry that the point about circularity arises as well for this weakened version of ETC. The claim here is that the agent knows that they are taking means ψ that will (with sufficient likelihood, we can say (see Pavese and Beddor forthcoming)) lead to success at

²³ Among other things, one might point out that there are relevant alternatives neither can rule out. Or, as it might be pressed, reliability past a certain threshold is required for knowing, and where this threshold (whatever level it is set – e.g., 90%) is higher than it is plausible to suppose the reliability threshold is for exercising control more generally. For further discussion on this point, see [Omitted].

²⁴ Note that, as Shepherd and Carter (forthcoming) argue, these cases make trouble for Frost's (2019) view that the control-involving execution of action is a capacity to know particular actions. The basic reason is that, while Rapinoe often intentionally makes penalty kicks, and Rose often intentionally makes 12-foot putts, their success is too risky to attribute knowledge of these actions. So it makes little sense to think of Rapinoe's success as an exercise of a capacity to know. A similar point applies to the Sosa-inspired account considered in section 2.1. For Rapinoe and Rose can manifest a competence to make penalty kicks or 12-foot putts, even though their success is too risky to attribute knowledge.

²⁵ Thanks to a referee for pressing this point.

ϕ -ing, and the explanation of their knowledge that they are ψ -ing is that they have knowledge of the means to ϕ , where ψ -ing is one such means.

A deep problem for this claim is that it is possible to iterate Rapinoe-Rose cases, such that the means the agent takes also involve significant risk. Rapinoe may take means to making a free kick that involve putting a certain spin on the ball, and she may only have a 65% chance of getting the spin right. So Rapinoe does not know that she is ψ -ing.

The ETC proponent might respond by claiming that there is some more basic thing that the agent is doing, and they know that they are doing the more basic thing. Rapinoe knows, let us say, that she is kicking the ball. But there is no reason to think that all cases of controlled behavior involve basic components that involve no risk at all. Many means we take – including basic actions – may very well have middle-grade chances of success. Rapinoe may have practiced her kick so much that it has become, for her, a basic action.²⁶ But it is entirely possible that this basic action misfires some of the time – enough of the time that attributing knowledge that she is kicking the ball is ruled out.

And yet, in such cases, it seems plausible to think that agents exercise at least some control when they take these means. Consider a case in which an agent can raise their chances of getting something right (from 40% to 60%, say) if they try harder, or focus their attention in a certain way. In such a case it seems plausible to say that when they focus, they exercise more control than when they do not focus. But, because of the risk involved in the behavior, it does not look like the taking of such means, even in the good case, will always qualify as an action that the agent knows that they are performing.²⁷

²⁶ Admittedly, the kicking case is not the most natural illustration of the point. But there will be many such cases, since the range of basic actions available to skilled human agents is large.

²⁷ A referee raises the possibility that consideration of cases of controlled action in the absence of conscious awareness might strengthen our argument. Consider, for example, the visual form agnosia patient D.F. In spite of severe deficits in her capacity to recognize objects or their form, she retains the ability to perform actions that depend upon recognizing the form of objects. In an early study, Goodale et al. (1991) report that although D.F.'s performance on perceptual recognition tasks was 'highly variable and included gross errors such as judging horizontal to be vertical,' her action control capacities were preserved such that when D.F. was asked to post a

3. Control and knowledge-how

One might be inclined to think, based on what we've suggested so far, that what is centrally wrong with the epistemic theory of control defended by Beddor and Pavese is that it links one's control of ϕ -ing with *knowledge that one is ϕ -ing*. Perhaps an epistemic condition might be saved were we to opt only for the second leg of Beddor and Pavese's condition, and maintain that control requires *just* knowledge-how to ϕ (at t), rather than knowledge that they are ϕ -ing at t .

It seems promising, after all, that control and knowledge-how share some relationship. We have already seen how the control constraint links control and intentional action. Now consider a very similar claim linking intentional action and knowledge-how.

(IAK): Necessarily, if $S \Phi$ s intentionally, S knows how to Φ .

As Piñeros Glasscock observes, '(IAK) has received extensive support among philosophers of action' (2021, 3). So both control and knowledge-how seem, to many, intimately (for some, necessarily) involved in intentional action.

One can also make very similar normative claims regarding control and intention, and knowledge-how and intention. For the latter kind of claim, consider Josh Habgood-Coote's (2018) argument that knowledge-how (rather than knowledge-that) serves as a norm governing intention in a way broadly analogous to Williamson's (1996; 2000) view that propositional

card through slots of differing orientation, her performance was good. They write, 'analysis of video records of each reaching movement revealed that, like the controls, D.F. began to orient the card correctly even as her hand was being raised from the start position in this task' (1991, 155). Doesn't this indicate that D.F. exercised control in posting the card through orientation X , even though she lacked knowledge that she was doing so at the relevant time? Analysis of this kind of case has proven delicate, with some arguing that D.F. does after all have information about what she is doing under a certain mode of presentation (Mole 2009), and with others arguing that such cases do not demonstrate control by unconscious vision (Phillips 2021). A full consideration of the philosophical import of absence of awareness cases (e.g., cases involving blindsight, or rapid reflexes, or modification of intention by sensorimotor mechanisms) would require us to sort through difficult issues regarding the applicability of the notion of control to behaviors in absence of awareness cases (e.g., cases involving blindsight, or rapid reflexes, or modification of intention by sensorimotor mechanisms), the potential relationships between know-how and conscious awareness, and between know-how and unconscious mental states. In addition, one would need to defend the very viability of unconscious action cases – an issue that has lately generated a range of sophisticated discussion (see, e.g., Shepherd 2016, Wu 2020, Shepherd and Mylopoulos 2021, Kozuch 2022). We do not think it is possible to do such work in this paper.

knowledge is the norm of assertion. Suppose, for example – to use a case from Habgood-Coote – that one invites friends over for a dinner party and purports to intend to make the complicated dish *coq au vin*; if it came to light that one did not know how to make *coq au vin*, the intention expressed to the invited friends would seem problematic. The felicity of criticism in such a case is plausibly explained, Habgood-Coote thinks, by the idea that intention is normatively connected with know-how, such that one should intend to do something, ϕ , only if one knows how to ϕ . This suggestion linking know-how and intentional action is closely connected to Kieran Setiya's view that intending to do something is *epistemically justified* only when one knows how to do what one purports to intend to do; for Setiya, this is because knowledge-how provides epistemic entitlement for the beliefs one relies on in intentional action.²⁸

Of course, both Habgood-Coote's and Setiya's thinking that know-how is a norm on intention is compatible with intention persisting with a lack of know-how (just as, by way of analogy, Williamson will allow assertions to be genuine when not made on knowledge, even if when not made on knowledge they are defective.) Even so, notice that the idea that an intention to do something ϕ is defective if one lacks knowledge how to ϕ fits snugly with the idea that intention is defective if one lacks suitable control over one's ϕ -ing.

The question is, given their plausibly intimate relationship, how are control and knowledge-how related? If one wants an epistemic theory of control spelled out in terms of knowledge-how, one will claim that control *requires* knowledge-how – presumably, because knowledge-how is needed to explain control.

Epistemic Theory of Control - Version 2 (ETC-2): Someone is in control of their ϕ -ing at t only if (at t) they have knowledge of how to ϕ .

Might ETC-2 have a chance at doing better with cases than Beddor and Pavese's ETC?

²⁸ See Setiya (2008, 406–9; 2012, 300–304). For discussion, see Habgood-Coote (2018, 1701 and fn. 2).

Let's now consider this in some detail, beginning with a distinction. Know-how can be general or specific in a way that is broadly analogous to how abilities can be general or specific (Mele 2003). For example: we can distinguish knowing how to hit baseballs generally from knowing how to hit a baseball against this specific pitcher on this specific occasion.²⁹ With this distinction in hand, the question of whether being in control of one's ϕ -ing (as per ETC-2) requires that one have knowledge how to ϕ might be disambiguated either as (i) the claim that control over one's ϕ -ing at t requires (at t) general know-how to ϕ (viz., knowledge how to ϕ that is applicable across a broad range of ϕ -situations); or as (ii) the claim that control over one's ϕ -ing at t requires (at t) specific knowledge how to ϕ .

3.1 Control and general know-how

Let's consider first the 'general know-how' interpretation – and let's think about this in connection with the following case:

IDIOSYNCRATIC HITTER: Ichiro bats .350 — and thus plausibly exercises excellent control — when hitting against left-handed pitchers who throw fastballs. In fact, against left-handed pitchers throwing fastballs, Ichiro is the most proficient in the league. Against all other combinations, however, Ichiro is miserable (he bats just .150, considerably worse than the average player in the league). Tonight, fortunately for Ichiro, he faces a lefty who throws a fastball. Ichiro, a feared hitter in this scenario by his opponents, makes a hit.

In IDIOSYNCRATIC HITTER, the following diagnosis is plausible: (i) Ichiro is in control of his getting a hit, however, (ii) it is false that Ichiro has *general* know-how when it comes to hitting baseballs. (i) is plausible if we grant the control constraint, because it is plausible that Ichiro gets a hit *intentionally*. (ii) is true because Ichiro cannot reasonably direct his very specific skill suited to this occasion to the broad range of normal situations that are applicable to general know-how to hit. (Compare: we do not say that one has *general know-how to speak Finnish* if

²⁹ The distinction between general and specific abilities to A is often made with reference to whether the agent has the option to A in the specific circumstances. The distinction we are drawing here is thus slightly different: it concerns the coverage of the agent's knowledge how across a range of relevant situations.

one is in command of only one question/answer combination [“How are you?” = “mitä kuuluu?” / “Very well, thank you.” = “hyvä on Kiitos.”], and happens on this occasion to be asked only the question (mitä kuuluu?) for which one can provide an answer in Finnish (hyvä on Kiitos!))

It is worth emphasizing further that the fact that Ichiro lacks general know-how in IDIOSYNCRATIC HITTER despite exercising control is a diagnosis that does not depend on any antecedent substantive commitment to either intellectualism or anti-intellectualism as a thesis about the nature of know-how.

On the anti-intellectualist line (e.g., Ryle 1946, 1949; Noë 2005; Poston 2009; Carter and Pritchard 2015), the idea that Ichiro lacks general know-how in IDIOSYNCRATIC HITTER will be explained with reference to Ichiro’s lacking a general *ability* — viz., a disposition to hit (reliably enough) across a normal range of circumstances where one attempts or is expected to attempt to hit. He of course lacks this, with his .150 average in non-lefty non-fastball scenarios. On an intellectualist line (e.g., Williamson and Stanley 2001; Stanley 2011; Pavese 2021a; Brogaard 2012) — according to which know-how to ϕ is identified not with ability but with propositional knowledge (for some way, w , of ϕ -ing, that w is a way for one to ϕ)³⁰ Ichiro’s lacking general know how will be explained with reference to lacking propositional knowledge of a general way to hit balls reliably (and across a normal range of circumstances).³¹

Accordingly then, in IDIOSYNCRATIC HITTER, if Ichiro’s control over his hit against left-handed hitters (throwing fastballs) requires know-how on the part of Ichiro, it is not general know-how which his hopelessness in all other non-lefty non-fastball scenarios precludes.

³⁰ This is the standard template formulation; a much more detailed characterisation is presented in Stanley (2011). We are setting aside here the further complication that intellectualists generally appeal (in various ways) to modes of thinking under which propositional knowledge identified with knowledge how is appended — what Stanley and Williamson and Stanley refer to as a ‘practical mode of presentation’, and what Pavese (2015) calls a ‘practical sense’.

³¹ Note that this assessment doesn’t rely on any commitment to the thesis that safety is necessary for knowledge.

This suggests that if control requires knowledge how, it will be knowledge how that is as circumstantially restricted as control often is. That is fine, as far as it goes. But it raises complications.

3.2 Control and specific know-how

Is specific (i.e., circumstantially restricted) knowledge-how required for control? Here we only consider versions of knowledge-how that require propositional knowledge, for the following reason. Circumstantially restricted anti-intellectualist knowledge-how is a circumstantially restricted kind of practical ability. But so is control. Indeed, the most plausible relationship between control and anti-intellectualist knowledge-how of this sort is just identity.

So we are considering the relationship between circumstantially restricted (intellectualist) knowledge of a way to ϕ , and control in ϕ -ing. An epistemic theory of control would make knowledge-how necessary for control, presumably because knowledge-how provides a crucial element in any explanation of control. We stress necessity here since all sides can grant that knowledge of ways to ϕ can contribute to explanations of control in ϕ -ing in some or even many cases.

But it looks like the requirements on an epistemic theory of control are still too strong. One reason why comes from cases in which an agent comes to have a true but Gettiered belief regarding a way to do something. Cath (2011), for example, has someone named Charlie find true instructions regarding how to change a light bulb, but the fact that these instructions are true is due to a very lucky printing error. Cath judges that Charlie's belief regarding the way to change the light bulb does not amount to knowledge. Much of the following discussion in the literature surrounds the question of whether Charlie nonetheless knows how (in an anti-intellectualist sense) to change the light bulb (Stanley 2011, Carter 2012, Cath 2015, Cath 2019), but that is not our question here. If Charlie's belief about how to change the light bulb does not

amount to knowledge³², then (intellectualist) knowledge how comes apart from control, for clearly Charlie possesses (i.e., is able to exercise) control in changing a light bulb.

If one does not like that kind of case, consider another.

OLD METHOD, NEW CIRCUMSTANCES. Ichiro is a skilled hitter, and he is playing recreational baseball with friends, when a new pitcher shows up to play. This pitcher delivers the ball in a way Ichiro has never seen – the underhand style common to fastpitch softball. After seeing the pitcher warm up, it's Ichiro's turn. As he steps up to bat, he is agnostic about whether his normal way of batting – involving certain temporally synchronized patterns of attention directed to the pitcher's body language, and then to the flight of the ball, coupled with rapid decisional processes about whether and how to swing – will work.³³ Perhaps this new delivery will completely befuddle him. Perhaps a totally new method is required. Given his ignorance regarding how to approach this new pitcher, Ichiro decides to use his old method. It turns out, given Ichiro's natural talent (attentional capacities, speed of reflexes, etc.), that his old method works fairly well in these new circumstances. Ichiro sights the first pitch and makes contact.

Consider, first, that Ichiro did not know that his old method (way w) was a way to make contact with the ball until he had done it;³⁴ after all, he lacked a belief one way or the other; so, approaching the first pitch, Ichiro lacked intellectualist knowledge of how to make contact. He did, however, retain significant control with respect to making contact. This is why his making

³² Note that this point is compatible with granting that know-how might very well be present (despite a lack of knowledge-that) in a specifically environmental-luck twist on Cath's case, such as that presented by Carter and Pritchard (2015).

³³ If one wants to see softball pitcher Jenny Finch flummox one of the greatest baseball hitters of all time, see <https://www.youtube.com/watch?v=gm9iZnqGMvY>.

³⁴ Among other explanations for why this is so, Ichiro had undefeated mental-state defeaters for his belief that his old way was a way to make contact with the ball. Specifically, in the language of defeat, Ichiro seems to have what is called an *undercutting* defeater (see, e.g., Pollock 1986) – viz., a doubt that his old way was a reliable way to make contact with the ball (even if not a positive belief to the effect that his old way is in fact not reliable).

contact using his old method was no fluke. So Ichiro's control comes apart from his knowledge how in a way that undermines an epistemic theory of control.

(Interestingly, OLD METHOD, NEW CIRCUMSTANCES also casts doubt on the intellectualist version of the widely endorsed principle we mentioned above (IAK). If one wishes to hold on to that principle, then it looks like the thing to do is to jettison intellectualism.)

One might respond that Ichiro still knew how to do some things that put him in a position to succeed. That's true, and we don't deny that human agency (especially adult human agency) is supported in important ways by knowledge of how to do things. But that fact alone does not get one to an epistemic theory of control. And an epistemic theory of control that claims that control in ϕ -ing requires knowledge of how to do some things is not very illuminating. If we are to think that knowledge how does important explanatory work for control, we need to see some stable relationship between what the agent knows how to do and the agent's possession of control. We think the lesson here is that one should give up on an epistemic theory of control.

4. Explaining action-relevant knowledge via control

We began by observing the importance of non-accidentality to intentional action, and the natural place of control in our understanding of intentional action's non-accidentality. There are different accounts of control on offer, and our aim here was to examine specifically knowledge-based theories of control – theories that maintain that control is explained by knowledge. We have not denied that knowledge of various sorts can help explain control in particular instances. But a general knowledge-based theory of control fails.

Indeed, it seems to us that the deeper explanatory direction runs from control to the knowledge that is often relevant to action execution. Consider, for example, that the following is a plausible principle.

Control for know-how. An agent does not know how to A unless they possess – are disposed to exercise – sufficient control in A-ing in relevantly circumscribed circumstances.

This suggests the importance of control for know-how, but it does not yet say anything about control's explanatory force. But consider the following kind of case.

OVERCONFIDENCE. Jim has been told three ways to A (w, x, and y), and he believes, of all three, that they are good ways for him to A. But in fact only one is, for his abilities are overly taxed by the other two. (There are lots of ways to fill in the details. But say that A-ing is climbing a wall at the climbing gym, and the ways involve different routes. One route doesn't require too much grip strength, any overly difficult jumps, and so on. The other two routes require climbing moves that Jim would rarely be able to perform.)

In the above case, Jim only knows how to A via way w, but not via ways x and y. Why? Plausibly, because Jim only possesses sufficient control regarding the behaviors way w prescribes.

We are not the first to suggest that control, or something like it, can explain knowledge relevant to action execution. Dickie (2012) suggests that skill explains knowledge how. Shepherd (2021) suggests that control explains knowledge of action. Piñeros Glasscock (2021) suggests that 'possession of certain practical capacities, say, capacities to control an object in some characteristic way' (14) can help explain the acquisition of conceptually structured knowledge how. We raise the issue here not to develop it at length, but only to conclude our consideration of epistemic theories of control. The fact that control explains the presence of action-relevant knowledge is another reason to doubt that epistemic theories of control are on the right track. And it suggests that capturing the non-accidentality inherent to intentional action requires less epistemic resources than epistemic theories of control maintain.

5. Concluding remarks

Given that intentional action and knowledge (both knowledge-that and knowledge-how) are, respectively, importantly non-accidental, it should be no surprise then that philosophers of action might hope to account for the non-accidental character of intentional action with reference to knowledge. We have interrogated here one very specific strategy for doing this which has been taken up recently – and which maintains that the kind of *control* required for intentional knowledge is itself knowledge-involving. While Beddor and Pavese and others offer

interesting lines of argument for making good on this idea (viz., that control requires knowledge), we've argued that it ultimately doesn't hold water; if control required for intentional action 'blocks' accidentality, it's not by requiring knowledge, either knowledge-that or knowledge-how.³⁵

³⁵ 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 and the Leverhulme Trust's *A Virtue Epistemology of Trust* project (RPG-2019- 302). Shepherd's research is supported by European Research Council Starting Grant 757698 (for project ReConAg), awarded under the Horizon 2020 Programme for Research and Innovation, as well as the Canadian Institute for Advanced Research's Azrieli Global Scholar programme on Mind, Brain, and Consciousness.

References

- Adams, Frederick, and Alfred Mele. 1989. "The Role of Intention in Intentional Action." *Canadian Journal of Philosophy* 19 (4): 511–31.
- Aguilar, Jesús H. 2012. "Basic Causal Deviance, Action Repertoires, and Reliability." *Philosophical Issues* 22: 1–19.
- Anscombe, G. 1957. *Intention*. Harvard University Press.
- Beddor, Bob, and Carlotta Pavese. 2021. "Practical Knowledge Without Luminosity." *Mind*, October. <https://doi.org/10.1093/mind/fzab041>.
- Bengson, John and Marc A. Moffett. 2011. "Two Conceptions of Mind and Action: Knowledge How and the Philosophical Theory of Intelligence". In John Bengson & Marc Moffett (eds.), *Knowing How: Essays on Knowledge, Mind, and Action*. Oxford University Press. pp. 3-55.
- Brogaard, Berit. 2012. "Knowledge-How." *Knowing How: Essays on Knowledge, Mind, and Action*, 137–60.
- Brueckner, Anthony and M. Oreste Fiocco. 2002. "Williamson's Anti-luminosity Argument." *Philosophical Studies* 110(3): 285-293.
- Carter, J. Adam. 2012. On Stanley's Intellectualism. *International Journal of Philosophical Studies* 20: 749 – 62.
- Carter, J. Adam, and Neil McDonnell. 2018. "Welcome to the Machine." *The Philosophers' Magazine*, no. 81: 33–39.

- Carter, J. Adam and Ted Poston. 2018. *A Critical Introduction to Knowledge-How*. London: Continuum.
- Carter, J. Adam, and Duncan Pritchard. 2015. "Knowledge-How and Epistemic Luck." *Noûs* 49 (3): 440–53.
- Carter, J. Adam, Duncan Pritchard, and Joshua Shepherd. 2019. "Knowledge-How, Understanding-Why and Epistemic Luck: An Experimental Study." *Review of Philosophy and Psychology* 10 (4): 701–34.
- Carter, J. Adam and Ernest Sosa. 2022. "Metaepistemology." *Stanford Encyclopedia of Philosophy*.
- Cath, Yuri. 2011. Knowing how without knowing that. In *Knowing How: Essays on Knowledge, Mind, and Action*, eds. Bengson, J., Moffett, M., 113–35. New York: Oxford University Press.
- Cath, Yuri. 2019. "Knowing how." *Analysis* 79(3): 487–503.
- Chisholm, Roderick M. 1977. *Theory of Knowledge*. Prentice Hall, Englewood Cliffs, NJ,.
- Colaço, David, Wesley Buckwalter, Stephen Stich, and Edouard Machery. 2014. "Epistemic Intuitions in Fake-Barn Thought Experiments." *Episteme* 11 (2): 199–212.
- Dickie, Imogen. 2012. "Skill Before Knowledge." *Philosophy and Phenomenological Research* 85(3): 737-745.
- Di Nucci, Ezio. 2013. Action, deviance, and guidance. *Abstracta* 7(2): 41-59.
- Frost, Kim. 2019. A metaphysics for practical knowledge. *Canadian Journal of Philosophy*, 49(3), 314-340.
- Ginet, Carl. 1975. *Knowledge, Perception, and Memory*. D. Reidel Publishing Co.
- Goldman, Alvin I. 1979. "What Is Justified Belief?" In *Justification and Knowledge*, 1–23. Springer.

- Goodale, Melvyn A., A. David Milner, L. S. Jakobson, and D. P. Carey. 1991. A neurological dissociation between perceiving objects and grasping them. *Nature* 349, 6305: 154-156.
- Habgood-Coote, Joshua. 2018. "Knowledge-How Is the Norm of Intention." *Philosophical Studies* 175 (7): 1703–27. <https://doi.org/10.1007/s11098-017-0931-4>.
- Hetherington, Stephen. 2011. *How to Know: A Practicalist Conception of Knowledge*. John Wiley & Sons.
- Hetherington, Stephen. 2013. Knowledge Can Be Lucky. In Matthias Steup & John Turri (eds.), *Contemporary Debates in Epistemology*. Blackwell. pp. 164-73..
- Horvath, Joachim, and Alex Wiegmann. 2016. "Intuitive Expertise and Intuitions about Knowledge." *Philosophical Studies* 173 (10): 2701–26.
- Iverson, Grant L., Brian L. Brooks, and James A. Holdnack. 2008. "Misdiagnosis of Cognitive Impairment in Forensic Neuropsychology." *Neuropsychology in the Courtroom: Expert Analysis of Reports and Testimony*, 243–66.
- Kenny, A. J. P. 1976. *Will, Freedom and Power*. Blackwell.
- Kozuch, Benjamin. 2022. Conscious vision guides motor action—rarely. *Philosophical Psychology*, 1-34.
- Mele, Alfred, and Moser, Paul K. 1994. Intentional action. *Noûs* 28(1): 39-68.
- Mele, Alfred. 2003. Agents' abilities. *Noûs* 37(3): 447-470.
- Mole, Christopher 2009. Illusions, demonstratives, and the zombie action hypothesis. *Mind*, 118(472), 995-1011.
- Newstead, Anne. 2006. "Knowledge by Intention? On the Possibility of Agent's Knowledge." In *Aspects of Knowing*, 183–201. Elsevier.
- Noë, Alva. 2005. "Against Intellectualism." *Analysis* 65 (4): 278–90.
- Paul, Sarah. 2011. Deviant Formal Causation. *Journal of Ethics and Social Philosophy*, 5(3), 1-24.

Pavese, Carlotta. 2015. "Practical Senses." *Philosophers' Imprint* 15.

———. 2021a. "Know-How, Action, and Luck." *Synthese* 198 (7): 1595–1617.

———. 2021b. "Knowledge and Mentality." *Philosophical Perspectives*.

Pavese, Carlota, and Beddor, Bob. Skills as Knowledge (forthcoming). *Australasian Journal of Philosophy*.

Phillips, Ian 2021. Blindsight is qualitatively degraded conscious vision. *Psychological Review*, 128(3), 558-584.

Piñeros Glasscock, Juan S. 2020. "Practical Knowledge and Luminosity." *Mind* 129 (516): 1237–67.

———. 2021. The puzzle of learning by doing and the gradability of knowing how. *Philosophy and Phenomenological Research*. On-line first, DOI: 10.1111/phpr.12832

Poston, Ted. 2009. "Know How to Be Gettiered?" *Philosophy and Phenomenological Research* 79 (3): 743–47.

Pritchard, Duncan. 2007. "Anti-Luck Epistemology." *Synthese* 158 (3): 277–97.

———. 2012. "Anti-Luck Virtue Epistemology." *The Journal of Philosophy* 109 (3): 247–79.

Rabinowitz, Dani. 2011. "The Safety Condition for Knowledge." *Internet Encyclopedia of Philosophy*.

Ramachandran, Murali. 2009. "Anti-Luminosity: Four Unsuccessful Strategies". *Australasian Journal of Philosophy*, 87(4): 659-673.

Rödl, Sebastian. 2011. "8. Two Forms of Practical Knowledge and Their Unity." In *Essays on Anscombe's Intention*, 211–41. Harvard University Press.

Ryle, Gilbert. 1946. "Knowing How and Knowing That: The Presidential Address." *Proceedings of the Aristotelian Society* 46: 1–16. <https://doi.org/10.1093/aristotelian/46.1.1>.

———. 1949. *The Concept of Mind*. Hutchinson & Co.

Setiya, Kieran. 2008. "Practical Knowledge." *Ethics* 118 (3): 388–409.

———. 2012. "XIV—Knowing How." In *Proceedings of the Aristotelian Society* (Hardback), 112:285–307. Wiley Online Library.

Shepherd, Joshua 2016. Conscious action/zombie action. *Noûs*, 50(2), 419-444.

———. 2017. Half-hearted action and control. *Ergo* 4(9): 257-276.

———. 2021. *The Shape of Agency: Control, Action, Skill, Knowledge*. Oxford University Press.

Shepherd, Joshua, and J. Adam Carter. Forthcoming. "Knowledge, Practical Knowledge, and Intentional Action." *Ergo: An Open Access Journal of Philosophy*.

Shepherd, Joshua, and Myrto Mylopoulos 2021. Unconscious perception and central coordinating agency. *Philosophical Studies*, 178(12), 3869-3893.

Small, Will. 2012. "Practical Knowledge and the Structure of Action." In *Rethinking Epistemology*, 133–228. de Gruyter.

Sosa, Ernest. 2010. "How Competence Matters in Epistemology." *Philosophical Perspectives* 24 (1): 465–75.

———. 2015. *Judgment and Agency*. Oxford: Oxford University Press.

———. 2021. *Epistemic Explanations: A Theory of Telic Normativity, and What It Explains*. Oxford University Press.

Srinivasan, A., and J. Hawthorne. 2013. "Disagreement Without Transparency: Some Bleak Thoughts." In *The Epistemology of Disagreement.*, edited by David Christensen and Jennifer Lackey, 9–30. Oxford University Press.

Stanley, Jason. 2011. *Know How*. OUP Oxford.

Turri, John, Wesley Buckwalter, and Peter Blouw. 2015. "Knowledge and Luck." *Psychonomic Bulletin & Review* 22 (2): 378–90.

Williamson, Timothy. 1996. "Knowing and asserting." *The Philosophical Review* 105(4): 489-523.

———. 2002. *Knowledge and Its Limits*. Oxford University Press on Demand.

———. 2017. Acting on knowledge. In *Knowledge First: Approaches in Epistemology and Mind*, Carter, J.A., Gordon, E.C., and Jarvis, B. (eds.), 163-183. Oxford University Press.

Williamson, Timothy, and Jason Stanley. 2001. "Knowing How." *The Journal of Philosophy* 98 (8). <https://ora.ox.ac.uk/objects/uuid:1fcf732b-731b-47e6-943d-aaf34905c02a>.

Wu, Wayne 2020. Is vision for action unconscious?. *The Journal of Philosophy*, 117(8), 413-433.