

KNOWLEDGE, PRACTICAL KNOWLEDGE, AND INTENTIONAL ACTION

JOSHUA SHEPHERD

*Carleton University
University of Barcelona*

J. ADAM CARTER

University of Glasgow

We argue that any strong version of a knowledge condition on intentional action, the practical knowledge principle, on which knowledge of what I am doing (under some description: call it A-ing) is necessary for that A-ing to qualify as an intentional action, is false. Our argument involves a new kind of case, one that centers the agent's control appropriately and thus improves upon Davidson's well-known carbon copier case. After discussing this case, offering an initial argument against the knowledge condition, and discussing recent treatments that cover nearby ground, we consider several objections. One we consider at some length maintains that although contemplative knowledge may be disconnected from intentional action, specifically practical knowledge of the sort Anscombe elucidated escapes our argument. We demonstrate that this is not so. Our argument illuminates an important truth, often overlooked in discussions of the knowledge-intentional action relationship: intentional action and knowledge have different levels of permissiveness regarding failure in similar circumstances.

1. Introduction

Our target is the connection between knowledge of action and intentional action¹—between an agent's knowledge of what they are doing and their doing

1. The usage of the term 'intentional action' may invite some contention, since theorists understand the term in different ways. We have in mind a category of agentive behavior that bears a close relationship to what the agent intends to do, and that is usually explained by reference to an agent's reasons for behaving in such a way, as well as by reference to the agent's control with respect to behaving in that way.

Contact: Joshua Shepherd <joshualshep@gmail.com>
J. Adam Carter <adam.carter@glasgow.ac.uk>

it intentionally. Primarily (though not exclusively), here we are interested in whether this knowledge is necessary for intentional action. Consider the following formulations of a necessity principle.

Practical Knowledge Principle (PKP): '(Necessarily) If an agent is F-ing (intentionally and under that description), she knows that she is F-ing (intentionally and under that description).' (Piñeros Glasscock 2019: 1238)

Knowledge Thesis (KT): 'An agent is doing something intentionally only if he knows that he is doing it.' (Vekony et al. 2020: 1)

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 .' (Beddor & Pavese 2021: 6)

There are important differences between these formulations. But we focus on the idea these formulations share—that knowledge of what I am doing (under some description: call it A-ing) is necessary for that A-ing to qualify as an intentional action. In this paper, we argue that this idea is false.

We do so by introducing a new kind of case, one that, we argue, centers the agent's control appropriately and thus improves on Davidson's well-known carbon copier case. In Section 2 we introduce the case, offer an initial argument—which leans on the idea that intentional action and knowledge have different levels of permissiveness regarding failure in similar circumstances—consider friendly recent arguments due to Piñeros Glasscock (2020) and Schwenkler (2019), and consider the objection that our argument relies too heavily on the notion of safety. In Section 3 we consider how the argument might apply to certain understandings of practical knowledge—knowledge that Anscombe famously claimed is 'the cause of what it understands' (2000: 87–88). We contend that the argument is also troublesome for a knowledge condition that appeals to practical knowledge. In Section 4 we consider options for elucidating the thought that though a knowledge condition formulated in terms of necessity fails, knowledge and intentional action remain intimately connected in some way.

2. A New Case, and the Initial Argument

A number of philosophers, persuaded by Davidson's carbon copier case, have thought knowledge unnecessary for intentional action. Here is the case:

in writing heavily on this page I may be intending to produce ten legible carbon copies. I do not know, or believe with any confidence, that I am succeeding. But if I am producing ten legible carbon copies, I am certainly doing it intentionally. (Davidson 1980: 92)

Perhaps Davidson's judgments are correct. But his case invites a certain kind of reply—namely, that the agent is not intentionally producing the ten legible copies (Small 2012; Thompson 2011; Wolfson 2012).²

Why deny that the agent is intentionally producing the copies? We find a similar idea in certain authors (e.g., Thompson 2011; Small 2012), the quick version of which is that the copier's success is too lucky to count as intentional action. Thompson claims that for the copier, 'the making of the inscription is like the buying of a lottery ticket' (2011: 210). Small compares the copier to an agent with no training and no speed attempting to run a four-minute mile, and suggests that if the copier succeeds, then he does so *through luck* (2012).³ The availability of such claims, we submit, is due to the fact that Davidson's case does not center the agent's control over behavior in the right way. As a result, a deeper issue does not get proper attention.

Here is a case involving baseball that brings the right issue to light. For those not versed in baseball, we focus on batting. The batter stands 60 feet, 6 inches from a mound upon which stands the pitcher. The pitcher throws a small white ball towards the batter, and the batter attempts to hit the ball with a bat. This is not always easy, since the pitch comes in at varying speeds (often anywhere from 80 to 100 mph), with different kinds of spin and curvature.

Say that the greatest hitter of all time (call him Pujols) approaches the plate and forms an intention to hit a home run—that is, to hit the ball some 340 feet or more in the air, such that it flies out of the field of play. Say that Pujols also believes that he will hit a home run, and that he has the practical belief as he is swinging, that he is hitting a home run. As it happens, Pujols's behavior, from setting his stance and eyeing the pitcher, to locating the pitch, to swinging the bat and making contact with the ball, is an exquisite exercise of control. Pujols hits a home run, and so his belief that he is doing just that is true.

Does Pujols *intentionally* hit a home run? Our intuition is that he does. (If he doesn't, one should probably say that no major leaguer ever does.) Does Pujols *know* at any point during the action that he will hit a home run, or that he is hitting a home run? Our intuition is that he does not. These intuitions suggest the formulation of an initial argument against the necessity of knowledge for intentional action. All it takes, after all, is one case.

Is there anything structurally unsound about this case? It doesn't seem so. Can Pujols really intend to hit a home run? It seems clear that he can. It is controversial

2. Stathopoulos (2016) offers a different kind of reply to Davidson's case, involving the claim that the agent is not intentionally producing the ten legible copies because having done so would be an achievement, and achievements are not performed intentionally.

3. To be fair, these claims are embedded in subtle discussions of the nature of practical knowledge, and charity requires further reflection on the context. We consider Thompson's view in more detail in Section 3, and Small's view in more detail in the same section.

whether agents can intend the impossible (Thalberg 1962; Ludwig 1995), but not whether agents can intend to do what they frequently do, and indeed, do reliably enough to merit praise. (Don't forget that we are talking about the greatest hitter of all time.) Can Pujols believe that he will hit a home run, and that he is hitting a home run? Again it seems clear that he can. One might object that home runs are too rare—too close to buying a lottery ticket—and that not even the greatest hitter of all time hits them intentionally.⁴ That judgment is plausible regarding some agentive achievements. We are reticent to judge, for example, that the greatest golfer of all time intentionally makes a hole-in-one, the probability of which, for even a professional, is approximately 3,000:1.⁵

So let us switch the case to getting a base hit—that is, hitting the ball into the field of play such that it evades defenders for long enough that the batter can run to (at least) first base. Albert Pujols (the actual person) may not be the greatest hitter of all time, but he was probably the best hitter in baseball during the first ten years of his career, viz., between 2001 and 2011, when he was with the St. Louis Cardinals. Pujols hit for an average of .359 during the 2003 season, when the average MLB player that year batted just .263. (His home-run-per-at-bat percentage was much lower that year—.073.) Could 2003 Pujols have intentionally got a base hit, something he did 212 times in his 157 games? Or were all of those hits 'too lucky'? We have a strong intuition, and we expect most fans of baseball would agree, that Pujols intentionally got many of those hits.⁶

We need not rest content with brute intuition here. Recall that Pujols's behavior is, by stipulation, an exquisite exercise of control. In our view this is a significant part of the explanation for why his behavior qualifies as intentional action. There are of course different uses of control in the literature, and different accounts (for discussion, see Mele & Moser 1994; Riggs 2009; Aguilar 2012; Broncano-Berrocal 2015; Wu 2016; Shepherd 2021). But our appeal to control

4. The two most frequent home-run hitters in MLB history, with at least 3,000 plate appearances, were Mark McGwire and Babe Ruth, respectively. The former hit a home run every 10.61 at bats and the latter every 11.76 at bats. Thus, even among these very elite sluggers, the probability of a homerun in a given at bat was <.1.

5. The odds are approximately 12,000:1 for an amateur player, according to hole-in-one data being tracked by Golf Digest since the 1960s (<https://www.golfdigest.com/story/gd200509-david-owen-aces>).

6. Note that this intuition is entirely compatible with our granting two concessions, both of which are relatively uncontroversial. First, it is compatible with the concession that there are other batting statistics expressible as rates (e.g., strikeout rate) that may be even better markers of hitting ability than batting average. This remains a live area of research in sabermetrics (see Albert 2004). Secondly, the intuition that Pujols intentionally made many of these hits during his peak season as measured by batting average can be embraced unproblematically alongside a concession that batting average is arguably an imperfect measure of total batting ability for the reason that batting average reflects only the official at-bats (thus, not walks) and also assumes that every base hit has the same value. For discussion on both of these points in the context of recent work in sabermetrics, see, along with Albert (2004), also Albert and Bennett (2003: chs. 6–8).

does not depend upon any particular account, and it does not require a notion of control that is particularly controversial. The basic ideas we need, and which most find very plausible, are that control [a] is a hallmark of intentional action,⁷ [b] results from the coordination of various agentive capacities, often structured by practice and intelligent planning, such that [c] control is luck reducing.⁸ In general, the greater the degree of control exercised, the less room for luck to play a role in explaining an agent's success.

So, when Pujols locks onto the pitch, and executes his swing, his behavior is the product of years of practice, heaped onto a good measure of physical power, grace and inborn talent, that guarantees a measure of reliability in the way that his body moves. There is a reason that Pujols consistently gets hits at a higher rate than competitors. His higher rate is not due to luck, but to control. But can we really ascribe success to control when Pujols only succeeds one third of the time (as indicated by his excellent batting average of roughly .333)? This is where appeal to our intuitions, and common practices of attributing intentional action, come into play. We submit that when intentional action is at issue, it is commonly the case that explanations that advert to control sit comfortably alongside the admission that in nearby cases, the agent fails. In part this is because fallibility is a hallmark of human agency, and agents engaged in difficult endeavors know that some amount of risk, and indeed luck, is involved in many agentive achievements.⁹ Intentional action is a success term, but its proper usage reflects our tacit sense that the cooperation of circumstance is often required to some degree—even for relatively simple actions.

Is the same thing true of knowledge? Supposing Pujols believes he is getting a hit, and is getting a hit, does Pujols know that he is getting a hit? With a failure rate that high, it seems to us that the answer is no.¹⁰

The underlying truth here is that intentional action and knowledge have different standards for safety—or if you like, different levels of permissiveness

7. There are different ways of cashing out 'hallmark.' We find it plausible that some degree of control is necessary for intentional action—that a completely uncontrolled event could not be an intentional action. One finds this idea endorsed and elucidated in different ways in, e.g., Mele and Moser (1994), Shepherd (2021), and Beddor and Pavese (2021).

8. Of course, the mean relative to which luck is reduced by control will plausibly differ across different performance domains where different levels of reliability are expected for competent performance (see, e.g., Sosa 2015: ch. 3 and 2017: ch. 12; Greco 2010: ch. 5; and Carter et. al 2015). For example, you plausibly lack the ability to ride a bike if you fall off the bike 4/10 times when you try to ride (while staying on 6/10 times)—and this is so even though competent hitting (in baseball) does not require *even* hitting base hits at a rate of 4/10 times.

9. For discussion on this point, see Pritchard (2016).

10. This point is consistent with the claim that in some cases, Pujols could know, in the midst of hitting, that he is getting a hit. Arguably this could happen if Pujols spots a pitch cleanly and early, knowing exactly how to approach it and being highly skilled at doing so.

regarding failure in similar circumstances.¹¹ The idea that knowing any fact—including facts about what we are doing—‘requires safety’ is often glossed, to a first approximation, as the idea that a subject, *S*, knows that *p* only if the agent couldn’t easily have been incorrect about whether *p*, given how *p* was formed. One standard way to unpack this, following Pritchard (2005; cf. Smith 2017), is modally: the idea is that *S*’s belief is safe in the sense knowledge demands just in case in most nearby possible worlds in which *S* 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: 81).¹²

Philosophers of action do not tend to use the term ‘safety’, but a non-accidentality condition on intentional action roughly (and only roughly) resembles the safety condition on knowledge. Here we highlight an important difference between knowledge and intentional action regarding the permissiveness of this condition. It is possible, and in our view frequent, for an agent to act intentionally under a description at least in part in virtue of the control they exercise, while failing to know that they act intentionally under that description, because the success of what they are doing is not sufficiently safe.

The baseball case helps illustrate this underlying truth, but there are many other such cases. All we need is to find cases that [a] stand out as clear examples of achievement in virtue of control, and thus as very intuitive examples of intentional action, while at the same time [b] stand out as fairly unsafe successes in the sense that failures dot at least some nearby possible worlds. Such is the case with many activities where an agent’s success rate grades out somewhere in the middle percentiles.

Michael Jordan hitting a game winning shot in basketball. This is something Jordan did at a roughly 50% clip, and a shot professional players typically make around 30% of the time.

Megan Rapinoe converting a penalty kick in soccer, something players do around 75% of the time, although Rapinoe’s percentage, which we do not know, is likely to be higher.

11. Although most find safety plausible as at least a necessary condition on knowledge, not everyone does (see Kelp 2009). But the more general point, regarding levels of permissiveness regarding failure, holds independently of details about a safety condition.

12. An advantage of this formulation in epistemology is that it rules out knowledge in standard Gettier cases (i.e., as, in those cases, the subject’s belief is false in very close possible worlds where we hold fixed the way the target belief was formed in the actual world). The modal strength of a safety condition depends, among other things, on what is held fixed in the actual world when we ask whether the subject continues to believe the proposition truly in near-by possible worlds. In some formulations, these further details are not made explicit. For example, according to Williamson, ‘If one knows, one could not easily have been wrong in a *similar case*’ (2000: 147, our italics). For some representative defences of a safety condition on knowledge, see, e.g., Luper-Foy (1984), Sainsbury (1997), Sosa (1999), Williamson (2000), and Pritchard (2002; 2005; 2007). For an overview of these proposals, see Pritchard (2008) and Rabinowitz (2011).

A golf pro, any golf pro, sinking an 8-foot putt, something the pro does a little over 50% of the time. Or sinking a 12-foot putt, something the pro does around 30% of the time.

Simone Biles, the world's greatest gymnast, landing a double-double dismount—that's two flips and two twists while in a tucked position, a move so hard almost no one can do it.

Chris Carpenter, or any great pitcher, painting the corner with a fastball (i.e., making a very accurate pitch). Coco Gauff, or any tennis pro, placing a serve right on the T. Anna Gasser (a snowboarder) landing a backside corked 1080 melon grab, a move that is easier to watch than to describe.

Perhaps one's eyes glaze at all the professional sports examples, involving world-class athletes. Everyday life offers plenty as well:

When Jerry takes the trash out, the city bin's lid does not open very wide. He often wishes to lift the bag over the lip of the bin without scratching his elbow on the side of the bin, and he always intends to do so. He succeeds only sometimes.

Beth is making eggs, and she hates spilling yolk. So she intends to crack the eggs without letting any leak down the side of the bowl and onto the counter. She often fails.

Rick intends to shave over a rough patch of his skin without drawing any blood. This is difficult, but not impossible, requiring a deft touch. He uses the same method every time. Sometimes he gets it right.

Morty has to pilot a spacecraft into the heart of an alien citadel, while being chased by enemy warships. (Not everyone's everyday life is as boring as, e.g., Jerry's.) He knows the way, and he is a capable pilot, but success is not assured.

And so on.

The initial argument, then, runs as follows. Once one focuses on activities where an agent's success rate grades out somewhere in the middle percentiles, examples of intentional action without knowledge proliferate. The chief reason is that intentional action and knowledge have different levels of permissiveness regarding failure in similar circumstances. Intentional action frequently occurs in the face of nearby failure; not so regarding knowledge. So, the initial argument concludes, knowledge of A-ing is not necessary for intentional A-ing.

2.1. *Fellow Travellers*

Before we consider important objections to our argument, it is worth noting two recent discussions that arrive at a conclusion similar to ours, albeit on the basis of slightly different kinds of considerations.

In chapter six of his book on Anscombe's *Intention*, John Schwenkler (2019) turns to cases of unsafe behavior that he takes to be 'decidedly unusual.'

Considering Anscombe's claim that 'A person does something intentionally only if she *knows* that she is doing this' (Schwenkler 2019: 180), he considers the possibility of 'an *unfavorable situation*,' which is 'one in which a person's power to act in a certain way is compromised, such that there is a significant likelihood of practical error or failure, or that her attempt will be thwarted or she'll unknowingly attempt something futile' (2019: 183).

Schwenkler states the following.

Being in an unfavorable situation will not always leave a person *altogether* unable to act as she intends, nor need it be the case that a person in an unfavorable situation who manages to do what she intends to is not acting *intentionally* in that way, and instead can act in that way only by accident. . . . there appear to be circumstances in which a person's power to do X is reliable enough for her to do X intentionally, and where she is in fact exercising this power in such a way that a "Why?"-question would be given application to her action, but still she is not reliable enough in this exercise for her belief that she is doing this to amount to *knowledge* of what she is doing. (2019: 183–84)

The relevant convergence between us and Schwenkler is that we agree that Anscombe's claim is not true as stated, and for an apparently similar reason, to do with differences in safety requirements. An important difference, to which we return at the end of this paper, is Schwenkler's emphasis that the problem case is decidedly unusual. Schwenkler goes on to argue that while Anscombe's claim is falsified 'in full generality,' its spirit survives *because* 'the circumstances in which [Anscombe's claim] . . . is not satisfied are decidedly unusual: it takes some philosophical ingenuity to construct them, and we don't often encounter situations like these in ordinary life' (2019: 190).

Juan Piñeros Glasscock (2020) offers a different kind of argument that he claims undermines PKP: '(Necessarily) If an agent is F-ing (intentionally and under that description), she knows that she is F-ing (intentionally and under that description)' (Piñeros Glasscock 2020: 1238). His argument draws on anti-luminosity arguments in epistemology—where a condition (e.g., feeling cold) is luminous if, whenever an agent is in that condition, they are in a position to know that they are in this condition. Anti-luminosity arguments purport to show that certain conditions are not luminous for the agent, and an anti-luminosity argument regarding intentional action purports to show that agents may be acting intentionally while not being in a position to know that this is the case.

Piñeros Glasscock's version depends upon placing the agent in a position of uncertainty regarding whether they intend to A, or merely aspire to A. In his case,

Sisyphus is cleaning floors with water that progressively becomes dirtier as he proceeds. Sisyphus knows he is intentionally cleaning the floors at the beginning. And he knows that he is not doing so at the end. But could Sisyphus's knowledge track what he is doing—cleaning, or merely moving dirty water around—at every temporal increment? Consider moving along the timeline incrementally, where each second or millisecond we stop to consider whether Sisyphus is intentionally cleaning, and whether he knows that he is. If knowledge requires safety, and if Sisyphus knows he is intentionally cleaning the floor at time t , Piñeros Glasscock reasons that he must be intentionally cleaning the floor at $t+1$ (otherwise his knowledge that he is intentionally cleaning at t is unsafe). Now apply PKP to the case. We get the result that Sisyphus is intentionally cleaning the floor at all times in the sequence. But Sisyphus isn't: '(PKP) leads to contradiction, so it must be rejected' (2019: 1249).

Let us zoom in on the case—call it case X —in which Sisyphus is intentionally cleaning the floors, but does not know this. Piñeros Glasscock follows Small (2012) in holding that when the agent lacks sufficient control over A -ing, the agent does not intend to A , but instead merely aspires to A . Piñeros Glasscock claims that when a very nearby case Y involves a mere aspiration, the agent may act intentionally in case X , but fail to know that she is doing so, since knowledge requires safety that covers case Y , and intentional action does not.

In that kind of case, even when the agent forms a belief on the basis of an intention that ensures that she is acting intentionally, her belief might not constitute knowledge because there is a nearby case where she merely has an aspiration, so that she is not acting intentionally. (Perhaps she acts 'aspirationally'.) At some point in the series, the case of Sisyphus has this structure, and this explains why his beliefs about what he is doing, even when true, fail to constitute knowledge. (2020: 1256)

We do not accept the point about control and intending. For reasons Ludwig (1995) gives, it seems like agents may sometimes intend the impossible. But this point can be extracted from the case. If so, Piñeros Glasscock's point is similar to ours—failures of safety may undermine knowledge without undermining intentional action. Notice, however, an important difference. Piñeros Glasscock's anti-luminosity argument has it that in case X , the agent has sufficient control, while in case Y , the agent does not. The continually changing circumstances change the agent's control progressively as the water dirties.¹³ Our point is distinct in that

13. One way to see the argument here is that Piñeros Glasscock holds that the best defense of PKP posits a strong connection between control and knowledge, such that the possession of control provides grounds for knowledge in cases of intentional action. If this is the best defense of PKP, then Piñeros Glasscock's argument is well-designed to undermine it, given the posited loss of control over the course of the case. Thanks to Juan Piñeros Glasscock for discussion.

we need no case Y—so our cases do not depend upon a contrast in the agent’s control between X and Y.¹⁴ The point is actually that the agent may have an amount of control that is the same between case X and Y, such that in case X she succeeds in A-ing, and in case Y she fails to A. Case X is a case of intentional A-ing, and case Y, obviously, is not.

It seems to us that Piñeros Glasscock’s Sisyphus case still does not center the agent’s control in the right way. As a result, one might wonder why we should consider the agent’s behavior in case X an instance of intentional action. If her control is this fragile, one might think, why not deny intentional action here? We submit that our cases bolster Piñeros Glasscock’s assertion in case X. If it can be shown that intentional action in cases of far-from-foolproof control are ubiquitous—and our cases do show this—there is no general problem with Piñeros Glasscock’s assertion that the agent’s behavior in case X is an intentional action.

2.2. *An Initial Objection*

An initial objection to our argument presses on our reliance upon the notion of safety. The idea that propositional knowledge *as such* must be safe, the objector says, isn’t as sacrosanct as it is being made out to be. So why can’t the defender of a knowledge condition on intentional action push back against Pujols-style examples by simply denying that propositional knowledge must *always* be safe? Such a move offers the defender a basis to diagnose Pujols-style cases as cases where intentional action is present, but so is knowledge, because even Pujols’s practical belief (that he is getting a hit) is unsafe, *this* alone needn’t prevent Pujols from knowing that he is getting a hit.

In reply, one initial worry is that this denial—that propositional knowledge must be safe in the special case where the propositional knowledge is that one is F-ing—is going to be arbitrary. *However*, the above objection might simply press more strongly by denying that knowledge *ever* strictly requires safety, even if it is often compresent with it. This is tantamount to simply denying what is perhaps the most standard way to unpack what Pritchard (2012) calls the ‘anti-luck’ platitude in epistemology. But let us suppose one *does* deny this platitude; this is a theoretical cost, presumably, but might—as the thought goes—the cost be worth it if one can save a knowledge condition on intentional action?

14. Beddor and Pavese (2021) argue that a reformed knowledge condition on intentional action escapes Piñeros Glasscock’s anti-luminosity argument. But, as the considerations adduced in this paragraph demonstrate, our argument is different, and spells trouble for their reformed condition as well.

We think the answer here is ‘no’ because—setting aside what the theoretical cost is to denying that propositional knowledge must be safe—there are in fact *other conditions* on propositional knowledge one might appeal to in order to get the same result—viz., that Pujols-style cases undermine a knowledge condition by constituting cases of intentional action without the corresponding propositional knowledge. We will focus on two such conditions, which concern (i) aptness; and (ii) relevant alternatives. Regarding (i), according to Ernest Sosa (2015), a belief is known iff it is apt, and it is apt iff one’s believing truly manifests an intellectual *competence* of hers, where intellectual competences are (highly) reliable dispositions to believe truly in normal conditions.¹⁵ A common criticism of Sosa’s proposal is that one can have aptness *without* safety, as appears to be the case in fake-barn cases.¹⁶ Thus, the idea that knowledge requires aptness needn’t involve the commitment to knowledge requiring safety. Crucially, though, notice that Pujols’s belief that he is getting a hit (whilst he is swinging the bat) is not apt (and so not knowledge on Sosa’s proposal) even when he does get the hit, and when his getting the hit manifests a baseball-hitting competence of his. The reason here is very simple: the reliability threshold for *epistemic competences*—viz., the kind of belief forming processes the exercising of which is capable of generating knowledge—is going to be much higher than Pujols’s batting average! It is debatable how reliable a knowledge-generating epistemic competence must be; this is, in effect, a version of the generality problem for reliabilism.¹⁷ If we must put a number on it, it will be closer to .9 than to Pujols’s average which is less than .4.

A second idea in epistemology (besides aptness) which furnishes us with a principled reason to deny Pujols with the relevant knowledge *without* insisting that propositional knowledge requires safety adverts to *relevant alternatives theory* (e.g., Lewis 1996; Pritchard 2010). On this view, propositional knowledge doesn’t require of a thinker that she be able to rule out *all* error-possibilities incompatible with the target proposition. For instance, the relevant alternative theorist will deny that in order to know there is a hand in front of your face, you

15. Sosa unpacks normal conditions here in terms of one’s being in proper shape and properly situated to undertake the relevant kind of performance. To give a simple example, suppose you would believe unreliably what the colour of the wall is *when you have been drugged by a hallucinogen*. Does this fact—viz., that you would be reliable under these circumstances—count against your possessing perceptual belief forming competence? The answer, for Sosa, is ‘no’—and this is because the conditions under which good performance matters do not include drugged conditions. What matters for competence possession is just that one would perform reliably enough *when one is in fact* in proper shape and properly situated. Aptness then requires that one believe truly via a competence exercises *in* proper shape and when one is properly situated. For discussion, see, e.g., Sosa (2015; 2017).

16. See, for discussion, Sosa (2010). For the original presentation of fake barn cases in epistemology, see Ginet (1975).

17. For a recent and detailed overview of this problem, see Lyons (2019).

need be able to distinguish your hand from a Matrix projection,¹⁸ even if you *do* need to distinguish your hand from, for example, your foot. What is apparent is that the *most salient* candidate for a relevant alternative to swinging and getting a base hit is swinging and *not* doing so. However, as one is intentionally swinging the bat, one plausibly *can't* rule out this (highly) relevant alternative; after all, in addition to striking out, one of the highest risks to failing to hit safely is to hit into a fielder's out, something one can't easily discern whether one is doing or not even when one is swinging and connecting with the ball. Thus, *even if* we make no reference to practical knowledge (of the sort that features in a knowledge condition) requiring safety, we can nonetheless offer a principled reason for what it is lacking epistemically in Pujols-style cases in the language of relevant alternatives.

The case against a knowledge condition on intentional action is thus starting to look more robust. But an important objection looms.¹⁹

3. Practical Knowledge and Intentional Action

Anscombe, who is a key source of inspiration for many who maintain that knowledge is necessary for intentional action, famously distinguished between contemplative and non-contemplative forms of knowledge. While contemplative knowledge mirrors the world, non-contemplative knowledge is 'the cause of what it understands' (2000: 87–88). A central case of non-contemplative knowledge, for Anscombe, is the case of practical knowledge—a special kind of self-knowledge of what the agent is doing. The important objection to our initial argument is that the argument makes most sense if applied to contemplative knowledge, but fails to take seriously the unique nature of non-contemplative, practical knowledge. Since this is the form of knowledge in view for most philosophers tempted to maintain a necessary connection between knowledge and intentional action, an argument that did not apply to this form would be, honestly, a bit of a sad sideshow.²⁰

18. The situation complicates for the relevant alternative theorist, however, when this kind of framework is paired with versions of attributor contextualism which permit conversational context to raise epistemic standards (Lewis 1996). On such views, for instance, if your epistemology teacher brings up the possibility of Matrix projections, their doing so can convert what was previously an irrelevant alternative into a relevant alternative. One mechanism by which this can be explained as happening is given via Lewis's 'rule of attention'. For discussion, see McKenna (2014).

19. Thanks to referees for clearly articulating versions of this objection.

20. In order to give the objection room to breathe, we do not consider in detail the objection that there is no legitimate category for a radically non-contemplative practical knowledge, as Anscombe conceived of it.

How, then, might our case, and our argument, apply to practical knowledge? Answering this question is not straightforward since, as theorists of practical knowledge are quick to admit, the notion itself is contested, as are interpretations of how Anscombe understood it.²¹ Debates are live regarding the sense(s) in which this knowledge has been said to be non-observational (Schwenkler 2015), as well as non-inferential (Setiya 2008), how this knowledge relates to perception (Pickard 2004; Grünbaum 2011), whether this knowledge should be construed as a judgment or not (Small 2012; Marcus 2018; Frost 2019; Stout 2019), how this knowledge relates to intention (Falvey 2000; Paul 2009), whether to understand this knowledge's causal role in terms of formal causation (see Moran 2004), or efficient causation, or both (see Schwenkler 2015), whether this knowledge could be understood as a form of knowledge-that, or knowledge-how, or whether it is *sui generis* (see Frost 2019). Answering this practical knowledge objection in a way that is comprehensive regarding all of these debates threatens to fracture along many of these points.

Our approach, then, is to trace a few lines of thought in recently influential discussions of practical knowledge. These discussions emphasize practical knowledge as distinct from contemplative forms of knowledge, and these discussions generally seek to maintain a knowledge condition on intentional action. We argue that the fundamental lesson of our initial argument applies to these lines of thought. It remains possible that some understanding of practical knowledge could escape the lesson. But we fail to see such an understanding in the literature, or on the horizon.

Preamble concluded, we begin with Michael Thompson's (2011) perspective on practical knowledge, which he develops by interaction with Anscombe.

21. It bears mention that the locution 'practical knowledge' is—somewhat confusingly—deployed in both epistemology and action theory to pick out two arguably distinct phenomena (though, cf. Small 2012). That is, it is sometimes used to refer to knowledge of what we are doing, where the content of this knowledge is, as Thompson puts it, of something in progress—and so the knowledge in question consists in one's knowing that one *is doing something* (in progress). 'Practical knowledge' is also often used by philosophers of *know-how* (e.g., Stanley 2011) to refer to (non-deontic, infinitival) knowledge-how—viz., as in "Hannah knows how to ride a bike." While it is deeply contested in the know-how literature whether know-how patterns with knowledge-that when it comes to safety of the sort that is taken to be incompatible with knowledge-that (see, e.g., Poston 2009; Cath 2011; Carter and Pritchard 2015), *generally* (but see Cath 2015) both sides of this debate about practical knowledge-cum know-how accept that propositional knowledge must be safe in exactly the kind of sense in which it is *not* safe in Gettier cases (see Stanley 2011: ch. 8). This bears importance because 'practical knowledge' in the *former* sense—e.g., knowledge of what you are doing—is arguably just a special case of propositional knowledge *as it features* in a knowledge condition on intentional action. For example, what is at issue in Piñeros Glasscock's (2019) PKP is whether when one intentionally Fs, she knows *that she is F-ing* (intentionally and under that description). This is knowledge of a fact, and as such, it is subject to safety-based constraints in a way that applies to the wider genus (propositional knowledge) of which knowledge that one is intentionally F'ing is arguably a species. See Small (2012) however, for an attempt to make sense of the above two strands of practical knowledge as intimately connected.

Thompson takes the overarching thesis of Anscombe's *Intention* to be this: 'self-knowledge . . . extends beyond the inner recesses of the mind, beyond the narrowly psychical, and into the things that I am doing' (2011: 200). Further, Thompson reads Anscombe as affirming, and seems himself to endorse, an account of intentional action that links it closely with this knowledge: 'what's up with me is an intentional action precisely where it is a content of specifically practical knowledge, otherwise not' (2011: 203).

What is *specifically* practical knowledge? Thompson contrasts Davidson, who tends to speak of action in the past tense and in the perfective aspect (actions considered as completed wholes, so to speak), with Anscombe, who makes room for action in the present tense and in the imperfective, and progressive aspect (actions considered as ongoing processes). This is crucial, for Thompson. As he says, 'the content of Anscombe's practical knowledge is progressive, imperfective, in medias res. Its character as knowledge is not affected when the hydrogen bomb goes off and most of what the agent is doing never gets done' (2011: 209). And further, 'there is practical knowledge only when the thing is precisely NOT done, not PAST . . .' (2011: 209).

Thompson suggests that this emphasis on the progressive aspect 'might be enough to bring out the defect in a celebrated argument of Davidson's (2011: 209). Thompson observes that in the normal case, agents would simply keep working at getting the ten copies down until the job was done, and they would know what they are doing the whole time. Davidson's one-off case is an odd variant: 'Davidson's man . . . must be under some strange mafia threat' (2011: 210). But, one might think, so what? One-off cases remain legitimate. What are we to say of them? Thompson suggests that in this case,

the making of the inscription is like the buying of a lottery ticket. You can say he made ten copies intentionally if you like, but it will not be an illustration of the topic of Anscombe's book, any more than lottery-winning is when you bought the ticket with that aim. (2011: 210)

This is an interesting reaction, splitting as it does in two. The latter reaction seems to assert Thompson's view (and his view of Anscombe): practical knowledge is only knowledge of what is in progress, and thus what we should say in the case where the agent succeeds in making ten copies is beside the point—it is not a case of practical knowledge because practical knowledge does not concern completed actions. And one might think that this is not a view our case is designed to handle—for if the key question, for us, is whether Pujols's (completed) successes qualify as intentional actions, then aren't we stuck in the Davidsonian mode Thompson rejects?

But it is possible to consider Pujols-style cases in the Anscombian mode. As Pujols is in the process of swinging, we can ask—does he have specifically practical knowledge that he is getting a base hit? One thing we might want to know, in thinking through what to say, is what relationship is supposed to hold between what Pujols is in the process of doing and what might be the case after he is done. In other words—does it matter if Pujols has little chance of succeeding, or might we say that he has practical knowledge that he is getting a hit come what may?

Frost (2019) considers a ‘Thompson-style’ view that attributes practical knowledge come what may, and makes a strong argument that this could not have been Anscombe’s view. We agree with Frost, and we find such a view a non-starter. One reason is that, as Frost writes,

It’s plausible that knowledge of any kind involves non-accidental agreement of the mind of the knower with the object of knowledge. It is hard to see how there is such agreement if Anscombe’s knowledge is ‘the same’ even when she fails. (2019: 317)

A second is that practical knowledge is supposed to be at once knowledge of what I am doing as well as of what is happening. Consider the case in which Anscombe closes her eyes and intends to write ‘I am a fool’ on the chalkboard, but in which her performance is flawed. As Frost notes, ‘when she makes the relevant mistake in performance, the claim “I am writing ‘I am a fool’ on the board with my eyes shut” is *false*, because she is not achieving, and has not achieved, any of those results’ (2019: 318).

If the relationship between what the agent is doing and what might be the case after they are done does not matter, then why equate the writing of Davidson’s man with the lottery ticket buyer? That Thompson does so suggests there must be some relationship in play. And it is this relationship that Will Small (2012) elucidates in an interesting and useful way.

Small, for the record, is a forceful proponent of a knowledge condition—what he calls a cognition condition—on intentional action. As he has it,

Once we have uncovered the real ground and meaning of the cognition condition, which lies in the calculative and temporal structure of intentional action, we will see that the carbon-copier example poses no threat, because it trades on confusion as to how to understand the possibility of failing to act as one had intended. (2012: 142)

The calculative structure of intentional action, for Small, is one element that Anscombe emphasizes, and Davidson misses. Here he agrees with Thompson:

a focus on the internal, rational structure of action, an order of relationships between means and ends that the agent represents and works out via practical reasoning, as the agent progressively and rationally respecifies their intention, and executes their action, is crucial.²²

This calculative structure lends the form of the knowledge condition that Small endorses. The agent in the midst of doing something is ‘engaged in a stretch of calculatively-articulated intentional activity’ and this agent knows what they are doing ‘because he understands how and why the elements of his action, which are themselves intentional actions, combine such as to amount to him intentionally producing the whole thing’ (2012: 164).²³ Small summarizes the condition as follows.

[I]f you are acting intentionally, you have knowledge in intention of:

- (i) what you are doing (e.g. replenishing the house water-supply),
- (ii) why you are doing it (e.g. because you’re poisoning the inhabitants),
- and (iii) how you are doing it (e.g. by operating this pump). (2012: 164)

How does this condition hold up in the face of problem cases? Look first at the carbon-copier case, and focus on the imperfective, progressive aspect. Could it be the case that the carbon copier is intentionally making the copies, even though he does not know he is doing so? Small thinks not, and his interesting claim here is that ‘the cognition condition on intentional action (in progress and in prospect) must encompass knowledge of eventual success, properly understood’ (2012: 177). This link between current activity and eventual success helps fill out the picture Thompson attributes to Anscombe.²⁴ So, what gives the agent knowledge of eventual success?

Here Small brings in the notion of an *accident*. For the agent to know that they will succeed, it must be no accident if they do, and ‘some accident’ if they do not

22. We can agree with all of this, and indeed, we find it insightful. This quote, from Small, is an interesting and exciting point, inviting further reflection on rational respecification.

[T]here is a kind of constant correction that goes on throughout the course of an intentional action, as the agent responds to the miniature successes and failures, obstacles and alternative possibilities both foreseen and unforeseen, that he encounters in what he’s doing and what he’s acting on and with. What I am bringing out is that this constant correction amounts to rationally respecifying one’s intention: the answer to the question ‘How?’ is constantly being finessed. (2012: 158)

23. Small speaks here of ‘knowledge in intention,’ as though the intention were the vehicle for the agent’s knowledge. That might be a positive view similar to Campbell’s (2018). But Small’s intention is not a static state. It is rather something being rationally respecified, suggesting that this activity of constant rational respecification is doing much of the work for Small. How to marry these things within Small’s framework is not entirely clear to us.

24. Indeed, Small presents an interesting argument that if we do not link the doing of B to the success of having done B, we end up with skepticism ‘about the very idea of an event’s being in progress’ (2012: 179).

(2012: 173). A part of the story here involves an appeal to knowledge-how. The agent who knows what they are doing is bolstered by their knowledge of how to do it, and their knowledge of the relationship between what they are doing and success. We thus get what sounds like an addition to the knowledge condition, for Small claims that for an agent to be doing B intentionally, they must know not only that they are doing B intentionally, but also that their eventual success will be no accident.

But now this account meets Pujols. Or it meets any of a number of cases involving mid-level degrees of success. In these cases the agents sometimes succeed, and sometimes fail. Can it be said that it will be some accident if they fail? We would suggest not. Can it be said that they know that it will be no accident if they succeed? We would suggest not. But can it be said, contrary to our initial intuition, that agents in these cases are not really intentionally A-ing, given the imperfect link between what they are doing and success? We would suggest not.

This may seem like a surface-level clash of intuitions. Pretty disappointing, after how far we've come. Fortunately, though, we can dig a little. It seems to us that what is really going on here is a difference in framework that narrows at just these kinds of cases, building some pleasant theoretical pressure. The choice is between a framework that makes room for fineness of grain regarding an agent's degree of control, and thus for lots of shading around notions of success and failure, and a framework that emphasizes—for reasons well worth exploring, although not so much here (though see Section 4)—a brighter line between success and failure.

Recall Thompson's insistence that the carbon copier is buying a lottery ticket. That's a telling comparison. (And look, if the lottery gives us the chances the carbon copier has, we're in.) Similarly, Small compares the copier to a slow, out-of-shape agent attempting to run a four-minute mile. Is there really no room in between those cases? Of course, we think there is, and we think the cases we have emphasized bring this room out.

But before we press this perceived advantage, let us briefly consider a development of the framework we oppose that might have something to say. This is due to Frost (2019), who develops his account of practical knowledge in part by considering, side-by-side, a failure and a success. He draws on a well-known passage in Anscombe (2000: 82), in which the agent closes her eyes and writes 'I am a fool' on a blackboard. Somewhat confusingly, Anscombe remarks:

. . . intentions [can] fail to get executed. That intention for example would not have been executed if something had gone wrong with the chalk or the surface, so that the words did not appear. And my [practical] knowledge would have been the same even if this had happened. If then my knowledge is independent of what actually happens, how can

it be knowledge of what does happen? Someone might say that it was a funny sort of knowledge that was still knowledge even though what it was knowledge of was not the case! On the other hand Theophrastus' remark holds good: 'the mistake is in the performance, not in the judgment'. (Anscombe 2000: 82)

One of Frost's innovations is to interpret Anscombe's appeal to Theophrastus as the appeal of a partially mistaken interlocutor. So Frost's Anscombe does not hold that practical knowledge is present alike in cases of success and failure. But the interlocutor gets one thing right.

What does the interlocutor get right? Anscombe's practical knowledge 'would have been the same' in one sense, because, as I mentioned before, practical knowledge can be construed, without much distortion, as a *capacity*. Like capacities generally, the capacity that is practical knowledge is independent of any particular exercise of the capacity. Consider salt's capacity to dissolve in water. This capacity is independent of whether a given sample of salt ever exercises the capacity in question. Similarly for practical knowledge. The capacity to write 'I am a fool' etc. such that the action exhibits an order of practical reasoning is independent of whether that capacity is exercised (perfectly or imperfectly) on some occasion. (Frost 2019: 326)

This distinction between perfect and imperfect exercises of the capacity (or capacities) for practical knowledge—which is, for Frost, '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' (2019: 326)—is key. For human agents sometimes make mistakes.²⁵

Now, certainly in some cases an agent's failure can be blamed on a mistake. But one point of the Pujols-style cases we have raised is that human agents sometimes know exactly how to behave, and make no specific mistake, and still fail. Sometimes they behave in indistinguishable ways, and succeed. (And sometimes—indeed, we would say most of the time—human agents behave imperfectly, but there is room for error, and they succeed.) This is a part of a commonsense understanding of some activities as chancy. And if the modal space allows for this, then action theory should be sensitive to it. One option is to insist that knowledge and action are always aligned—revise judgments of intentional action in conditions

25. Small also invokes a distinction between perfect and flawed exercises of a capacity:

that a skilled golfer might have missed her putt, placing a birdie irrevocably beyond reach, does not impugn the status of the putt that she holes through the exercise of her skill: that her capacity is fallible does not mean that every exercise of it is flawed (2012: 202).

a bit too chancy to allow for knowledge. This is a bad move, in our view. For the deep truth we are tracing in this paper is that the notion of intentional action is permissive of risk, as it should be if our action theory is based upon the animal rather than the angel, that is, upon a fundamentally fallible class of agent.

It follows that if we are not willing to revise the notion of knowledge, the relative difference in permissiveness regarding the possibility of failure—even in the absence of mistake—delivers the result that knowledge and intentional action come apart in a key class of cases. This key class falls in between the cracks if our theory of action only admits certain successes and lottery-level accidents. We are suggesting that human action often inhabits the moderately risky spaces in between.²⁶

Poetic conclusion! But unfortunately we are not done yet. For a new approach to practical knowledge seems designed to accommodate our perspective, while maintaining a knowledge condition on intentional action. This approach appeals to Sarah Moss's (2018) recent work on probabilistic knowledge. First, a brief set-up.

Moss (2018) defends the idea that probabilistic cognitive states—credences, partial beliefs, probabilistic beliefs—can qualify as knowledge. So, to take just one example (and to be quick about it), an agent might believe that it is likely that they will succeed. Moss argues that this belief can be true if it is, indeed, likely that they will succeed. And this belief can be safe if the agent was not likely to have a false belief regarding their likelihood of success.

Let us assume—with the note that Moss argues extensively for this in her book—that such beliefs can qualify as knowledge. If so, there is room for an

26. This point, and our arguments in this paper, leave open the possibility of a fallback, namely, that intentional action requires that the agent have knowledge of what they are doing under some description. A number of authors who reject a strong knowledge condition on intentional action have nonetheless endorsed a version of this option (Davidson 1980; Setiya 2007). Now, while we regard this option as a significant concession, one might still think the option is important. As a referee notes, this option keeps live the idea that action is necessarily self-conscious in that it involves some self-knowledge of what is being done. And this option may be attractive to those who, like Frost, hold that 'the question "What is *the* content of practical knowledge?" is a bad question, because this question presupposes that practical knowledge has one (unique) content in every case of intentional action' and who also hold that 'A better question would be "Which *contents* are (contents of) apt expressions of practical knowledge in various cases?"' (2019: 330)

Unfortunately, considering the idea that that intentional action requires that the agent have knowledge of what they are doing under some description in depth will require much more space, for one needs to think through views on the proper vehicle(s) for knowledge (whether belief, or intention, or a *sui generis* state), and whether such vehicles are necessary for intentional action. And one needs to set up and think through a range of problem cases, for example, cases that involve highly habituated behavior, cases that involve skilled and rapid reactions to changing circumstances, cases that involve guidance by sensorimotor mechanisms that are relatively inaccessible to awareness, or guidance by (arguably) unconscious perception. Perhaps, with the stronger version of a knowledge condition defeated, theorists will be able to see that such cases deserve more philosophical attention.

argument that the agent's knowledge of what she is doing can take the form of probabilistic knowledge.

Carlotta Pavese (2020) makes this argument, and insists that probabilistic knowledge can save a knowledge condition in the face of difficult cases. Usefully for our purposes, she considers two kinds of cases of middle-grade control. The first involves archery.

Consider Archie who attempts to shoot a target. An expert archer such as he is, he does what he would do in any similar situation in which the wind is quiet. But suppose a gust of wind could easily divert his shot, although it does not. His shot can be intentional, it seems, and he might even know how to hit the target, even though his belief about how he can hit the target in that circumstance is unsafe. After all, the way he shot in that circumstance might easily not have been successful, if the gust of wind had intervened. (Pavese 2020: 352)

Pavese counsels a move away from full to probabilistic knowledge in this case.

Probabilistic knowledge to the rescue! The proponent of probabilistic knowledge could object that Archie's belief is probabilistic: it is the probabilistic belief that shooting in a certain way is sufficiently likely, given certain normal circumstances, to lead to successfully hitting the target. We might think of this probabilistic belief as a conditional credence: a credence that assigns sufficiently high probability to hitting the target by shooting in a certain way, conditionally on certain normal circumstances being in place. This conditional credence can be safe in this circumstance, and even be knowledge, for in normal circumstances it would not easily be the case that shooting in that way would not likely be successful. (2020: 352)

Such reasoning motivates a condition on intentional action that holds that 'If S successfully and intentionally fs at t, then at t s knows, for some means y of f-ing, that oneself is sufficiently likely to f by y-ing' (2020: 353). Pavese notes that what counts as sufficiently likely may vary according to the task and the circumstances, and that seems fair (see fn. 9). But what do we think about this condition on knowledge?²⁷

27. Note that we have moved some distance from Anscombe-inspired the views on practical knowledge that view it as a distinct form of knowledge. This form of knowledge is, arguably, a sort of contemplative knowledge. Here we consider it on the terms Pavese offers. One might, however, consider whether a view like Frost's could make room for probabilistic practical knowledge. This would be practical knowledge as a probabilistic capacity or pair of capacities to increase one's chances at doing something, and to have probabilistic knowledge of particular actions. This view

We think that even if one grants the possibility of probabilistic knowledge—and the claim that one can have probabilistic knowledge of what one is doing—its necessity for intentional action is dubious.

Consider Ticha, the pessimistic basketball player. Ticha significantly under-rates herself and her chances, even though she is quite a good shooter. So she systematically forms beliefs about her chances that are false, believing that success is unlikely when it is likely. When Ticha lines up a shot that has—if the percentages are to be believed—a 50% chance of success, Ticha believes that the chances are closer to 25%. Ticha makes the shot, say. Was Ticha intentionally making the shot, and did she intentionally make it? Plausibly, yes. Did Ticha have probabilistic knowledge along the way? Plausibly, no, since it seems her probabilistic belief was false.

If one worries that beliefs about the unlikelihood of success undermine intentionality, reverse the case. Marcus is an ‘irrational confidence guy’ (for many basketball fans, this is an actual type, with many funny examples). What kind of guy is this? Marcus makes around 35% of his shots, but he believes that every one is going in, (and is happy to tell you this, and to behave in ways that confirm this). Marcus sometimes makes shots intentionally, but his belief is false.

The conclusion of Section 3 is thus similar to the conclusion of Section 2—the argument against the necessity of knowledge for intentional action, even of a specifically and epistemically unique practical knowledge, looks increasingly robust.

4. The Essence of Intentional Action

We have argued that knowledge of what the agent is doing is not necessary for that doing’s being an intentional action. Others have recently done so as well, but there is an interesting divergence here that merits discussion. We saw above that Schwenkler argues that while Anscombe’s claim is falsified ‘in full generality,’ its spirit survives *because* ‘the circumstances in which [Anscombe’s claim] . . . is not satisfied are decidedly unusual: it takes some philosophical ingenuity to construct them, and we don’t often encounter situations like these in ordinary life’ (2019: 190). What does it mean to say that the spirit of Anscombe’s claim survives? Schwenkler writes this: ‘self-knowing action is the *paradigmatic* form of intentional agency, so that when someone acts intentionally but without knowledge of her action there must be some special circumstance that accounts for why this is’ (2019: 190).

would need filling out, but an initial worry is that this view needs an explanation of the tight concordance between one’s probabilistic capacity to do things and one’s probabilistic knowledge, and it is not clear one is available. That is to say, our irrational confidence example would seem to be a problem for this view as well.

This is an interesting claim, and it suggests a line of thought one finds in other theorists as well as Schwenkler (e.g., Velleman 1992; Setiya 2007; Gibbons 2010; Wolfson 2012; and to some extent Piñeros Glasscock 2020). Put abstractly, the idea is that while knowledge may not be necessary for intentional action, it is intimately connected to it *in some way*. How to spell this out in full is a nice topic for a different paper (or book). Here we can do little more than sketch options, hoping to instigate further reflection down the line.

One option is to develop a view on which the knowledge-intentional action connection is in some way normative. Piñeros Glasscock, for example, suggests (without fully endorsing) something reminiscent of Velleman (1992)²⁸—that ‘the function of the will (understood as the capacity to act intentionally) is to yield practical knowledge’ (2020: 1262). On this kind of view, intentional action without knowledge is possible, but normatively defective, in the sense that it is a case of the will failing to perform its function. We are not attracted to this view, but it is certainly a live possibility.

A second option is to develop a view that holds that in spite of exceptional cases, knowledge is a part of the essence of intentional action (where essential properties, obviously, are not the same thing as properties had necessarily).²⁹ This is one way to understand an implication of generic statements like ‘cats have tails’ or ‘winners never quit.’ Some cats lack a tail, and winners occasionally quit, but these exceptions do not falsify the generic statement. Perhaps the relationship between knowledge and intentional action is something like this. As Bailey and van Elswyk observe,

[Generics] add explanatory depth to generalizations. In contrast to a quantifier phrase like ‘All *F*s are *G*’, which merely states the quantity of *F*s within a domain that are also *G*, ‘*F*s are *G*’ conveys that there is a special or intimate connectedness between being *F* and being *G*. (Bailey & van Elswyk in press)

Now, getting the semantics of generic statements right is difficult—a number of competing proposals remain afloat. But one plausible thought is to construe generics in terms of normalcy or typicality (Asher & Morreau 1995; Eckardt 2000; Nickel 2008; 2016; Bailey & van Elswyk in press). On such a view, again quoting Bailey and van Elswyk, ‘A generic is about the normal, typical, or characteristic *F*s. So being of kind *F* carries a special connectedness to being *G* because to be a normal *F* is to be *G*’ (in press).

28. That is, Velleman’s explanation for the pervasiveness of knowledge of action in terms of intrinsic desires to know what one is doing, which in turn lead one to act in ways that generate this knowledge.

29. Thanks to a referee for a suggestion along these lines.

In one way, construing a knowledge-intentional action connection in terms of normalcy or typicality pushes the issue back a step. For what is it to be a normal or characteristic instance of a kind? This turns out to be complicated. Proponents of normalcy-based accounts of generics do not universally agree. Nickel points out that, plausibly, there may be ‘cases in which there are many, equally good ways of being normal’ (2008: 643).

But the suggestion may be fruitful, for at least two reasons. First, and just as a kind of technician’s point, the literature on generics may offer us new connections, and new ways of thinking about the senses in which knowledge may or may not be typical of, or otherwise intimately connected to, intentional action. Second, and perhaps more importantly, this construal of the knowledge-intentional action relationship suggests a pressing need to connect our reflections on the relationship between knowledge and intentional action to broader currents. For an account of the essence of intentional action will, most likely, have to confront difficult issues about the nature of agency and the various (e.g., practical, moral, aesthetic) normative structures that inform our assessments of agency and action.

This is no small task, and from where we sit it is unclear how exactly this set of reflections might go, in part because there are many theoretical options available, and many perspectives on the nature of agency already in the background, at various levels of articulation.

As a way out, allow us to offer—with apologies, since there is clearly room for nuance and additional qualification here—an imagined contrast between a perspective on the essence of intentional action that sees knowledge of what one is doing as of the essence of what one is intentionally doing, and a perspective that sees the opposite, perhaps replacing knowledge of what one is doing with controlled execution of what one is doing as of the essence. One can imagine the knowledge perspective lining up more cleanly with a view of the nature of agency that is more angelic, emphasizing powers of rationality and the importance of self-consciousness, and viewing the typical case of intentional action as one in which the agent’s success is very close to guaranteed, is the result of the perfect exercise of some agentive capacity. It is no secret that Aquinas’s view of practical knowledge was a major influence upon Anscombe’s, and Schwenkler (2015) notes that Aquinas³⁰—along with a range of important medieval thinkers—deploys the notion of practical knowledge primarily as a way of capturing the kind of knowledge God has of his creation. As Schwenkler puts it,

This idea, that one may know an object not just by having one’s mind reflect what is anyway there, but also by being the one who through her

30. One might develop a version of this perspective independently of Aquinas, of course: perhaps along Hegelian lines (Taylor 2010; Rödl 2018).

knowledge brings that very object into existence, is at the center of Aquinas's conception of practical cognition. (2015: 11)

To make that kind of knowledge essential for intentional action is to take the angelic perspective on the nature of agency and action.

By contrast, one can imagine that the perspective that sees knowledge as inessential to intentional action lines up more cleanly with a view of the nature of agency that is more animal, emphasizing the limits of our powers of execution, planning, and perception, and thus emphasizing the need for agency to involve a special kind of mental structure, a structure of specific sorts of attitudes (as in, e.g., Castañeda 2012 or Brand 1984) as well as a layered cognitive architecture (as in Mele 1995 or Pacherie 2008), and in addition a range of tricks and techniques (not only the formation of habits and 'motor schemata' but the capacity to abut one's fallibility by strategic planning [Bratman 2018], including strategic back-up planning [Paul 2021]), ways to raise our chances of success in an unfriendly world. This contrast is not the common one between a Davidsonian and an Anscombian perspective on agency. For the current perspective is free to see practical knowledge as critical for agency, perhaps by treating knowledgeable action as a sub-class of intentional action (as Shepherd 2021 does explicitly), and thinking of practical knowledge as something that enhances an agent's control, or an agent's chances of success. And such a perspective is free to take on board many Anscombian points, for example the points discussed above regarding the calculative and temporal structure of action. One key difference between the animal and angelic perspectives will be the former's emphasis on imperfection, leading to a tolerance of risk and the absence of knowledge in many cases, and an emphasis on degrees of control and the chancier elements of human activity, as critical to understanding agency.

5. Conclusion

We have argued that any version of a strong knowledge condition on intentional action, on which an agent A-s intentionally only if they know that they are A-ing, is false. One reason is that intentional action and knowledge have different levels of permissiveness regarding failure in similar circumstances.

Our argument runs through a kind of case that is common, and that centers features of agentic control in a way that Davidson's carbon copier case does not, and in a way that recent discussions of the relation between knowledge and intentional action do not fully appreciate. We have argued that this has implications for broader claims about the place of knowledge in our understanding of intentional action.

- Carter, J. Adam, Benjamin W. Jarvis, and Katherine Rubin (2015). Varieties of Cognitive Achievement. *Philosophical Studies*, 172(6), 1603–23.
- Castañeda, Héctor-Neri (2012). *Thinking and Doing: The Philosophical Foundations of Institutions* (Vol. 7). Springer Science & Business Media.
- Cath, Yuri (2011). Knowing How Without Knowing That. In J. Bengson and M. A. Moffett (Eds.) *Knowing How: Essays on Knowledge, Mind and Action* (113–35). Oxford University Press.
- Cath, Yuri (2015). Revisionary Intellectualism and Gettier. *Philosophical Studies*, 172(1), 7–27.
- Davidson, Donald (1980). *Essays on Actions and Events*. Oxford University Press.
- Eckardt, Regine (2000). Normal Objects, Normal Worlds, and the Meaning of Generic Sentences. *Journal of Semantics*, 16(3), 237–78.
- Falvey, Kevin (2000). Knowledge in Intention. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 99(1), 21–44.
- Frost, Kim (2019). A Metaphysics for Practical Knowledge. *Canadian Journal of Philosophy*, 49(3), 314–40.
- Gibbons, John (2010). Seeing What You’re Doing. *Oxford Studies in Epistemology*, 3, 63–85.
- Ginet, Carl (1975). *Knowledge, Perception, and Memory*. D. Reidel.
- Greco, John (2010). *Achieving Knowledge*. Cambridge University Press.
- Grünbaum, Thor (2011). Perception and Non-Inferential Knowledge of Action. *Philosophical Explorations*, 14(2), 153–67.
- Kelp, Christoph (2009). Knowledge and Safety. *Journal of Philosophical Research*, 34, 21–31.
- Lewis, David (1996). Elusive Knowledge. *Australasian Journal of Philosophy*, 74(4), 549–67.
- Ludwig, Kirk (1995). Trying the Impossible: Reply to Adams. *Journal of Philosophical Research*, 20, 563–70.
- Luper-Foy, Steven (1984). The Epistemic Predicament. *Australasian Journal of Philosophy*, 62(1), 26–50.
- Lyons, Jack C. (2019). Algorithm and Parameters: Solving the Generality Problem for Reliabilism. *Philosophical Review*, 128(4), 463–509.
- Marcus, Eris (2018). Practical Knowledge as Knowledge of a Normative Judgment. *Manuscrito*, 41, 319–47.
- McKenna, Robin (2014). Normative Scorekeeping. *Synthese*, 191(3), 607–25.
- Mele, Alfred R. (1995). *Autonomous Agents: From Self-Control to Autonomy*. Oxford University Press.
- Mele, Alfred R. and Paul Moser (1994). Intentional Action. *Noûs*, 28(1), 39–68.
- Moran, Richard (2004). Anscombe on ‘Practical Knowledge’. *Royal Institute of Philosophy Supplements*, 55, 43–68.
- Moss, Sarah (2018). *Probabilistic Knowledge*. Oxford University Press.
- Nickel, Bernhard (2008). Generics and the Ways of Normality. *Linguistics and Philosophy*, 31, 629–48.
- Nickel, Bernhard (2016) *Between Logic and the World: An Integrated Theory of Generics*. Oxford University Press.
- Pacherie, Elisabeth (2008). The Phenomenology of Action: A Conceptual Framework. *Cognition*, 107(1), 179–217.
- Pavese, Carlota (2020). Probabilistic Knowledge in Action. *Analysis*, 80(2), 342–56.
- Paul, Sarah K. (2009). How We Know What We’re Doing. *Philosopher’s Imprint*, 9(11), 1–24.

- Paul, Sarah K. (2021). Plan B. *Australasian Journal of Philosophy*. <https://doi.org/10.1080/00048402.2021.1912126>
- Pickard, Hannah (2004). X—Knowledge of Action without Observation. *Proceedings of the Aristotelian Society*, 104(1), 205–30.
- Piñeros Glasscock, Juan S. (2020). Practical Knowledge and Luminosity. *Mind*, 129(516), 1237–67.
- Poston, Ted (2009). Know How to Be Gettiered? *Philosophy and Phenomenological Research*, 79(3), 743–47.
- Pritchard, Duncan H. (2002). Resurrecting the Moorean Response to the Sceptic. *International Journal of Philosophical Studies*, 10(3), 283–307.
- Pritchard, Duncan H. (2005). *Epistemic Luck*. Oxford University Press.
- Pritchard, Duncan H. (2007). Anti-Luck Epistemology. *Synthese*, 158(3), 277–97.
- Pritchard, Duncan H. (2008). Sensitivity, Safety, and Anti-Luck Epistemology. In J. Greco (Ed.), *The Oxford Handbook of Scepticism* (437–55). Oxford University Press.
- Pritchard, Duncan H. (2010). Relevant Alternatives, Perceptual Knowledge and Discrimination. *Noûs*, 44(2), 245–68.
- Pritchard, Duncan H. (2012). Anti-Luck Virtue Epistemology. *The Journal of Philosophy*, 109(3), 247–79.
- Pritchard, Duncan H. (2016). Epistemic Risk. *The Journal of Philosophy*, 113(11), 550–71.
- Rabinowitz, Dani (2011). The Safety Condition for Knowledge. *Internet Encyclopedia of Philosophy*. Retrieved from <https://iep.utm.edu/safety-c/>
- Riggs, Wayne (2009). Luck, Knowledge, and Control. In Pritchard, Haddock and Millar (Eds.), *Epistemic Value* (204–21). Oxford University Press.
- Rödl, Sebastian (2018). *Self-Consciousness and Objectivity*. Harvard University Press.
- Sainsbury, Richard Mark (1997). Easy Possibilities. *Philosophy and Phenomenological Research*, 57(4), 907–19.
- Schwenkler, John (2015). Understanding ‘Practical Knowledge’. *Philosophers Imprint*, 15(15), 1–32.
- Schwenkler, John (2019). *Anscombe’s Intention: A Guide*. Oxford University Press.
- Setiya, Kieran (2007). *Reasons without Rationalism*. Princeton University Press.
- Setiya, Kieran (2008). Practical Knowledge. *Ethics*, 118(3), 388–409.
- Shepherd, Joshua (2021). *The Shape of Agency: Control, Action, Skill, Knowledge*. Oxford University Press.
- Smith, Martin (2017). *Between Probability and Certainty: What Justifies Belief*. Oxford University Press.
- Small, Will (2012). Practical Knowledge and the Structure of Action. In Günter Abel and James Conant (Eds.), *Rethinking Epistemology* (133–228), Vol. 2 of Berlin Studies in Knowledge and Research. De Gruyter.
- Sosa, Ernest (1999). How to Defeat Opposition to Moore. *Philosophical Perspectives*, 13, 141–54.
- Sosa, Ernest (2010). How Competence Matters in Epistemology. *Philosophical Perspectives*, 21(1), 465–75.
- Sosa, Ernest (2015). *Judgment & Agency*. Oxford University Press.
- Sosa, Ernest (2017). *Epistemology*. Princeton University Press.
- Stanley, Jason (2011). *Know How*. Oxford University Press.
- Stathopoulos, Alexander (2016). Knowing Achievements. *Philosophy*, 91(3), 361–74.

- Stout, Rowland (2019). Practical Reasoning and Practical Knowledge. *Canadian Journal of Philosophy*, 49(4), 564–79.
- Taylor, Charles (2010). Hegel and the Philosophy of Action. In Arto Laitinen and Constantine Sandis (Eds.), *Hegel on Action* (22–41). Palgrave-MacMillan.
- Thalberg, Irving (1962). Intending the Impossible. *Australasian Journal of Philosophy*, 40(1), 49–56.
- Thompson, Michael (2011). Anscombe's *Intention* and Practical Knowledge. In Anton Ford, Jennifer Hornsby, and Frederick Stoutland (Eds.), *Essays on Anscombe's Intention* (p–p). Harvard University Press.
- Vekony, Romy, Alfred Mele, and David Rose (2020). Intentional Action without Knowledge. *Synthese*, 199, 1231–43.
- Velleman, David (1992). What Happens when Someone Acts. *Mind*, 101(403), 461–81.
- Williamson, Timothy (2000). *Knowledge and Its Limits*. Oxford University Press.
- Wolfson, Ben (2012). Agential Knowledge, Action and Process. *Theoria*, 78, 326–57.
- Wu, Wayne (2016). Experts and Deviants: The Story of Agentive Control. *Philosophy and Phenomenological Research*, 92(2), 101–26.