

Intentional Action and Skill Proximity

Abstract

According to one popular analysis of intentional action that we dub the "Simple Knowledge View," knowing how to ϕ is necessary for intentionally ϕ -ing; absent such knowledge, one's successful ϕ -ings would involve too much luck and too little control. But the Simple Knowledge View is far too restrictive; it excludes from the class of intentional actions those performed in novel circumstances by co-opting proximal skills. Here, we articulate a theory of skill proximity and motivate a more cautious analysis of intentional action, the "Skill Proximity View." According to it, intentionally ϕ -ing only requires that the agent know how to ψ , where ψ is sufficiently skill-proximal to ϕ . Our view strikes a plausible balance between explaining the control characteristic of intentional action, on the one hand, and accommodating the luck that such action plausibly tolerates, on the other.

Word Count: 9,999

Keywords: Knowledge-How
Intentional Action
Skill
Control
Luck

Few theses in action theory have more adherents than the Simple Knowledge View,¹ hereafter referred to as “SIMPLE”:

SIMPLE: If S ϕ -s intentionally, S knows how to ϕ .²

But SIMPLE, in all its elegance and parsimony, is false. The aim of this paper is to present a counterexample to it (Sec. 1); and then motivate a more cautious and theoretically preferable position, the Skill Proximity View, hereafter referred to as “SKILL” (Secs. 2-3), which, as a first approximation to be developed, states:

SKILL: If S ϕ -s intentionally, S knows how to ψ , where ψ is sufficiently skill-proximal to ϕ .

SKILL is more cautious than SIMPLE in the sense that it is a strictly weaker position; as we will go on to explain, for an agent who knows how to ψ , ψ will always be sufficiently skill-proximal to itself. And SKILL is theoretically preferable to SIMPLE for three, related reasons (Sections 4-5). First, SKILL avoids the counterexamples to SIMPLE. Second, SKILL is better positioned to explain the phenomenon of *learning how*. Third, SKILL brings into relief certain nuances in the relationship between the various forms of control characteristic of intentional action and the kind and degree of luck such forms of control can tolerate.

1. SIMPLE

1.1. Intentional action requires know-how: some motivations

The motivations for the SIMPLE primarily concern a theoretical need to rule out cases of accidental or too-lucky successes as intentional.

Hawley (2003: p.27) offers the following three widely cited cases taken to support this rough but intuitive idea:

Avalanche: Sally is caught out in an avalanche. Having no idea of how to escape them, but mistaking the incoming snow for water, she begins to make swimming motions, frantically paddling her way to safety.

¹ So-called because of the resemblance it bears to the “Simple View” of intentional action attacked by Bratman (1984, 1987), according to which ϕ -ing intentionally requires an intention to ϕ . See also di Nucci (2009, 2010), McCann (2010, 2011), and Amaya (2018) for extended discussion.

² A partial list of proponents includes Ryle (1949); Stanley & Williamson (2001); Hawley (2003); Hornsby (2004, 2011); Stanley (2011); Setiya (2012); Pavese (2018, 2020, 2021a, 2022); Carter and Navarro (2017), Shepherd and Carter (2021), and Beddor and Pavese (2022). Some authors (e.g., Hawley 2003; Setiya 2012) endorse a biconditional version of this principle (regarding what S is able to do intentionally, rather than what S does intentionally), but I will set aside that stronger idea here.

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

Cake: Shelley, having no clue how to make a cake under normal circumstances, attempts to make one by putting together everything that happens to be in her kitchen cabinets. Luckily for her, her cabinets contain precisely the right ingredients (and in the right proportions) for her to succeed.

In each case, there is both a failure to act intentionally and a lack of knowledge-how that explains it.³ In **avalanche**, Sally survives by sheer dumb luck, mistaking snow for water and exercising her swimming know-how. Her failure to escape intentionally is explained by her lack of knowledge of how to escape avalanches; her knowledge of how to swim seems more or less irrelevant.⁴ In **annoyance**, since Susie misconstrues the situation, she plausibly doesn't know how to annoy Joe, and this explains her failure to do so intentionally. And in **cake**, Shelley's circumstances conspire so as to guarantee her success through haphazard methods, ones that would fail in almost any other kitchen.⁵ That Shelley doesn't know how to bake a cake is evident once we consider the culinary monstrosities that would spawn from counterfactual kitchens with different ingredients.

These verdicts are more or less anchor points for many theorists of action, and so is SIMPLE, which plausibly explains them. Again, according to this principle:

SIMPLE: If S ϕ -s intentionally, S knows how to ϕ .

Because Sally, Susie, and Shelley lack knowledge of how to do what they succeed in doing, their successes are not attributable to them as intentional actions. They are, in a sense that SIMPLE purports to capture, defective for involving too much luck and too little control. What about cases in which one knows how to ϕ and successfully ϕ s, but not by manifesting one's knowledge of how to ϕ ? Consider:

Sibling: Sally's sibling Sam is likewise caught out in an avalanche. Unlike Sally, Sam knows how to escape avalanches (she just took an extensive winter survival

³ It's worth noting that these examples also lend intuitive support, independently, to the idea that ability possession is insufficient for know-how – a point that is generally accepted by both intellectualists as well as anti-intellectualists about know-how. For discussion, see Pavese (2021b, §5).

⁴ Bengson and Moffett (2012) explain Sally's failure to act intentionally in terms of her having an inadequate grasp of the fact that swimming motions are a way to escape avalanches.

⁵ For a structurally similar case to Hawley's (2003) **cake** case, see Sosa's (1997, 418) tomato ripening case.

training course), but mistaking the incoming snow for water, she begins to make swimming motions, frantically paddling her way to safety.

To the extent that one thinks Sally merely unintentionally escapes the avalanche, one should think that Sam likewise merely unintentionally escapes the avalanche. After all, Sam's know-how played no role in her successful escape.⁶ And yet SIMPLE, taken on its own, does not predict that Sam's action was defective in any way. We think that the intuitive appeal of SIMPLE really rests on the plausibility of a deeper explanatory connection between intentional action and know-how, namely that know-how is necessary for intentional action because *manifesting*—not just possessing—the relevant know-how is control-securing and luck-precluding. Thus, if SIMPLE is understood as merely imposing a necessary condition on action, it is silent on **sibling**, while if SIMPLE is understood in more robust, explanatory terms, it correctly predicts **sibling** is on a par with **avalanche**. We think the second way to understand SIMPLE is more apt, but nothing below hangs on this (although we discuss it further in Section 4).

Being able to adequately characterize and explain the incongruity between the control that practical agents exhibit over their behavior so as to render it intentional, on the one hand, and the luck with which they succeed, on the other, is among the chief *desiderata* of any plausible theory of action. Luck is, in some intuitive sense, control-*diminishing*, and control is likewise luck-*excluding*. Consider:

Darts: Brimming with the confidence of a few beers, Tim signs up to play his first ever game of darts, fully intending to hit the bullseye on every throw. On his first turn he looks at the bullseye, makes some arm movements that, for all he knows, resemble the arm movements of genuine dart players, and sends the dart on a wing and a prayer. Lo and Behold! Tim hits the bullseye. As it turns out, he goes on to whiff every subsequent throw and loses badly.

Darts is not a case in which Tim intentionally hits the bullseye; it is not even a case in which Tim intentionally hits the dart board!⁷ There is, to be sure, more agency to Tim's hitting the bullseye in **darts** than there would be if Tim had only hit the bullseye after sneezing while noting a dart's heft. SIMPLE explains why: Tim doesn't know how to do what he succeeds in doing. The best we can say, perhaps, is that he intentionally *tried*—something he might not have done had Tim believed (bizarrely) the feat to have been

⁶ Indeed, Sam's know-how played no more role in her success than does Pavese's (2022, 375) 'unlucky clown', who knows how to do a certain kind of tumble but then performs the tumble only through an accidental trip. We return to Pavese's own positive view in §3.1.

⁷ Note that this assessment does not concern whether Tim was able to form the relevant intention here, only whether Tim counts as hitting the bullseye (or dart board) intentionally.

impossible.⁸ But there remains a gap between intentionally trying⁹ to hit the bullseye and intentionally hitting the bullseye.¹⁰

To fill in the details a bit, we can even suppose that Tim, despite never having played darts himself, had watched a few games played on television, so he was not totally in the dark as to what a proper throw looked like; he could at least spot good technique as well as any living-room coach. So we needn't think that Tim was entirely self-deceived or otherwise mistaken (at least not in the way that Sally was) about how one generally goes about hitting bullseyes.

Notice that SIMPLE explains **darts** even with these added details; on the plausible assumption that one doesn't know how to do something just because one has an ability to recognize it done well,¹¹ Tim's failure to hit the bullseye intentionally is plausibly explained in terms of his lack of knowledge of how to hit it.¹² Even though he is not, as was suggested above, totally in the dark as to what a proper throw looked like, whatever knowledge he possessed in that regard was not available for the purposes of guiding his throwing.

Against this backdrop, moreover, Tim's inability to hit a bullseye intentionally seems fairly robust. Even if Tim's track record were better than it actually was, with Mr. Magoo-like luck, say, and Magoo-Tim were to have hit one bullseye after another, each of his successes would have been unintentional. And the explanation is no different: despite his string of dumb luck, Magoo-Tim lacks knowledge of how to hit a bullseye, and so, according to SIMPLE, doesn't hit the bullseye intentionally.

SIMPLE offers a perfectly adequate explanation of **darts**, along with cases like **avalanche**, **annoyance**, and **cake**, ones where the agent is more or less skillless and clueless. And if we understand SIMPLE as making an explanatory claim (beyond the mere statement of a necessary condition on intentional action), it can explain trickier cases like **sibling** too, where the agent possesses but fails to manifest her know-how.

1.2. Problems for SIMPLE

⁸ For discussion on this point, see Buckwalter et al. (2021).

⁹ Compare: one could intentionally try to win the lottery and win the lottery without intentionally winning the lottery.

¹⁰ For discussion on the significance of this distinction, see, e.g., Anscombe (1963, 52-3) and Paul (2009, 3-4).

¹¹ On this point, see, e.g., Bengson and Moffett's (2011, Ch. 1) distinction between knowing how to do something and knowing how *one* does something; in the case of skiing, for instance, a skier might know how to do a tricky jump in a way that the coach doesn't even if the coach knows how *one* does the tricky jump, and is able to teach it.

¹² Though see §4.1 for an anticipated challenge of this assessment on behalf of proponents of SIMPLE, and our reply to it.

As we move further away from cases of skilllessness, cluelessness, and beginner's luck, while at the same time *remaining* within the space of scenarios in which an agent doesn't (yet) know how to do what they succeed in doing, SIMPLE begins generating the wrong kinds of predictions.

Suppose Steve is a master of making and flying paper airplanes. Possessing such mastery involves a number of layers of skill. One involves picking the right piece of paper—considerations of thickness and weight are paramount. Another involves delicate folding—Steve did origami in his youth, which prepared him well. A third and obvious layer involves actually flying the plane: within a fairly robust range of conditions, Steve can, holding his paper airplane at about cheek-height, elbow aimed in the direction he wishes to throw, extend his elbow and snap his wrist with just the right force and timing so as to send the plane flying towards his target, typically a stop sign (assuming no errant gusts of wind). Now consider how Steve fares at darts:

Darts-2: Steve hears about and decides to join the lively weekly darts tournament at *Che's Lounge*, and having never played a game of darts before in his life, Steve signs up to play, fully intending to hit the bullseye on this and every subsequent throw. Moreover, Steve knows that he might be in a position to do well at darts by recruiting other, paper airplane-related skills. On his first turn he looks at the bullseye, makes some arm movements that, for all he knows, resemble the arm movements of genuine dart players, and sends the dart on a wing and a prayer. Lo and Behold! Steve hits the bullseye.

Steve's case (**darts-2**) is similar to Tim's (**darts**) in that neither Steve nor Tim know how to hit the bullseye; each is a total beginner *when it comes to darts*. This is so, to be clear, even though Steve (though not Tim) has a background that includes repeatedly extending his arm and contracting muscles in a dart-throwing-like way. But since neither knows how to hit a bullseye, SIMPLE generates the result that both Tim and Steve fail to hit the bullseye intentionally.

And yet Steve's case differs from Tim's in that Steve—given his airplane throwing background—does *something* intentionally, in throwing the dart, that Tim does not: Steve intentionally hits the bullseye.¹³ Hear us out.

When Steve takes aim at the bullseye with his dart, to a first approximation, he knows (perhaps only tacitly) that he might recruit certain of his other skills to meet this

¹³ One might, for instance, think that only maestros of darts hit bullseyes intentionally; at best Steve intentionally *hits the dart board*. We can make the same point with a slightly weaker claim: even if Steve fails to hit the bullseye intentionally, he nevertheless hits the *dart board* intentionally, whereas for Tim, *both* of the corresponding actions are unintentional because too lucky.

novel objective.¹⁴ Given Steve’s knowledge of the similarity between what he knows how to do and what he is attempting to do, he seems to be in a position to reasonably conclude in favor of throwing the dart *thusly*, more or less like a paper airplane. And his throwing in that way is, given the skills he recruits, sensitive to the actual features of his throwing environment, say the distance from his target, the position of the target on the wall, ambient conditions, the heft and shape of what he’s throwing, and so on. In other words, these other skills Steve recruits are non-accidentally connected to accurate dart-throws. These two features—(i) Steve’s knowledge that he might recruit other skills, and (ii) the non-accidental connection between his manifesting those other skills and his success—afford Steve a certain kind of control over his attempts that prompts us to think that he succeeds in doing something *on purpose* that Tim, similarly situated, would only succeed in doing *by accident*, even if Steve does not, strictly speaking, count as knowing how to hit the bullseye.¹⁵

Because SIMPLE incorrectly predicts that Steve’s hitting the dart board is, like Tim’s hitting the dart board, unintentional, SIMPLE is false. What has gone wrong?

Even though neither Steve nor Tim knows how to do what they succeed in doing, what Steve has and Tim lacks is knowledge of how to do *something else*, and in particular something else that is not too distant from throwing a dart in a way that tends towards bullseyes (Secs. 2-3). It is because Steve knows how to do something else suitably related to what he succeeds in doing that his success counts as intentional. And it’s precisely this sort of skill-proximal relationship to which SIMPLE is insensitive.

We are not making the untenable claim that intentional action and knowledge-how stand in *no interesting relationship*; to think so would be to snub one’s nose at the recent history of the theory of action.¹⁶ Instead, we claim only that intentional action and knowledge-how stand in a more nuanced relationship than SIMPLE captures, and so more nuanced than has been appreciated thus far.

Let us now turn to introducing and fleshing out SKILL.

2. SKILL

Perhaps the most that can be said for SIMPLE is that it characterizes central or well-ordered cases of intentional action. We might even go so far as to say that SIMPLE explains intentional action *par excellence*. But even that way of putting things risks being misleading, since there is nothing obviously deviant or defective about the kinds of

¹⁴ We are not claiming that it is a necessary condition on Steve’s acting intentionally that he *knows* that he might recruit other skills to meet this novel objective, although such knowledge may be part of a set of jointly sufficient conditions.

¹⁵ In Section 4, we distinguish our position from Pavese’s (2020) *probabilistic knowledge-how* account, according to which knowing how to hit the bullseye requires probabilistic knowledge that the means one employs are *sufficiently likely* to result in hitting the bullseye.

¹⁶ See, e.g., Hawley 2003; Pavese 2022, Hornsby 2004, 2011; Setiya 2012.

intentional actions that fall outside its scope, such as in **darts-2**. And the problem magnifies when we consider that cases like **darts-2** are hardly idiosyncratic; after all, our practical needs often require us to call upon our dispositions, acquired and refined for other purposes, when attempting to meet novel objectives. We might even expect that across the spectrum of novel objectives (playing darts for the first time), the *default* strategy we have at our disposal will be to simply call on (like Steve) what by our lights *looks* familiar and reliable as a way to do something similar.

In this section, we'll explain and motivate SKILL, a weaker and more plausible thesis than SIMPLE. Recall that SKILL says:

SKILL: If S ϕ -s intentionally, S knows how to ψ , where ψ is sufficiently skill-proximal to ϕ .

There is a technical point that needs to be addressed in order to motivate this thesis. The variables ϕ and ψ range over act-types, which we can think of as bodily movements¹⁷ under descriptions. Particular tokens of these act types are the token bodily movements that, when caused in the right way, serve as candidates for actions (of that type) that agent performs intentionally.

One such act-type might be that of *hailing a cab*; in New York, various token movements of walking to the curb and raising one's arm above one's head fall under the description, 'hailing a cab.' Raising one's arm above one's head from the curb is a *means by which* one hails a cab in New York. Employing these means involves tokening certain movements that, when non-deviantly related to one's actually hailing a cab, render the latter intentional. Tokens of walking to the curb and giving passing cars a thumbs down are not; giving passing cars a thumbs down from the curb is not a means by which one hails a cab in New York. If one hailed a cab by tokening these other movements, one's success would be merely accidental.

Another such act-type might be *bowling a strike*. In a bowling alley, various token movements of approaching the lane while holding a bowling ball, swinging one's arm behind one's body and then forward in a pendulum motion, and releasing the ball down the lane while planting the opposite foot fall under the description, 'bowling a strike.' These movements are the means by which, when non-deviantly causally related to one's actually knocking down all the pins—as they are in paradigmatic cases—one intentionally bowls a strike.

For any given act-type, there will typically be many interesting generalizations over token movements by which to intentionally bring it about. Call these "ways" or

¹⁷ The restriction to *bodily movements* is a simplifying assumption; *deliberating*, for instance, is sometimes treated as a mental action, and deliberating need not involve any bodily movements (Mele 1997, 2021; Peacocke 2021).

“means” of tokening an act-type. And for simplicity, we’ll proceed as though there are certain *canonical* means of tokening an act-type, unless context requires we drop that simplifying assumption; as such, we may assume an agent skilled in ϕ -ing manifests this skill by employing *the* (canonical) means of ϕ -ing.

With these clarifications in place, consider that SKILL appeals to a notion of similarity between act-types, that of being “sufficiently skill-proximal.” This invites two related questions: What determines skill proximity between act-types? And what makes two act-types sufficiently skill-proximal? Most of our time will be spent presenting an answer to the first and more general question; once that has been accomplished, the second question becomes much more tractable.

2.1. Skill proximity

Let’s start with the first and more general question. Skill proximity is a kind of modal proximity—a proximity determined by facts about what *would* (tend to) bring about what. In particular, skill proximity between two act-types is a matter of the extent to which the particular movements in a token of one act-type (the means by which one succeeds in tokening an act of that type) would tend to bring about actions of another type in contexts over which the latter is defined. The rough idea, spelled out further below, is that the act-types are skill proximal *to the extent that the means of bringing about one will reliably serve to bring about the other*.

It is easy to see the potential payoff of exploiting this idea: if intentionally ϕ -ing involves knowing how to do *something sufficiently skill proximal to ϕ -ing*, then we immediately have the resources to explain why Steve did something intentional in **darts-2** when he hit the bullseye (or more weakly: the dart board). And the usefulness of these resources generalizes over to other cases like **darts-2**, where we—as so often—appropriate the means we do in intentionally performing novel action types.

But to really see the payoffs from SKILL, skill proximity conditions need to be precisified; and since skill proximity is (as we’re relying on it) under the scope of know-how in SKILL, the conditions for sufficient skill proximity will need to be understood *in relation to know-how*.

Naturally, one would prefer to see sufficient skill proximity understood in terms friendly to one’s own theory. To that end, we’ll discuss two general strategies for explaining skill proximity; the first will likely appeal to “intellectualists,” those who think that skill is a special kind of *de se* propositional knowledge.¹⁸ The second will likely appeal to anti-intellectualists, who tend to be broadly Rylean in their commitments; skills are (perhaps multi-track) dispositions or abilities. We won’t argue that one way of going is

¹⁸ For a related intellectualist-friendly account of skill, on which skill is a disposition to form relevant knowledge-wh, see Stanley and Williamson (2017).

preferable to the other, and we instead hope that this menu of compatible explanations will make SKILL all the more appealing.

Here is one broadly intellectualist way to understand skill proximity. The proximity of two skills is determined by the proximity of *means by which* those skills are manifested. The basic idea is that an agent's skill in one act type consists in her possessing certain *knowledge of means*, the employment of which constitute (or cause¹⁹) an agent's tokening an action of that type. One skill is then proximal to another to the extent that the knowledge of means employed in one can be *co-opted* for the other. Thus, In slogan form: for the intellectualist, skill proximity is grounded in means-knowledge co-option or 'transferability' across action types.

Recall Steve, who knows how to accurately throw paper airplanes. When Steve throws accurately, he does so by employing means that are in some sense general, and in some sense particular: in general, the means include knowing how to position his body relative to his target, the angle at which to hold the plane, the force with which he extends his arm and snaps his wrist in the direction of the target. These will, within certain margins, remain fixed on any throw. But Steve also knows how to accommodate and modify his technique on any given throw in the face of particular environmental conditions like humidity and wind. We can say that Steve knows how to accurately throw paper airplanes *by means of M*, where 'M' stands in for this long list of movement-patterns Steve employs when he manifests this knowledge-how.

In front of a dart board, not all of Steve's knowledge of the means of throwing paper airplanes accurately carries over in the sense that not all of M is *relevant*. He does not, for instance, have to worry about how the weather might affect his throw when he's inside a bar; nor does he have to worry about what he knows about what makes paper airplanes *glide* at certain angles; this knowledge is inapplicable to heavier (non-gliding) darts. But a proper subset of Steve's knowledge of means *will* carry over, namely, the subset involving how to position his body relative to his target, or how to modulate the force of extending his arm and snapping his wrist, etc. Call this "M*." Employing M* will fairly reliably give rise to Steve's hitting the bullseye; when Steve positions himself, takes aim, and throws by employing means that are conducive to hitting bullseyes, and he hits them because of the means he employed, his success is non-accidentally related to something he knows how to do. And this is why throwing paper airplanes accurately and hitting bullseyes are fairly skill-proximal act-types; M* is a means of reliably hitting bullseyes, and it is a subset of M, the means Steve employs when accurately throwing paper airplanes.

In contrast, Tim lacks knowledge of M; he is, after all, knowledgeable about darts only to the extent that he knows a good throw when he sees one, but he is otherwise

¹⁹ See Kelley (2022) for an argument that we can act non-basically through causal rather than constitutive means.

skillless. We can add that Tim lacks knowledge of M^* too (Tim has never even heard of paper airplanes). So Tim enjoys no proximal skills he could carry over to the act type ‘hitting the bullseye’ so as to render his occasional success intentional.

To sharpen the (intellectualist-friendly) approach to skill proximity further, let’s consider a further character, Nour, who is like Steve, but who always ends their paper-airplane throw with an additional kind of ‘flourish’ with the wrist, a movement that we can understand as just one additional movement beyond what Steve does (under the description of M) when throwing paper airplanes. Call Nour’s means of throwing the paper airplane “ $M+$,” which simply includes all of Steve’s means, M , plus the extra flourish. We can further characterize a proper subset of $M+$ (i.e., $M+^*$) in terms of relevance (*vis-à-vis* darts): the means of throwing the airplane that would relevantly carry over to dart-throwing.²⁰ Since $M+^*$ will *not* include the Nour’s flourish (irrelevant to darts) but everything else that features in M^* , $M+^*$ and M^* are coextensive. Would it follow (on our intellectualist-friendly gloss of SKILL) that, since M^* and $M+^*$ are co-extensive, Nour intentionally hits the bullseye (when employing $M+^*$) provided Steve intentionally hits the bullseye (when employing M^*)?

An affirmative answer here will be problematic; this is apparent once we add two further details. First, (i) suppose that in all or most near-by worlds where Nour employs $M+^*$ when throwing the dart, they also employ the flourish; and (ii) that while $M+^*$ are means that would bring about hitting the dart board with a dart, the means consisting in $\{M+^* + \text{the flourish}\}$ are much less reliable. The flourish, suppose, often ‘sabotages’ the dart throw but not the airplane throw. When Nour uses $\{M+^* + \text{the flourish}\}$, then, even when they succeed in hitting the bullseye using the means consisting in $\{M+^* + \text{the flourish}\}$, their successful dart throw not only fails to manifest know-how (to hit the dart board), but further, Nour (employing means $\{M+^* + \text{the flourish}\}$) fails to have hit the dart board intentionally.

The case of Nour suggests a tweak to the intellectualist account of skill proximity as it would feature in SKILL. Rather than to say that one skill is then proximal to another to the extent that the knowledge of means employed in one can be *co-opted* for the other, the intellectualist should say that one skill is then proximal to another to the extent that the knowledge of means employed in one can be *robustly co-opted* for the other; knowledge of means, A , can be robustly co-opted for knowledge of means B , just in case *not easily* would one co-opt one means for another only by also employing additional means the employment of which would, in conjunction with the co-opted means, issue in failure.

With this tweak in place, we can characterize SKILL, along intellectualist lines, as follows:

²⁰ This will be a subset involving how to position the body relative to the target, or how to modulate the force of extending the arm and snapping the wrist, etc.

SKILL_{INTELLECTUALISM}: If S ϕ -s intentionally, S knows how to ψ , where (i) ϕ is sufficiently skill-proximal to ψ ; and (ii) ϕ and ψ are skill-proximal to the extent that S's knowledge of means M_ϕ employed in ϕ is *robustly co-optable* for S's knowledge of M_ψ employed in ψ .

SKILL_{INTELLECTUALISM} distinguishes Steve as hitting the dart board intentionally from both Tim and Nour who don't. Even more, *SKILL_{INTELLECTUALISM}* lines Nour up with Tim rather than Steve *without* any commitment to the idea that *anyone* who uses a (de facto dart-sabotaging) flourish in paper-airplane throwing would thereby fail to hit the dart board intentionally via co-optable means. Just suppose Nour* is like Nour with an important exception: whereas Nour uses the flourish in paper-airplanes (unlike Steve) and carries this flourish over to dart-throwing, Nour* knows that, given the difference in weight between the paper airplanes and the darts, it would be a mistake to employ means $\{M^{+*} + \text{the flourish}\}$ rather than just M^{+*} . Since M^{+*} was coextensive with M^* , there is no barrier to Nour* counting as hitting the dartboard intentionally when employing means M^{+*} .

Having laid out this broadly intellectualist account of skill proximity – as this would feature in *SKILL* – let's consider now how a broadly *anti-intellectualist* gloss of skill proximity might go. Start with the idea that an agent skilled in ϕ -ing possesses certain reliable or safe dispositions to successfully ϕ when she tries or intends to ϕ .²¹ Steve knows how to accurately throw paper airplanes insofar as he possesses certain canonical dispositions– call them “D” – to accurately throw paper airplanes when he tries or intends to do so, and his accurate throws are attributable to him as intentional actions when and because they are manifestations of D. As in the intellectualist explanation above, not all of D will be relevant to Steve's hitting a bullseye with a dart, but some important subset of D, call it “D*,” will; manifesting D* will be fairly conducive to hitting bullseyes.

We can see how then, for the anti-intellectualist, skill-proximity will be a matter of overlapping dispositions to succeed when one tries (as opposed to a matter of

²¹ By ‘safe’ we mean just that the exercise of the relevant disposition would result in successfully ϕ -ing in a suitably wide class of worlds where we hold fixed those conditions where reliability at ϕ -ing matters for the purpose of attributing that particular skill-type. Pottery skill, for instance, implicates safe dispositions to smooth and burnish functional and aesthetically pleasing clay in worlds where one tries with the right kind of clay available along with a working kiln, etc. but it does not implicate any such disposition to secure such results across worlds with abnormally high ambient moisture, lack of availability of relevant tools, etc. Moreover, we should expect that different domains of performance will vary (for one to possess that kind of skill) with respect to just *how* safe those dispositions must be which are implicated by skill possession in that domain. Skill at the point guard position in basketball implicates safe dispositions to make passes to teammates without too often turning the ball over, *even* while tolerating a turnover or two per game. Skill at electrical engineering on the other hand implicates dispositions with a more demanding modal profile – viz., such that in nearly all worlds where one tries in appropriate conditions, one's design is not dangerous/faulty.

overlapping knowledge of means). And, as with the case of *SKILL_{INTELLECTUALISM}*, we don't want it to be the case that *mere* overlapping dispositions suffice for the kind of skill proximity that matters for acting intentionally, given that (as our case of Nour suggests) a set of dispositions might 'carry over' – viz., by being co-opted by an agent only by also importing over additional dispositions (from D) which (in conjunction with D*) are not success-conducive. Thus, as the 'robust' qualifier on co-optability is needed in *SKILL_{INTELLECTUALISM}*, the same will be the case for *SKILL_{ANTI-INTELLECTUALISM}*. Accordingly, and by parity of reasoning from *SKILL_{INTELLECTUALISM}*, we get:

SKILL_{ANTI-INTELLECTUALISM}: If S ϕ -s intentionally, S knows how to ψ , where (i) ϕ is sufficiently skill-proximal to ψ ; and (ii) ϕ and ψ are skill-proximal to the extent that S's safe ϕ -ing dispositions are *robustly co-optable* for S's safe ψ -ing dispositions.

Again, these two sketches of intellectualist and anti-intellectualist treatments of skill proximity are meant to highlight the overall neutrality and appeal of SKILL, not to make any substantive claims about SKILL's connection to the intellectualism debate.

2.2. Sufficiently skill-proximal

So far, we've suggested that skill proximity is a matter of the extent to which one's canonical knowledge of means (safe dispositions) employed in one skill can robustly "carry over" to or "overlap with" another so as to give rise to reliable success in a related activity. This puts us in a position to address the second question mentioned above: if overlapping knowledge of means (safe dispositions) makes for skill proximity, what metrics would then determine *sufficient* skill proximity?

Consider that the sorts of activities one is engaging in will determine just how reliably one has to succeed in order to count as intentional. Take, for example, recent work by Romy Jaster (2020), building on Manley and Wasserman (2008). Jaster argues that an agent has the ability to ϕ intentionally only if the agent ϕ 's in a sufficiently high proportion of the relevant situations in which she intends to ϕ .²² To be able to bake a loaf of bread, one must at least be able to do better than chance in the relevant intention-

²² Jaster's view of agent's abilities is doubly contextualist; context determines which situations are relevant *and* the weight that various relevant situations are given (see especially her Section 4.3). This way of theorizing about abilities has the advantage of being able to capture different senses of ability in terms of which situations are relevant. For instance, to assess whether an agent has the general ability to jump a fence, one holds fixed the "stable, mostly intrinsic features of the agent"; to assess whether an agent has what Geach (1957) called a "particular ability" to jump a fence, one holds fixed the "totality of features in an agent's circumstances" (116-117). It may seem that particular abilities do not really admit of the proportionality-across-relevant-situations analysis because particular abilities are "all-in," so to speak. But particular abilities fit the analysis trivially: to succeed across a suitable proportion of the relevant intention-situations is to succeed in all of them, namely the one you are actually in.

situations; if, say, 90% of my attempts resulted in a collapsed brick of dough, it would not be appropriate to say that I was able to do so intentionally; my successes would be merely accidental. Intuitively, the explanation of this fact is that I wouldn't count as knowing how.

In other sorts of activities, a sufficiently high proportion of the relevant intention-situations may be low in some absolute terms.²³ For instance, Steve Kerr, the record-holder for three-point percentage in the National Basketball Association, made 45.4% of his three-point shots over his career (Steph Curry, arguably the best current shooter, sits at a “mere” 42.77%). And outfielder Ty Cobb had the highest career batting average in Major League Baseball history at 36.6%.

Given that the activity-type sets some lower bounds on just how reliable one must be to count as knowing how,²⁴ one might think that the activity-type likewise sets a lower bound on how safely proximal skills must give rise to success in order for an agent to succeed intentionally. We can thus introduce the *safety* of one means, given another act type. If M_ψ is a means to ψ , the safety of M_ψ , given ϕ is determined by how often one succeeds in ϕ -ing by M_ψ -ing. In these terms, the relationship between M_ψ , given ψ , sets a plausible upper bound on safety; nothing is more skill proximal to ψ than ψ itself. (Note that, for this reason, SIMPLE comes out as the special, limiting case of SKILL.)

According to this proposal, the higher the safety of M_ψ , given ϕ , the more skill proximal ψ and ϕ will be. And the more skill proximal ψ and ϕ are, the more pressure there is to deem successful ϕ -ings that employ M_ψ intentional. We've gestured at one way to set a threshold that would determine sufficiency: the type of activity one is engaging in often supplies some lower bound below which successes are too fortuitous to count as intentional. It's possible that not all activities provide us with such clear-cut lower bounds, and so there may sometimes be a penumbra of indeterminacy around whether an agent brings to bear skills the safety of which ensures sufficient proximity.

But this tolerance for a certain amount of indeterminacy is a feature of *any* plausible account of the relationship between what an agent can do intentionally and what she knows how to do, not a bug. To get a sense of why it is a feature, consider an observation from Jaster (2020):

Someone who has driven a car once or twice and manages to get from A to B without running someone over, but drives extremely poorly otherwise, may be counted as having the ability to drive in one context, but not in another, for instance. When asked whether she can drive, she may truly say “yes” when a car

²³ For discussion on this point, see, e.g., Greco (2010: 60) and Carter, Jarvis, and Rubin (2015: 1065).

²⁴ The matter of just *how* a given activity type sets such bounds may depend on just how formalised the norms governing what counts as good performance (within that activity type) are codified. For discussion, see Sosa (2010).

needs to be moved from one space in an empty parking lot to another, say, but she may truly say “no” when a car needs to be moved from one neighborhood of a crowded city to another. (27)

There are many activities on the spectrum between moving a car from one space in an empty parking lot to another, and moving from one neighborhood of a crowded city to another. For instance, moving a car from one space in a crowded parking lot to another, or moving from the airport’s waiting lot to the arrivals terminal, fall somewhere between the extremes Jaster identifies.

In these middling cases, it’s not obvious whether the driver has the ability to get from A to B intentionally; there may be no fact of the matter. But across all these cases, middling and extreme alike, the driver’s skills and know-how remain fixed; what shifts is whether, according to facts about the activity and context, those fixed skills and know-how engender enough success across modal space for us to speak truly when we that she could intentionally get from A to B.

Any plausible account of the relationship between skills and intentional action will be saddled with some indeterminacy, but this indeterminacy is just a reflection of the extent to which facts about activity and context settle a threshold for success. When and to the extent that they do, possessing *sufficiently* proximal skills is just possessing skills the safety of which exceeds this activity- and context-dependent threshold.

3. Taking stock

Let’s take stock. In §2, we introduced the Skill Proximity View, captured by SKILL:

SKILL: If S ϕ -s intentionally, S knows how to ψ , where ϕ is sufficiently skill-proximal to ψ .

SKILL is an improvement on SIMPLE because it can accommodate cases in which an agent, lacking knowledge of how to do the task at hand, employs other know-how to succeed. SIMPLE seemed to predict that **darts** and **darts-2** were both cases of unintentional success. But Tim (**darts**) was a total novice succeeding by beginner’s luck, and Steve (**darts-2**) was skilled in a proximal activity, that of accurately throwing paper airplanes, and he employed those proximal skills in succeeding.

We argued that skill proximity could be understood in two ways, one friendly to intellectualists about know-how, the other friendly to anti-intellectualists about know-how; both approaches to know-how, accordingly, (via reference to *SKILL_{INTELLECTUALISM}* and *SKILL_{ANTI-INTELLECTUALISM}*) can deal with cases like **darts-2** in a way previously unavailable, as well as with more complex cases like Nour and Nour*. In order to get these results, on the intellectualist telling, skill proximity was a matter of overlapping knowledge of

means.²⁵ And in the anti-intellectualist telling, skill proximity was a matter of overlapping dispositions to succeed when one tries. *Sufficient* skill proximity was relative to activity; activity-types set lower bounds on how reliably one must be able to succeed to count as knowing how, and sufficient skill proximity piggybacks off of this activity-dependence.

3.1. Partners in Guilt

There is precedent to the idea that we can succeed in doing one thing intentionally by intentionally doing something else, and in that sense SKILL has partners in guilt. Our aim here is to situate our thesis among its partners. Consider three proposals due to Kieran Setiya (2008), Carlotta Pavese (2020), and Adam Carter (2019). Start with Setiya.

Setiya's position is captured by MEANS:

MEANS: If S ϕ -s intentionally, S knows how to ϕ , or S knows how to ψ , where ψ is a means to ϕ .²⁶

According to MEANS, the relationship between intentional action and knowledge-how is disjunctive: ϕ -ing intentionally requires knowing how to ϕ , or it requires knowing how to ψ , where ψ is a means to ϕ in the context.

Like SKILL, MEANS is consistent with there being cases in which an agent ϕ -s intentionally without knowing how to ϕ , but these two views place importantly different constraints on *what else one must know how to do*. MEANS only requires that the agent knows how to perform an action that in fact counts as a way to succeed, while SKILL requires that the agent know how to perform a skill-proximal action. What does this difference come to?

Unlike SKILL, MEANS places no constraints on whether the agent must be aware that what she knows how to do will count as a way of succeeding, nor on whether what she knows how to do will afford her sufficient sensitivity to the actual features of her environment that matter to her having control over what she succeeds in doing. For all MEANS says, an agent might not only intentionally buy a ticket in a fair lottery, but also intentionally win that lottery, if they know how to buy tickets and exercise that knowledge in buying the ticket turns out to be the winner. MEANS seems to expose intentional action to too much luck and too little control.

Still, our aim is not to be combative. To the extent that one found a view like MEANS attractive because it was a less demanding alternative to SIMPLE, SKILL articulates constraints that any such less demanding view must satisfy; charitably, then, SKILL might be thought to extend or supplement MEANS by imposing further

²⁵ Additional details were made explicit, e.g., concerning co-optability (as pertinent to cases like that of Nour).

²⁶ See Setiya's discussion of the "bomb" case (2008: 404).

constraints on its second disjunct so as to exclude success through too much luck and too little control from counting as intentional.

Pavese's (2020) position differs from both SKILL and MEANS—and is plausibly a version of SIMPLE—in that she accepts that knowing how to ϕ is necessary for ϕ -ing intentionally, although she relaxes the conditions under which an agent counts as knowing how to ϕ by offering an account of probabilistic knowledge-how:

PROB: If S ϕ -s intentionally, S knows how to ϕ , where S has such knowledge only if S knows, for some means ψ of ϕ -ing, that oneself is *sufficiently likely* to ϕ by ψ -ing.²⁷

According to PROB, the relationship between intentional action and knowledge-how is probabilistic: ϕ -ing intentionally requires knowing how to ϕ , and knowing how to ϕ is a matter of knowing that one can probably ϕ by employing means ψ .

SKILL is more demanding than MEANS but less demanding than PROB. To see why PROB is *too* demanding, note that it faces the same problem that SIMPLE does, that of too tightly wedding what one does intentionally to what one knows how to do. Recall **darts-2**; Steve, the protagonist, knows that he might recruit his paper airplane skills to hit a bullseye. But this does not, all on its own, put Steve in a position to know that recruiting his paper airplane skills will *probably* result in his hitting the bullseye. The latter is, perhaps, something Steve could come to know after a few successful attempts, but there is a lot of daylight between an agent's *knowing that proximal skills may carry over*, and an agent's *knowing that proximal means are sufficiently likely to bring about success*.

It could, of course, still be true that some cases in which an agent employs proximal skills are ones in which they *also* have probabilistic knowledge-how. But to collapse these two notions would, we think, either render probabilistic know-how too cheap, or the phenomenon of intentional success in novel circumstances too demanding.²⁸

Finally, consider Carter's **pole vault** case:

²⁷ See Pavese's discussion of the principles *Intentionality/Probabilistic Knowledge* and *Intellectualism about Know-How* (2020: 353). We set aside considerations bearing specifically on the intellectualist aspects of her proposal.

²⁸ For all we've said, it may be true that when agents know how to ϕ *and* knows that ϕ -ing is a way to (probably) ψ , then the agent knows how to ψ in a way that satisfies PROB. For instance, I might have never pushed a vase off a table before intentionally, but I know that reliance on means I've used to push similar sized objects off of similar surfaces will suffice to do the trick

But the agents we're concerned with—like Steve—lack this quasi-inferential probabilistic know-how because they do not satisfy the second conjunct. On the minimal assumption that reflective agents acquire track-record evidence after repeated attempts, there will be a point at which the agent is in a position to know that ϕ -ing is a way to (probably) ψ , and so count as knowing how to ψ . Thanks to [REDACTED] for raising this point.

Pole vault: Paul is hoping to make his high-school pole vaulting team. In order to qualify, Paul must demonstrate that he has the ability to jump over the competition bar set at 10ft. Paul's dishonest nemesis told Paul the competition bar for tryouts was three feet higher — 13ft — in an effort to dissuade Paul from attempting to make the team. Paul nonetheless showed up to tryouts and, setting the bar to 13ft rather than 10ft, proceeded to jump over the bar (impressing his coach and his nemesis). (Carter 2019: 2496)

Carter suggests, attending to cases like **pole vault**, that the relationship between what one does successfully and *which* abilities one manifests is nuanced; the manifestation of an ability is not necessarily a matter of the subject's doing what the ability is an ability to do.²⁹ In particular, **pole vault** is a case in which the agent exercises his ability to jump over the bar at 10ft by doing something *else*, namely jumping over the bar at 13ft.

Although Carter is not interested in the relationship between intentional action and knowledge-how *per se*, his proposal might be understood as having implications for that relationship. In particular, we might think that his account, suitably modified, generalizes to the relationship between what one succeeds in doing *intentionally*, and which abilities one intentionally manifests. Consider **hole vault**:

Hole vault: Hal is a skilled pole-vaulter trying his hand for the first time at "hole-vaulting," a new feat of skill and grace put together by the twisted minds at *cirque du soleil*. Hole-vaulting requires many of the same skills as pole-vaulting, but rather than vault over a pole at a given height, one must vault through a hole of a given shape at a given height. Hal, having some understandable reservations about whether he possesses sufficient aerial bodily control, proceeds to jump through the hole.

Hal's jumping through the hole is no accident; despite his understandable reservations, he employs a great deal of transferable pole-vaulting skill in his hole-vaulting success. To the extent that **pole vault** prompts us to think that one can manifest an ability to ϕ by ψ -ing (rather than ϕ -ing), **hole vault** should prompt us to think that one can manifest an ability to intentionally ϕ by intentionally ψ -ing (rather than intentionally ϕ -ing). Moreover, when we extend Carter's view in this way, we see that it is plausibly supported by SKILL; pole-vaulting is, in these circumstances, sufficiently skill-proximal to hole-vaulting for Hal to intentionally hole-vault by manifesting his pole-vaulting know-how.

To sum up: Seitya, Carter, and Pavese have each offered theories that are weaker than SIMPLE, and so in that sense SKILL has partners in guilt. We suggested that SKILL occupied a plausible middle-ground between MEANS, which was insufficiently

²⁹ *Pace* philosophers like Millar (2009).

demanding insofar as it permitted cases with too much luck and too little control to count as intentional, and PROB, which was overly demanding insofar as it prohibited cases in which the agent didn't know that her means were sufficiently likely succeed but only that they might carry over.

Ecumenically, SKILL encodes a conception of intentional action that is importantly informed by—and builds upon—the views that flank it, all of which seek to strike a plausible balance between accounting for the control characteristic of intentional action and the luck such action tolerates.

Let's now consider several potential criticisms against SKILL on behalf of those sympathetic to SIMPLE, and then, after addressing them, explain why SKILL plays an illuminating role in thinking about (i) learning how and (ii) the connection between intentional action and luck.

4. Objections and replies

4.1. The know-how defense

One way to defend SIMPLE is to resist the cases that motivate departing from it. For instance, one might think that Steve (**darts-2**) *did know how* to hit a bullseye, and it was precisely the fact that he manifested this extant knowledge-how that accounted for his intentionally hitting it. Perhaps what Steve lacked, according to this defense, was not knowledge of how to do what he did, but knowledge *that* he knew how.³⁰

This imagined objector grants that Steve hit the bullseye intentionally but insists that it is because he antecedently knew how to hit the bullseye. How plausible is it that Steve knew how to do what he did? To assess this, it's important to distinguish different things it might be taken to mean, and in so doing help clarify what's really at issue between SIMPLE and SKILL.

Surely Steve antecedently knew how to perform whatever *basic* actions he actually performed: his moving his arm *thusly* (in a particular throwing motion), say. But the ability to perform that or similar basic actions can't be what's at issue, since beginners like Tim plausibly possess those abilities too, yet fail to hit the bullseye intentionally. And anyway, *hitting the bullseye* is not a basic action.³¹ So the question at issue concerns the plausibility of Steve's antecedently possessing knowledge of how to perform the non-basic action he performs.

Here, we want to be careful not to enter into a battle of intuitional hearsay. Steve's possessing proximal skills either entails that he knows how to hit the bullseye, or it doesn't. But it seems false that merely possessing skills proximal to bullseye-hitting

³⁰ For some discussion more widely of cases where one seems to be able to use abilities they 'didn't know they had', see Moon (2018).

³¹ If Small (2019) is correct, it might be more appropriate to say that *hitting the bullseye* is not a basic action for *amateurs* like Steve and Tim, although it might be for experts.

entails that one knows how to hit bullseyes, and more generally that possessing ϕ -proximal skills entails one knows how to ϕ , since this would seem to imply also that Sally's possessing proximal avalanche-escaping skills entails she knows how to escape avalanches.³² And it was supposed to be a fixed point among contemporary theorists of action that Sally's successful avalanche escape was unintentional precisely because she didn't know how to do what she did.

But even if the entailment doesn't hold in general, don't we owe an explanation of why someone might be attracted to that idea in the first place?

The underlying appeal of thinking that Steve's possessing proximal bullseye-hitting skills entails he knows how to hit the bullseye might be that Steve's possessing those skills is *part of what puts him in a position to know how* to hit a bullseye (perhaps only after a few successful attempts). As we stressed in Section 3.1, there is a lot of daylight between an agent's *knowing that proximal skills may carry over*, and an agent's *knowing that proximal means are sufficiently likely to bring about success*. But surely, for sufficiently rational and self-aware agents like Steve, possessing the former knowledge is an important part of how he is posed to gain the latter.

This preserves an intimate connection between knowing how to ϕ and possessing skills sufficiently proximal to ϕ without collapsing the two. And it preserves the original line on the Sally case by not predicting that she knows how to escape an avalanche (we could certainly grant she has part of what one needs to be in a position to know how to escape an avalanche).

More realistically, since Steve already knows how to do something generally in the ballpark of hitting the bullseye, he is, perhaps in the course of hitting it, *extending* what he knows how to do. At the end of the game he might have put together a general 'recipe' for hitting bullseyes, a recipe he can knowingly deploy in future games. But to accept that he initially learns how to do it as he goes is to deny that he knew before he began.

There are, of course, cases where one realizes that one knew how to do something all along. Perhaps my lack of confidence prevents me from realizing I really do know what I'm doing, and it's only upon succeeding that I can overcome my own self-doubt. But these are the special cases, not the norm; hindsight biases shouldn't make us cavalier.

4.2. The no-intentional-action defense

³² Moreover, note that if merely possessing skills proximal to bullseye-hitting entailed that one knows how to hit bullseyes, this would presumably be by implication from a more general principle to the effect that: any subject knows to do anything X for which a Y skill is a proximal skill to Xing and S has Y. But that principle (along with weak premises, and given that skill proximity isn't a transitive relation) quickly overpredicts know-how by proximal skill case iteration.

Another way to defend SIMPLE is to insist that Steve *didn't* intentionally hit the bullseye. There are really two, distinct ways to develop this objection. According to one, Steve merely unintentionally hit it. According to another, Steve's hit is a "middling" or "non-intentional" action.

Consider the claim that Steve merely unintentionally hit the bullseye. Support for it might be put as a question: why think Steve hits the bullseye intentionally in virtue of his proximal skills, when Sally (**avalanche**) and Sam (**sibling**) possess but manifest skills proximal to avalanche-escaping, and yet escape their respective avalanches merely unintentionally?

In Section 1.1, we suggested that SIMPLE, all on its own, was silent on **sibling**, but that it was also plausibly underwritten by an explanatory claim: know-how is necessary for intentional action *because manifesting know-how is control-securing and luck-precluding*. Here, we want to say something similar, namely that even if SKILL, taken on its own, does not predict that Sally and Sam's actions are defective, there is a deeper explanatory connection underlying SKILL that does. In particular, proximal skill is necessary for intentional action *because manifesting a proximal skill because it is proximal is control-securing and luck-precluding* (albeit in a more nuanced way than SIMPLE). And this underlying explanatory principle clearly distinguishes Steve from Sally and Sam insofar as only Steve manifests proximal skills *because they are proximal skills*; Sally and Sam manifest proximal skills by dumb luck.

A bit more intuitively, in Section 1.2, we stressed that Steve's throw exhibits (1) knowledge of a similarity between what he knows how to do and what he's trying to do, and (2) sensitivity to the actual features of his throwing environment, and (1) and (2) afford him sufficient control over his attempts so as to render his success attributable to him as an intentional action. Sally and Sam satisfy neither condition; each mistakenly thinks that the oncoming avalanche is a big wave, and their respective frantic swimming motions are only accidentally related to their escaping the avalanche in the sense that they are not responsive to the actual features of the world that render them appropriate. Sally and Sam thus lack the kind of control over what each does that would be needed to view their respective successes as intentional.

The second way to develop this objection is to deny both that Steve hit the bullseye intentionally and that he hit it unintentionally. Instead, one might think, Steve's hitting the bullseye is what Mele and Moser (1994) have called a "middling action," or a "non-intentional" one.

The problem with this suggestion is that non-intentionality, at least as Mele and Moser discuss it, paradigmatically applies to the foreseeable consequences of one's intentional actions. Consider a famous case due to Harman (1976) in which a sniper, firing his gun in trying to kill a soldier, knowingly alerts the enemy to his presence. While alerting the enemy to his presence is something the sniper *does*, it is not, according to Mele and Moser, something the sniper does intentionally. Still, because the sniper is fully

aware that he will bring about this event in firing his gun, there is pressure to explain the nature of this action in different terms than, say, one of unwittingly alerting the burglars by illuminating the room.³³

It's difficult to see how Steve's hitting the bullseye would count as non-intentional if that category is understood by way of similarity to actions that are the foreseeable consequences of intentional actions. There are two clear points of divergence. The first is that Steve has no business being highly confident that his way of throwing it is a way of hitting the bullseye, whereas Harman's sniper has every reason to be highly confident that his shooting will alert the enemy. Second, even if what Steve can foresee is that his way of throwing *just might* hit the bullseye, this is precisely what he intends to do. In contrast, the sniper only intends to kill the soldier, not also to alert the enemy.

5. Learning how and agential contribution

We conclude with a more positive defense of SKILL in terms of its relationship to learning how, and in terms of the perspective it affords us on the relationship between luck and an agent's contribution to her purposive behavior.

Learning how to do something is a way of extending one's knowledge-how. Sometimes, learning how to do something is a matter of accidental success; one sees how one's actions affect the world and one learns how to act so as to bring about certain desired outcomes. Famously, penicillin was discovered by accident; Sir Alexander Fleming returned from a holiday to find that one of his Petri dishes contained a mold that prevented *Staphylococcus* bacteria from growing. Fleming and his team sought to then isolate pure penicillin from the mold. (It proved untenable, and it wasn't until years later that an Oxford laboratory established successful methods of production.)

Other times, learning how to do something is the result of deliberate individual practice. When Tim learned how to flip eggs in a pan, the process took a great deal of patience and spare eggs. But each attempt aimed at a kind of mastery—or sufficient reliability—in bringing about an antecedently known outcome. Tim knew what he wanted to do, and he tinkered with various ways of doing it so as to reliably produce the desired outcome.³⁴

Whether one learns how by observing accidental successes or by deliberate individual practice, learning how begins with knowing how. Of course, the know-how with which one begins is not what one learns how to do. It is, rather, knowledge of how to do something *else*, something appropriately related to what one learns how to do.³⁵

³³ Davidson (1980)

³⁴ And still other times, learning how is a matter of testimony (Hawley (2010)).

³⁵ This point has parallels in theoretical as opposed to practical learning. For instance, proponents of knowledge norms on inquiry (e.g., Willard-Kyle 2022) stress the importance of possessing and using propositional knowledge we have in the service of acquiring knowledge of propositions of which we are ignorant. A scientist, for instance, comes to learn that some claim of interest - X - is true by first learning

Fleming knew how to culture cells in Petri dishes and Tim knew how to make certain grabbish and wrist-flickish motions. Learning how to produce penicillin mold and to flip eggs was a matter of recruiting and re-purposing prior know-how to new ends.

In general, it is not true that learning how to do something *just happens* to us, certainly not in the way that non-actions just happen to us. Nor is it generally true that, in learning how to do something, our successes are unintentional or merely foreseeable consequences of what we do intentionally. Granted, some successes in learning how are accidental; perhaps Tim's first successful egg-flip was beginner's luck. But tinkering with various ways of flipping tended towards successes and away from failures; at some point in learning how, one begins to succeed intentionally, even if a bit shakily (because one does not yet know how).

One of the more interesting and substantive upshots of SKILL is that it captures the sense in which agents still learning how to do something can succeed on purpose (and non-defectively)³⁶. After all, succeeding on purpose reinforces the employment of certain means over others, ones which become reliably employed to bring about success across a wider range of situations and ground an agent's know-how. If SIMPLE were correct, learning how to do something would be a process devoid of intentional success, and would seem to inappropriately undermine the sense in which the learning agent contributes something to her own behavior. The agent's contribution to her purposive success in learning how would diminish to the point of being uninteresting.

What the agent contributes to intentional success is a certain kind of *control*. Luck tends to be control-undermining; but to think that anything short of knowing how to do what you in fact do renders success too lucky to be intentional is to adopt an overly demanding conception of the control characteristic of intentional action. It is perhaps the control characteristic of intentional action *par excellence*, but this is not the yardstick by which we measure agents with lesser degrees of mastery or skill, those still practicing and learning.

It is a fixed point of action theory that too much luck is incompatible with intentional success. There is no doubt that part of SIMPLE's long-standing appeal is its ability to exclude too-lucky successes from one's theory of intentional action. What counts as a too-lucky success, in this view, is success without knowing how to do what one in fact does. In this paper, we've challenged this thesis and argued for an alternative view, one that is grounded in the relationship between proximal skills. SKILL is likewise luck-excluding, albeit in a more nuanced way than SIMPLE. What counts as a too-lucky success, according to SKILL, is success without knowing how to do things sufficiently

that various X-adjacent facts are true. More mundanely, an inquirer might gain predictive knowledge of the behaviour of a pet in a new circumstance by track-record observation of that pet's behaviour in old circumstances.

³⁶ *Pace* authors like Habgood-Coote (2018), according to whom knowledge-how is the norm of intention.

skill-proximal to what one in fact does. One of SKILL's chief advantages over SIMPLE is its ability to explain intentional successes for agents still learning how, and for agents attempting – as we do so often – to meet novel objectives.³⁷

³⁷ [Acknowledgments removed]

References

- Amaya, S, 2018, "Two Kinds of Intentions: a New Defense of the Simple View," *Philosophical Studies* 175 (7):1767-1786.
- Anscombe, G.E.M., 1963, *Intention*, 2nd Edition, Oxford: Blackwell Press.
- Beddor, Bob, and Carlotta Pavese, 2022, "Practical Knowledge Without Luminosity," *Mind* 131: 917-934.
- Buckwalter, Wesley, David Rose, and John Turri, 2021, "Impossible Intentions," *American Philosophical Quarterly* 58 (4): 319–32.
- Bratman, M, 1984, "Two faces of intention," *The Philosophical Review*, 93(3), 375–405.
- Bratman, M, 1987, *Intentions, Plans, and Practical Reasoning*, Cambridge: Harvard University Press.
- Carter, J. Adam, 2021, "Exercising Abilities," *Synthese* 198: 2495-2509.
- Carter, J. Adam and Jesús Navarro, 2017, "The Defeasibility of Knowledge-How," *Philosophy and Phenomenological Research*, 95(3): 662–685.
- Carter, J. Adam., Jarvis, Benjamin W. & Rubin, Katherine, 2015, "Varieties of Cognitive Achievement," *Philosophical Studies* 172, 1603–1623.
- Davidson, Donald, 1980, *Essays on Actions and Events*, Oxford: Oxford University Press.
- Di Nucci, E, 2010, "Rational Constraints and the Simple View," *Analysis*, Volume 70, Issue 3, July 2010, Pages 481–486.
- Habgood-Coote, Joshua, 2018, "Knowledge-How is the Norm of Intention," *Philosophical Studies*, 125: 1703-1727.
- Hawley, K, 2003, "Success and Knowledge-How," *American Philosophical Quarterly*, 40(1): 19–31.
- Hornsby, J, 2004, "Agency and Actions," *Royal Institute of Philosophy Supplement*, 55: 1–23. doi:10.1017/S1358246100008614
- Hornsby, J, 2011, "Ryle's *Knowing-How* and Knowing How to Act," in Bengson and Moffett 2011c: 80–98.
- Kelley, Mikayla, 2023, "How to Perform a Nonbasic Action," forthcoming in *Noûs*, <https://doi.org/10.1111/nous.12440>.
- McCann, H, 2010, "Di Nucci on the Simple View," *Analysis* 70 (1):53-59
- McCann, H, 2011, "The Simple View again: a brief rejoinder," *Analysis*, Volume 71, Issue 2, April 2011, Pages 293–295
- Millar, Alan. 2009. 'What Is It That Cognitive Abilities Are Abilities to Do?' *Acta Analytica* 24 (4): 223–236.
- Moon, Andrew, 2018, "How to Use Cognitive Faculties you Never Knew you Had", *Pacific Philosophical Quarterly* 99: 251-275.
- Paul, Sarah K, 2009, "How We Know What We're Doing." *Philosophers' Imprint* 9(11).
- Pavese, C, 2022,. "Practical Knowledge First." *Synthese*, 200(5), 1-18.

- Pavese, C , 2021a, "Knowledge, Action, and Defeasibility", in *Reasons, Justification, and Defeaters*, edited by Jessica Brown and Mona Simion, chapter 8: 177–200, Oxford University Press, Oxford.
- Pavese, C, 2021b, "Knowledge How", *The Stanford Encyclopedia of Philosophy* (Fall 2022 Edition), Edward N. Zalta & Uri Nodelman (eds.), URL = [<https://plato.stanford.edu/archives/fall2022/entries/knowledge-how/>](https://plato.stanford.edu/archives/fall2022/entries/knowledge-how/)
- Pavese, C, 2020, "Probabilistic Knowledge in Action", *Analysis*, 80(2): 342–356.
- Pavese, C, 2018, "Know-How, Action, and Luck", in *Knowledge and Justification, New Perspectives*, special issue of *Synthese*, first online: 31 May 2018.
doi:10.1007/s11229-018-1823-7
- Ryle, Gilbert, 1949, *The Concept of Mind*, Chicago: University of Chicago Press.
- Setiya, Kieran, 2008, "Practical Knowledge", *Ethics*, Volume 118, 3: 388-409
- Setiya, Kieran, 2012, "Knowing How", *Proceedings of the Aristotelian Society*, 112(3): 285–307.
- Shepherd, Joshua, and J Adam Carter, *forthcoming*, "Knowledge, Practical Knowledge, and Intentional Action", *Ergo*, forthcoming.
- Sosa, Ernest, 1997, "Reflective Knowledge in the Best Circles", *The Journal of Philosophy*, 94(8), 410-430.
- Sosa, Ernest, 2010, "How Competence Matters in Epistemology", *Philosophical Perspectives* 24: 465-475.
- Stanley, Jason and Timothy Williamson, 2001, "Knowing How", *Journal of Philosophy*, 98(8): 411–444.
- Stanley, Jason and Timothy Williamson, 2017, "Skill." *Noûs*, 51 (4), 713-726.
- Stanley, Jason, 2011, *Know How*, Oxford: Oxford University Press.
- Velleman, David, 1992, "What Happens when Someone Acts?", *Mind*, 101 (403): 461-481.
- Willard-Kyle, Christopher, 2022, "The Knowledge Norm for Inquiry." forthcoming in *Journal of Philosophy*.
- Williams, John N., 2008, "Propositional Knowledge and Know-How", *Synthese* 165 (1): 107–25.