



Trying without fail

Ben Holguín¹ · Harvey Lederman² 

Accepted: 29 July 2024

© The Author(s), under exclusive licence to Springer Nature B.V. 2024

Abstract

An action is *agentially perfect* if and only if, if a person tries to perform it, they succeed, and, if a person performs it, they try to. We argue that trying itself is agentially perfect: if a person tries to try to do something, they try to do it; and, if a person tries to do something, they try to try to do it. We show how this claim sheds new light on questions about basic action, the logical structure of intentional action, and the notion of “options” in decision theory. On the way to these central ideas, we argue that a person can try to do something even if they believe it is impossible that they will succeed, that a person can try to do something even if they do not want to succeed, and that a person can try to do something even if they do not intend to succeed.

Keywords Trying · Basic action · Intentional action · Options · Iteration principles · Ability · Control · Decision theory · Philosophy of action · Action theory · Propositional attitude psychology

1 Introduction

There is a vision of the divine that centers on a perfect correspondence between divine act and divine will. The divinity, it is said, does everything it wills, and wills everything it does. In a folksier manner of speaking, the divinity does everything it tries to do, and tries to do everything it does.

We mere mortals, of course, are not like this at all. We do not do everything we try to do. A tired parent tries to stay calm, but fails when they fly into a rage instead. A thrillseeker tries to get high, but fails because the dealer sold them poison, not their favored drug. A pacifist tries to avert a revolution, but fails when their speech instead incenses the crowd. And, as these examples show, we also do not try to do

✉ Harvey Lederman
harvey.lederman@utexas.edu

¹ Johns Hopkins University, Baltimore, USA

² University of Texas at Austin, Austin, USA

everything we do. The parent flies into a rage, the thrillseeker poisons themselves, and the pacifist starts a war—although none of them was trying to do what they did.

The expression of the human will in action can be distorted in ways that make us quite unlike the divine. But is there a domain that is safe from such misfires—a domain in which the divine spark in our human will is guaranteed to shine through?

A central goal of this paper is to argue that there is. We characterize a class of doings that we call *agentially perfect*, doings that (to put it roughly) by their very nature ensure a perfect match between attempt and success. More exactly, ϕ -ing is *agentially perfect* if and only if: necessarily, if a person tries to ϕ , they ϕ , and necessarily, if a person ϕ s, they try to ϕ . We argue that trying itself is agentially perfect: necessarily, if a person tries to ϕ , they try to try to ϕ , and necessarily, if they try to try to ϕ , they try to ϕ .

We go on to develop some consequences of this claim. We argue that if trying is agentially perfect, then on a natural way of understanding “basic action”, all basic actions are tryings. Likewise, if trying is agentially perfect, then plausibly trying is essentially intentional in Anscombe’s (1957, §47) sense—that is, necessarily, anyone who tries to ϕ , intentionally tries to ϕ —yielding a striking divergence between the logic of intentional action and the logic of knowledge. Finally, we suggest that if trying is agentially perfect, then it opens the way to a new argument for the claim that tryings play the role of “options” in decision theory.

As a warmup to these central points, we begin by exploring a slightly different theme. Descartes writes of the will that it is “so free of its nature that it can never be constrained”.¹ We defend a related doctrine for trying. A common view is that what a person can try to do is highly constrained by their attitudes toward the success of their attempt: that if a person tries to do something, they must think they can succeed, they must want to succeed, and perhaps they must even intend to succeed. As we will explain, we believe that if these claims were correct, they would rule out the possibility that trying is agentially perfect. So we begin the paper with extended arguments against them, and in favor of a “Cartesian” conception of trying, on which trying is not constrained by such attitudes to success.

Section 2 draws attention to some different uses of ‘try’ and isolates a use that will be our focus throughout. Sections 3 and 4 defend the Cartesian conception of trying. We argue against various prominent claims about the connection between trying and believing (§3), between trying and wanting, and between trying and intending (§4). Section 5 isolates the property of agential perfection. Section 6 argues that trying is agentially perfect: necessarily, a person tries to ϕ if and only if they try to try to ϕ . Section 7 explores consequences of the agential perfection of trying for theories of basic action, for the logic of intentional action, and for the notion of options in decision theory. Section 8 concludes.

¹ *Passions of the Soul* I. XLI: “la volonté est tellement libre de sa nature, qu’elle ne peut jamais être contrainte”, Descartes (1996, XI.359), with translation in Descartes (1985–1992, vol 1, p. 343), cited by Albritton (1985).

2 On 'trying'

Our arguments will often rely on intuitive judgments about trying. As we elicit these judgments, it will be important to attend to different uses of the word 'try'. So, in this preliminary section, which also serves as an introduction to our overall approach, we begin with some key background points about trying and 'trying'.

We start with a claim that seems widely accepted these days, but is still worth stating explicitly: that trying is compatible with foreknowledge of success.² We find this intuitively clear in its own right. The mere fact that you knew you would succeed in raising your arm seems to pose no barrier to your having tried to raise it. But we also think these sorts of judgments are confirmed by reflection on the epistemology of ability: often, when you know that you are able to raise your arm, you also know that if you were to try to raise it, you would succeed (cf. Mandelkern et al. (2017)).

A subtler question is whether trying is compatible with effortlessness. That is: can a person who can ϕ without exerting any amount of physical or mental effort—a person who can ϕ at will, as it were—nonetheless count as having tried to ϕ ?

We think the answer is yes. We see nothing invariably problematic about describing a person who effortlessly succeeds in ϕ -ing as having tried to ϕ . And again there is an argument for this conclusion based on the epistemology of ability: often, when you know that you are able to effortlessly raise your arm, you know that, if you were to try to raise your arm, then whether or not you put any effort into it, you would succeed. So effortlessness must be compatible with trying, at least on one natural way of using the word 'try'.

Why the qualification about ways of using 'try' here? Because sometimes we say things like 'I can beat him in tennis without even trying'. Sometimes we order people to '*Actually* try!'. And sometimes we ask 'Is she going to have to try on this question, or will it be as easy as the last one?'. On the assumption that 'S is trying to ϕ ' never entails 'S is putting significant effort into ϕ -ing', it seems hard to make sense of these ways of talking.

What should we say about these more demanding uses of 'try'? One option is to give a pragmatic explanation, so that the literal meaning of 'try' needn't be adjusted to account for them. Another would be to conclude that 'try' is ambiguous or context-sensitive, and that on certain interpretations (though not all!) it expresses something along the lines of what is expressed by 'try hard' or 'try with sufficient effort'. The examples just given would be ones in which this sort of reading is salient.

This isn't the only kind of example that shows that 'try' has a complex array of uses. Suppose you promise your friend that you'll try to climb Mt. Everest. You train intensely and buy plane tickets to Nepal in advance of your scheduled climb. But

² The *locus classicus* of the view that trying is incompatible with foreknowledge of success is Wittgenstein (1950, §622), but see also Heath (1971) and Jones (1983) for sympathetic discussion. The canonical argument in favor of the compatibility of trying and foreknowledge of success was given by Grice (1989, p. 7) (originally delivered, 1967), but see also, e.g., O'Shaughnessy (1973), Armstrong (1973), Hornsby (1980), McGinn (1982, ch. 8), Schroeder (2001).

then disaster strikes in the form of an international pandemic, and you're forced to cancel your trip. Obviously, you aren't someone who has climbed Mt. Everest. But are you at least someone who has *tried* to climb it? On the one hand, if a stranger were to ask you 'Have you ever tried to climb Mt. Everest?', it would be pretty misleading to respond 'Yes I have'. (Imagine how the conversation might evolve: 'How far up the mountain did you get?'—'Well, I never actually made it to the mountain, but I did buy plane tickets to Nepal...') On the other hand, if your friend were to ask 'Did you even *try* to climb Everest?', it seems like you could speak truly in responding 'Of course I did—I bought tickets and everything!'

What should we make of these facts? As before, it could be that there's a pragmatic explanation. But also as before, it could be that for certain complicated, temporally extended actions—e.g., climbing Mt. Everest, writing a novel, starting a new career—whether one counts as having “tried” to perform the action in question is a context-sensitive matter. Perhaps in some contexts it suffices that one tried to perform certain preliminary steps (perhaps with the relevant intentions), while in other contexts one counts as having tried only if one manages to succeed in performing certain “core” actions—e.g., setting foot on the mountain, typing sentences, or leaving one's current job.

For the purposes of this paper it won't be so important whether 'try' has a simple lexical entry and these alternative “readings” are given a pragmatic explanation, or whether 'try' is genuinely ambiguous or context-sensitive. Our official target is whatever relation (or relations) are expressed by the very permissive readings of 'try'—that is, the readings of 'try' on which you count as having tried to climb Mt. Everest in virtue of having bought tickets to Nepal, and where *any* amount of effort could in principle be enough to count as having tried.³ It's clear that 'try' has such readings, whether or not it's context-sensitive or ambiguous.

But, while this is our official attitude, there is a stronger hypothesis that we find attractive, and which will help simplify the exposition below. The stronger hypothesis is that 'try' is indeed context-sensitive, and, furthermore, that there is a contextual resolution of 'try' that is minimal in the following sense: if for *any* context *c*, 'S is trying to ϕ ' expresses a truth in *c*, then 'S is trying to ϕ ' expresses a truth on the minimal resolution. Call the relation picked out by 'try' on this minimal resolution *pure trying*. According to our hypothesis, all trying entails pure trying.

This hypothesis has several nice features. But one reason to attend to it here is that, if it is correct, then pure trying has a distinguished place in the semantics of 'try', and would be a natural target for philosophical investigation. Of course, it is not in general true that if there is a unique weakest resolution of a philosophically interesting context-sensitive term, then we should investigate it. But in the case of trying, such a weak notion would be of special interest. An important question about the metaphysics of agency is what the minimal conditions are for its exercise. The weakest reading of 'try'—pure trying—is the barest form of trying, a place where agency gets its start.

³ In this sense our quarry is not the notion of *attempt* as used in the law, which requires the performance of certain canonical steps and not just mere mental preparation (think of what's required to be convicted of *attempted* murder). So we're aiming at a target quite different from that of (e.g.) Yaffe (2004, 2010).

Whatever one thinks about our stronger hypothesis and its consequences, we now ask the reader to understand our uses of ‘try’ for the rest of the paper as having a very permissive reading. In later sections of the paper we will sometimes step back to discuss how various “readings” of ‘try’ might be driving certain objections to our arguments. But for the most part we will simply talk about trying, and ask for the reader’s charity in understanding us as we intend.

3 Trying and believing

We now turn to arguing for the “Cartesian” conception of trying, on which what a person can try to do is comparatively unconstrained by their attitudes toward success. A good deal of our discussion here will build on arguments that have been known for some time now. But we will expand on these arguments and cast them in a new light. At the end of the section, we will explain how these ideas set up an important pillar of our defense of our central thesis, that trying is agentially perfect.

Suppose a person is trying to ϕ . What, if anything, follows about their beliefs as to whether they will ϕ ? Here are three possible answers, in order of decreasing strength:

- BELIEVE WILL If S is trying to ϕ , then S believes S will ϕ .
- BELIEVE PROBABLE If S is trying to ϕ , then S believes it is probable that S will ϕ .
- BELIEVE MIGHT If S is trying to ϕ , then it is compatible with what S believes that S will ϕ .

The first two principles are natural places to start, but they aren’t correct. People routinely try to do things they know have an enormously low probability of success: you can try to hit a hole-in-one, try to guess the combination of a safe, try to win the lottery, or try to write a bestseller. In each case you may know you’re almost certain to fail. But that’s no barrier to your trying.

The status of BELIEVE MIGHT is less obvious. It has many distinguished defenders.⁴ But we follow an equally distinguished tradition in rejecting it. We are moved by cases like the following. Suppose you’re in front of a brick wall. You are certain that no matter how hard you push it, it won’t fall over. Does it follow that you cannot try to push it over?⁵

We think not. We think you can try to push the wall over as part of an exercise regimen. Or you can try to push the wall over—perhaps as hard as you possibly

⁴ See, e.g., O’Shaughnessy (1973, 1980), Jones (1983), McCann (1986), and Adams (1995, 2007).

⁵ Cases like this are discussed in detail by Hampshire (1959, p. 134), Thalberg (1962, p. 54), McCormick & Thalberg (1967 p. 45), Harman (1986, p. 370), Ludwig (1992, 1995), Adams (1995), Hornsby (1995). For an example caught in the wild, see Elon Musk’s infamous request to his assistant Franz to “try to break the window” of Tesla’s supposedly bulletproof Cybertruck: <https://www.youtube.com/watch?v=LMWwImDX3ks>.

can—to prove that you are not strong enough to push it over. More generally, when you know you can't ϕ and want to prove that you can't, often the best way of doing so is to try as hard as you can to ϕ . You'll fail having done all you can to succeed, thereby illustrating your inability to ϕ .

These sorts of judgments seem to us about as clear as they come. But, as we have said, there is a rich tradition of rejecting them. Here we'll consider two kinds of objections.

Arguably the most prominent objection to such cases has been to claim that they are just not cases of genuine trying.⁶ Perhaps the agents are merely *pretending* to try, or *acting as if* they are trying. It's not always clear whether the proponents of this response intend it as an error-theory or instead as a theory of the pragmatics of ordinary 'try'-ascriptions. But either way we find it unsatisfactory.

First, suppose that you push on the wall as hard as you can and, to your surprise, it actually falls over. To our ears it would be rather strange for you to explain what happened with a speech like 'I wasn't trying to push the wall over, I was merely acting as if I was trying!'. The speech is no better if we substitute 'pretending to try' for 'acting as if I was trying'. Better would be to say something like 'Yes, I was trying to push the wall over—but only because I was sure I couldn't actually do it!'. This is especially clear if you were trying to push on the wall to demonstrate to shareholders of a wall-building company that this model of wall really could not be toppled. If in that context you merely pretended to try to push the wall over, the demonstration would be fraudulent: it is not at all impressive that someone pretending to try to push a wall over fails to do so. To prove that the wall is strong, you need to try to push it over.

Second, anyone who is trying as hard as they possibly can to ϕ is trying to ϕ . But there is no intuitive pull to the idea that one can only try as hard as one possibly can to ϕ if it is compatible with what one believes that one will ϕ . In fact, its being background knowledge that a person can't actually ϕ tends to make it especially natural to tell them to try *as hard as they possibly can* to ϕ .

Third and finally, the proposed alternative account of the intuitive judgments about appropriate uses of 'try' is problematically *ad hoc*. Consider a different verb—say, 'decide'. To our ears, if it is common ground that one is certain one will not succeed in pushing the wall over, one cannot felicitously say that one has *decided* to push it over. By contrast, one *can* felicitously say that one is going to act as if one has decided to push it over. If speakers were willing to reinterpret 'decide' as 'act as if one has decided' in these settings, then there should be no contrast in felicity between these two speeches. But there is. So we must not be willing to reinterpret claims featuring 'decide' in this way. But then what explains our alleged willingness to do so with 'try'? Absent a plausible and general story that explains why we would reinterpret reports about trying but not reports about deciding, the proposal is *ad hoc*.

⁶ See McCann (1986, n. 20), Adams (1986, p. 288), Adams & Mele (1992, p. 330), Adams (1997, p. 301).

A second way to resist our counterexamples to BELIEVE MIGHT is to admit that such examples do seem initially plausible, but to argue that the relevant judgments weigh against other equally strong ones and so must, on reflection, be rejected. In this case, the opposing judgment is the following simple one: there are lots of things it seems like we cannot even *try* to do, and the best explanation of this fact is something along the lines of BELIEVE MIGHT.

For instance: suppose there's a rock in front of you, and we ask you to try to lift it while staying still. You might reasonably reply 'I don't know how to do that', meaning that you don't know how to even *try* to lift it off the ground while staying still. This case might be thought to support BELIEVE MIGHT. For supposing the principle is valid, it would provide a good explanation of why it seems that you cannot try to lift rocks while remaining still: you can only try that which you take to be possible; you know (and thus believe) that your only means of lifting the rock is by physically picking it up; so you can't try to lift the rock while staying still.

But there is a better explanation of what's going on when people take themselves not to be able to try something in a case like this one, and this better explanation does not appeal to anything in the vicinity of BELIEVE MIGHT. The explanation is that the people in question simply haven't thought of a way to try to lift the rock while staying still. If they did think of one—through cogitation or the right sort of prompting—they'd realize that it's as easy to try to lift it while staying still as to try to do anything else.

To see this, suppose that instead of merely asking you to try to lift the rock while remaining still, we asked you to try to lift the rock using only your mind. Most people find it easy enough to do this. A typical response involves staring at the rock, scrunching up one's face, and imagining the rock gently lifting off the ground. (All to no effect, of course.) In doing this, it seems that you are trying to lift the rock using only your mind, and thus trying to lift the rock while staying still.⁷ Indeed, it seems that this is exactly the sort of thing you might do if you wanted to prove to someone (or even just to yourself) that you cannot lift the rock while staying still.

These sorts of considerations make us think that it's not just that you can try to do things when you believe you won't succeed, but that you can even try to do things when you know that it's metaphysically impossible that you'll succeed. For example: someone who knows that it is mathematically (and hence metaphysically) impossible to trisect an angle of 60° using only a straightedge and compass can still try to trisect it using only these means. They might try it to "prove" that they can't do it—as for instance in front of a class—or try it to gain intuition as they search for a new proof of this known result, or try it to see how good of an approximation to such a trisection they can get.⁸

This concludes our defense of the claim that one can try to do something even if one believes one will not succeed. We think this claim is significant in its own right. But it is also of central importance for our thesis that trying is agentially perfect

⁷ See Albritton (1985, esp. p. 245) for a related idea.

⁸ This example tells against other analyses of the cognitive condition on trying as well, e.g., Ludwig (1992, p. 268–9), Ludwig (1995), Hornsby (1995, p. 531–2).

(i.e., that a person tries to ϕ if and only if they try to try to ϕ). If a constraint like BELIEVE MIGHT were genuine, our thesis would not be true.⁹ For suppose that trying to ϕ really did require its being compatible with what one believes that one ϕ s. Now imagine that someone is (correctly) certain that squaring a circle is impossible, but (falsely) believes that BELIEVE MIGHT has false instances. Given these stipulations, it would follow that they can't try to square the circle: they believe it's impossible to square the circle, and by hypothesis anything one believes to be impossible is something one cannot try to do. But for all we've said they could still *try to try* to square the circle, since they believe BELIEVE MIGHT has false instances, and thus believe it's possible to try to do that which one believes to be impossible. So, if BELIEVE MIGHT were correct, there would plausibly be some things that one could try to try to do but not try to do (namely squaring the circle). If so, trying would not be agentially perfect.

Fortunately, as we have argued, BELIEVE MIGHT isn't correct, so this objection is defused in advance. But this style of objection doesn't arise just in the case of beliefs about one's success. Similar arguments could be made using other putative constraints on trying—most obviously, ones that tie what an agent can try to do to facts about what they want or intend to do. So, in the next section, we'll continue to lay the foundations for our defense of our main claim by arguing against these other constraints as well.

4 Trying, wanting, and intending

Suppose again that someone is trying to ϕ . Does it follow that they want to ϕ , in at least *some* natural sense of 'want'?¹⁰

Well, given that a person can try to ϕ even when they are certain that, no matter how hard they try, they won't ϕ , it isn't difficult to see how a person might try to ϕ even when they don't in any sense want to ϕ . The wall case demonstrates this plainly. The person who tries to push over the wall for exercise might in no intuitive sense want to push the wall over. In fact, they might strongly want *not* to push it over. Perhaps if the wall fell it would destroy a priceless work of art, or kill the wall-pusher, or kill the wall-pusher together with those inside. It's only because the person is certain they won't be able to push it over that they are willing to try.

These cases show that you can try to ϕ without wanting to ϕ when you are sure you won't ϕ . But what if you aren't sure? If you are trying to ϕ and it's compatible with what you believe that you will succeed in ϕ -ing, does it follow that there is some sense in which you want to ϕ ?

Once again we think the answer is no. Suppose you're playing your friend in chess. Your friend is a much stronger player—so much stronger, in fact, that you accurately take them to be a 99% favorite to win the match. However, you also know

⁹ Pollock (2002, p. 12) uses BELIEVE MIGHT in a similar way to argue that trying is not "infallibly performable". See also Hermes (2006, p. 64–5).

¹⁰ For related discussion see Mele (1990, 1991, 1994), Adams (1991, 1994a, b), and Hornsby (1995).

that there are two respects in which your friend is rather emotionally insecure. First, if they lose they'll be psychologically devastated. Second, if they detect that you aren't trying your hardest to beat them—and let's suppose they have a keen eye for a sham attempt—they'll again be psychologically devastated, whether or not they win. Knowing all this, and supposing that you enjoy playing chess but don't intrinsically care about whether you win or lose, can you (rationally) try to beat your friend in chess?

We think you can.¹¹ Indeed, it seems the best thing to do in these circumstances. Yes, if you manage to succeed the consequences will be bad. But the consequences will also be bad if you don't try your hardest to win. And the chance that you win conditional on trying your hardest is very low—low enough to make the gamble worth it. But none of this requires that you want to win; plausibly there is no sense at all in which you do. If that's right, then you can (rationally) try to ϕ without at all wanting to ϕ , even when you believe it's possible that you will succeed.

So much for the claim that trying to do something requires wanting to do it. What about the claim that trying to do something requires intending to do it?

Several authors have defended this idea.¹² But we've already seen strong reasons to doubt it. If one accepts the common view that if a person intends to ϕ , then it is compatible with what they believe that they will ϕ , then trying can entail intending only if BELIEVE MIGHT is true.¹³ Likewise, if one accepts the common view that if one intends to ϕ , then there is *some* sense in which one wants to ϕ , then trying can entail intending only if trying entails wanting. Given these common views, our arguments against BELIEVE MIGHT and the claim that trying entails wanting are also arguments against the claim that trying entails intending.

But even independent of general principles like these, the various cases discussed so far also present intuitive counterexamples to intention-based constraints on trying. A person who tries to push an unmovable wall down for exercise, or to show that it can't be done, intuitively just does not intend to push the wall down. And although you may be trying your hardest to beat your friend in chess, it seems clear that you do not intend to win. So trying does not entail intending.

This concludes our defense of the Cartesian conception of trying, according to which trying is not constrained by particular attitudes toward success. Our putative counterexamples notwithstanding, these constraints on trying have historically attracted a good amount of support. So, one might ask, why have they been so popular?

We suspect the answer has to do with a distinctive feature of our counterexamples: they all rely crucially on the fact that one can have motivating reasons to try to ϕ even when one has no motivating reasons to ϕ (and perhaps quite strong

¹¹ Hornsby (1995, p. 529 and n. 11) makes this point using a version of Bratman (1987)'s video game case.

¹² See Hampshire (1959, p. 107) McCann (1975, 1986, 1989), Adams (1986, 2007), Adams and Mele (1992), Ludwig (1992, 2021), Yaffe (2004, 2010), Grano (2011, 2017), Shepherd (2016, p. 422).

¹³ But see Thalberg (1962), Hedman (1970), McCann (1986, 1991), Ludwig (1992) and Buckwalter et al. (2021) for dissent.

motivating reasons *not* to ϕ).¹⁴ And we suspect that this kind of trying is in a certain sense exceptional, since we find it plausible that “paradigm” cases of trying are those in which a person is trying to ϕ *because* they take it to be a necessary means toward ϕ -ing. So perhaps something like the following principle holds:

PARADIGMATIC TRYING If S is rational and all S’s reasons for trying to ϕ are also reasons for ϕ -ing, then if S tries to ϕ , (i) it is compatible with what S believes that S will ϕ , (ii) S wants to ϕ , and (iii) S intends to ϕ .

To be clear, we are not suggesting that non-paradigmatic tryings are in any interesting sense rare. We suspect that the phenomenon brought out by the wall and chess cases is actually quite common, arising whenever we shoot for the stars (“Try to punch through him”, “Try to touch the ceiling”, “To see how close you can get, try to draw a perfect circle”, “Try to count to a million, you’ll fall asleep soon enough”...). However, we also suspect that the capacity for this kind of non-paradigmatic trying is a mark of cognitive sophistication, much in the way that the capacity to resist the testimony of one’s senses (in cases of suspected or known illusion) is a mark of cognitive sophistication. For much of the animal kingdom, trying may be a reliable indicator of belief, desire, and intention. But adult humans can think about the possible effects of trying in ways that many animals may not. For creatures like us, the connections between trying and these other states are more tenuous.¹⁵

5 Agential perfection

As we discussed at the end of §3, our arguments for the Cartesian conception of trying are necessary for responding to an important objection to our main claim: that trying is agentially perfect. We’ve now completed our preemptive response to that objection, so it’s time to turn to the main claim itself. We’ll start by bringing the

¹⁴ See McCann (1986, p. 206–208), Mele (1990, p. 252) (cf. Mele (1992, p. 61f)), and Hornsby (1995).

¹⁵ Similar points apply to O’Shaughnessy’s (1973, p. 369) claim that trying “is constituted by doing, intentionally and with just that purpose, whatever one takes to be needed if, the rest of the world suitably cooperating, one is to perform the action”, as well as to Ludwig’s (1995, p. 569) “least resistance principle”, which says that “A is a trying to B only if it is conceived by the agent to be the path of least resistance to his end...”. As restricted to paradigmatic tryings, both claims seem plausible enough. But one can find counterexamples in cases where one’s reasons for trying to ϕ are distinct from one’s reasons for ϕ -ing. Take the wall case from earlier, but now imagine the wall has a strong section and a weak section. The strong section is strong enough that no matter how hard you push on it, the wall won’t fall over, whereas the weak section is weak enough that any amount of pushing will cause the wall to collapse. Knowing this, you know that pushing the wall over will require pushing on the weak spot. But you still could try to push the wall over by trying to push it over in the strong spot, whether to get some exercise, or to show that it can only be pushed over in the weak spot. Similarly, an excellent tennis player can try to beat a weak player with their non-dominant hand. In doing so, they are trying to win, although not in a way that involves the path of least resistance.

property of agential perfection into focus. In the next section we'll argue that trying has this property.

We opened the paper with a picture on which the divinity's attempts are immune from failure: if the divinity tries, the divinity succeeds. By contrast, we humans routinely try and fail. But still, one might ask, is there a class of doings such that even we cannot fail to do them if we try?¹⁶ To have a name for this condition, we will say that:

ϕ -ing is *essentially successful* iff: necessarily, if S tries to ϕ , then S ϕ s.

As we illustrated earlier, many doings lack this property. People try to stay calm and fail. They try to get high and fail. They try to assuage a crowd and fail. None of these doings is essentially successful.

But arguably some doings are. Perhaps, for instance, trying to think about Confucius requires having Confucius in mind enough that you're already thinking about him. If so, then thinking about Confucius would be essentially successful.

The class of essentially successful doings is interesting in its own right. But our aim here is to characterize a domain in which people approximate the kind of divine agency with which we began, and so far we're still not there. For even supposing that thinking about Confucius is essentially successful, it can happen without being attempted—as when one wakes up in a sweat considering what the sage would say about one's relationship to one's parents, or finds oneself thinking about him in response to someone saying his name. Thus, even if a doing is essentially successful, it does not thereby rise to the level of divine agency that interests us here. The divinity does not just do whatever they try to do; the divinity also doesn't do something unless they try to do it.

These reflections motivate a second condition:

ϕ -ing is *essentially attempted* iff: necessarily, if S ϕ s, then S tries to ϕ .

In a way, it's easier to come up with at least initially plausible examples of doings that are essentially attempted. Maybe *defrauding the state* is such an example: perhaps you can't defraud the state unless you try to. More generally, Anscombe (1957, §47) argues that some doings are *essentially intentional*: if a person performs them, they perform them intentionally. Putatively included in this class are *greeting*, *marrying*, and *promising*. Though we have doubts as to whether these sorts of things really are essentially *intentional*, perhaps they are at least essentially *attempted*.

To be clear: none of these examples is "essentially attempted" on all ways of interpreting 'try'. A person can certainly defraud the state without having tried *hard* to do so; some people are just naturals. If 'try' is context-sensitive or ambiguous, and there are readings on which 'try' expresses what is expressed by 'try hard', then there would be readings on which defrauding the state is not "essentially attempted", since it can be done effortlessly. But as we emphasized earlier—and will from now

¹⁶ Here we use the unlovely expression 'doings' in place of 'actions', to avoid contested questions about the characterization of action. We intend 'doings' rather expansively, covering roughly anything that can be expressed by a verb phrase in English, including (for instance) dying, digesting, and drooling.

on be especially important—our focus is exclusively on very permissive uses of ‘try’: in particular, on those where it expresses *pure trying*. The fact that someone can defraud the state without having put much effort into it is thus no counterexample to the intended interpretation of the claim that defrauding the state is essentially attempted.

The properties of being essentially successful and of being essentially attempted are both, in our view, of great interest on their own. But we won’t say much about them separately. Instead, we’ll be interested in our promised notion of agential perfection, which we obtain by putting the two conditions together:

ϕ -ing is *agentially perfect* iff: necessarily, if S tries to ϕ , then S ϕ s, and if S ϕ s, then S tries to ϕ .

To further illustrate the contours of this notion, it is worth contrasting it with more familiar notions connected to ability and control.

For ability: to say that ϕ -ing is agentially perfect is not to say that necessarily, one is able to ϕ . To see why, consider the claim (which we will argue for in a moment) that trying is agentially perfect. This claim entails that if one tries to try to ϕ , one will succeed in trying to ϕ . But it does not guarantee that one will be able to try to try to ϕ , or be able to try to ϕ , and in general one may not be. A person might be unable to try to ϕ because a malevolent neuroscientist has implanted a certain kind of chip in their brain that causes a fatal stroke the moment it detects whatever neural events would normally precede their trying to ϕ (cf. e.g. Storrs-Fox (2019), Koon (2020), building on Frankfurt (1969)). Or they might be unable to try to ϕ because they aren’t in a position to grasp ϕ -ing under any relevant guise. For example, for some sufficiently large prime number n —which contains, say, more digits than there are atoms in the universe—the action of *proving that n is prime* is plausibly something a person cannot even try to do, since they might simply lack any relevant guise under which they can apprehend the question of whether n is prime. But even if these cases are counterexamples to the claim that we are always able to try, they are no counterexample to the claim that trying is agentially perfect.

Similar points apply to control: to say that ϕ -ing is agentially perfect is also not to say that necessarily, ϕ -ing is always *under a person’s control*. Whether something is under a person’s control often depends on the nature of the motivational states that would explain their doing it. When these states are the result of forces that are sufficiently alien to the person—coercion, brainwashing, psychosis, and so on—it will often not be true that what they did was under their control. This applies just as much to what the agent *tries* to do: someone who is brainwashed might try to do all sorts of things they would never have tried to do had they not been brainwashed. But it doesn’t follow that they didn’t try to try to do what they did. Trying can be agentially perfect even if it is not always under one’s control.

6 Trying is agentially perfect

With the notion of agential perfection clarified, we turn at last to our arguments for our main claim: that trying is agentially perfect. We will argue for this claim by arguing for a strictly stronger thesis about the relationship between trying and trying to try, namely:

IDENTITY To try to ϕ just is to try to try to ϕ .

IDENTITY implies that trying is agentially perfect, but the converse need not hold: it is conceptually possible that trying and trying to try are merely necessarily equivalent rather than identical. But while it is conceptually possible for there to be such a gap, it is hard for us to see a principled reason for believing in one. So here we'll take the agential perfection of trying to stand and fall with IDENTITY, and take arguments for either one of these theses to be arguments for both.

Our first argument for these claims begins from the fact that IDENTITY follows from a natural (albeit simplistic) picture of trying.¹⁷ On this simplistic picture, to try to ϕ just is to satisfy the disjunctive condition of *beginning or being in the process of ϕ -ing*. If this picture were correct, then IDENTITY would hold, because to begin or be in the process of beginning or being in the process of ϕ -ing just is to begin or be in the process of ϕ -ing. To see this, note that we are typically *not* inclined to give a certain answer to the question 'During what period of time had you begun or been in the process of quitting drinking?', but a different answer to the question 'During what period of time had you begun or been in the process of beginning or being in the process of quitting drinking?'. The latter question just seems like a bizarre way of asking the former question. Presumably this is at least in part because we do not distinguish the times at which one counts as beginning or being in the process of something from the times at which one counts as beginning or being in the process of beginning or being in the process of it. As such, the simplistic picture plausibly implies that to try to ϕ just is to try to try to ϕ .

This simplistic picture is, however, too simple. Sometimes we can begin or be in the process of ϕ -ing without trying to ϕ : if one begins to fall by accident, or is in the process of doing so, one is typically not trying to fall. Similarly, one may be in the process of digesting one's food even if one isn't trying to do so.

Still, the attractions of the simplistic picture suggest that it captures something important about trying. There is a structural analogy between trying to ϕ on the one hand and beginning or being in the process of ϕ -ing on the other: the former is something like the "agential" version of the latter. This structural analogy, together with the fact that the analog of IDENTITY holds for beginning or being in the process, is what underlies much of the intuitive force behind IDENTITY.

Our second, more direct argument concerns linguistic judgments about iterated 'try' reports. Consider first the claim that trying is essentially successful, that is:

¹⁷ For related ideas, see Hermes's (2006, pp. 63–4) discussion of McCann (1986), as well as Hermes's (2021, p. 346).

COLLAPSE Necessarily, if S is trying to try to ϕ , then S is trying to ϕ .

In contexts where ‘try’ gets its minimal reading, it is hard to see what a counterexample to this principle would look like. For example, if in such a context you were to ask me ‘Did you try to lift that rock with your mind?’, it does not seem I could felicitously reply with ‘No, but I did *try to try* to lift it with my mind’. If I second-order tried to lift the rock but failed to first-order try to lift it, then either my failure to first-order try was by design or it wasn’t. If it *was* by design, then presumably I must have had some sense in advance of what it would look like to second-order try to lift the rock with my mind while failing to first-order trying to lift it with my mind. But what could such a sense even come to? On the other hand, if my failure to first-order try was *not* by design, then presumably it was due either to some kind of ineptitude on my part, or to some amount of uncooperativeness from the world. But again, what could such ineptitude or uncooperativeness even look like here? Would I have succeeded in first-order trying had I scrunched up my face differently? In the absence of compelling answers to these questions, we find it hard to see how COLLAPSE could have false instances.

Similar points apply to the claim that trying is essentially attempted, that is:

ITERATION Necessarily, if S is trying to ϕ , then S is trying to try to ϕ .

An argument in favor of ITERATION is that it explains why, if someone tells you to try to try to ϕ , you can comply by simply trying to ϕ —at least in contexts where ‘try’ gets its minimal reading. Suppose I tell you ‘You should really try to try something you know to be impossible; it’s liberating’. A good way of complying with this advice is to go and try to push down the wall, or to try to trisect a 60° angle. But this wouldn’t be a good way of satisfying the request if trying did not entail trying to try.

These two arguments make us think that IDENTITY and the claim that trying is agentially perfect have a kind of default status. But the two arguments on their own wouldn’t have gotten us all the way to believing these claims, and we don’t expect them to have done so for you either. The full case for the claims rests not just on the positive arguments in their favor, but also on the fact that, once one attends to the various readings of ‘try’ and takes care with the logic of trying, what might have seemed to be obvious arguments against them turn out to be weak. It is this fact, together with the positive arguments just presented, that moves us on balance to endorse IDENTITY and the agential perfection of trying.

In defense of the claim that the objections are weak, we’ll present the four we find most compelling, and argue that they fail.

First, O’Shaughnessy (1973) argues that claims about whether one tries to try are unintelligible.¹⁸ But we think this can’t be right: such claims may be hard to assess, but they’re not unintelligible. Indeed, as we just saw, a request along the lines of

¹⁸ McCormick and Thalberg (1967), Hunter (1987), also voice worries along these lines. Cf. Mele (1992, p. 64).

‘You should really try to try the impossible; it’s liberating’ can be understood and carried out.

A second and more serious worry about ITERATION (and thus about IDENTITY) arises from cases like the following.¹⁹ A yoga instructor tells their students not to try to breathe from their diaphragm, but instead to breathe from their diaphragm naturally and automatically. A student tries to follow the teacher’s directions, but fails: they breathe through their diaphragm in a way that is conscious and unnatural. It seems plausible that the student tried to breathe through their diaphragm, and that in virtue of trying to do this failed to follow their teacher’s instructions. But it also seems plausible that the student did not try to do this: that is, that they did not try to try to breathe through their diaphragm. So, the objection goes, the student is proof that one can try without trying to try.

We can see two ways of supporting the claim that the student does not try to try to breathe through their diaphragm, but we think neither of them works. A first way starts from the intuitive observation—which we are happy to concede—that it is natural to describe the student as trying *not* to try to breathe through their diaphragm. It then adds in the general principle that if you’re trying not to ϕ , then you’re not trying to ϕ . From the observation and the principle together, it follows that the student is not trying to try to breathe through their diaphragm.

But we reject the principle. People can knowingly both try to ϕ and try not to ϕ , and even be rational in doing so. Suppose that, from long experience, you’ve learned that the best way for you to hit the bullseye in a game of darts is for you to try not to hit it, and, in particular, to try instead to hit a point about a foot above it. In a tense game, when you need a bullseye to win, you may try to hit the bullseye by trying not to hit it. Since, in such a case, you are both trying to hit it and trying not to, we deny that trying not to do something implies that one is not trying to do it. Accordingly, the student may be trying not to breathe through their diaphragm, but it doesn’t follow that they are not trying to breathe through it.²⁰

The second way of motivating the claim that the student does not try to try to breathe through their diaphragm is on the basis of a direct intuition: it’s supposed to just be intuitively obvious that they aren’t trying to try to do this.

But we’re skeptical that there is any direct intuition to this effect. We suspect that those who claim to have this intuition are interpreting ‘try’ as meaning something like *consciously try*. On that reading, it does seem correct to say that the student is consciously trying to breathe through their diaphragm (against the teacher’s instructions), but also that they aren’t consciously trying to consciously try to breathe through their diaphragm. Crucially, though, not all trying is conscious trying: tennis players can try to put certain kinds of spins on their shots without thinking about it, musicians can try to play in tune without thinking about it, and writers can try to

¹⁹ This particular case comes from correspondence with Adam Elga. Thanks to Elga and Simon Goldstein for discussion here.

²⁰ Ulysses’ encounter with the Sirens provides another possible example of the same form. Plausibly, Ulysses was simultaneously trying not to heed the Sirens (that’s why he was tied to the mast) and trying to heed them (that’s why he was struggling to free himself).

spell their words correctly without thinking about it. So the fact that the student consciously tries without consciously trying to consciously try is no threat to ITERATION on its intended reading, which concerns pure trying, not conscious trying.

Moreover, if ‘try’ *isn’t* interpreted as meaning something like ‘consciously try’, there is no longer any direct intuition to the effect that the student is not trying to try to breathe through their diaphragm. In fact, when one focuses on such readings, one can argue directly that the student *is* trying to try to breathe through their diaphragm. After all, the student is trying to follow their instructor’s advice, which is to try to breathe through one’s diaphragm automatically rather than consciously. So the student is trying to try to breathe through their diaphragm automatically rather than consciously, and thus, plausibly, trying to try to breathe through their diaphragm *tout court*.

A third objection to ITERATION and IDENTITY claims that simple animals lack the conceptual wherewithal required to try to try to do anything. This objection seems to rest on two ideas: first, that trying requires a rich ability to conceive of the objects of one’s attempts under a particular description; and second, that animals lack such an ability. We are skeptical of both.

With respect to the first, suppose a bear escapes from the zoo and puts a girl’s smartphone into its mouth. A distraught uncle calls out to the zookeeper: “Sir, please help me, the bear is trying to eat my niece’s smartphone!”. Does the truth of this ascription show that in the relevant sense the bear has the concept of a smartphone, or of a niece? Presumably not. So we doubt that trying requires a rich ability to conceive of the objects of one’s attempts under a particular description.

With respect to the second idea, it is far from clear that animals lack the concept of trying. We routinely say things like ‘The dog knows that the cat will try to take his food’, ‘The spider knows that the fly is trying to escape’, and so on. Some may be inclined to give an error-theory about these speeches. But as far as we can tell, any argument for this kind of error-theory is no less compelling as an argument for an error-theory about first-order ‘try’-ascriptions to animals—in which case animals would still pose no threat to ITERATION. So we’re happy to say that animals have the conceptual wherewithal to try to try, at least if they ever try at all.

The fourth objection to IDENTITY is the one we raised at the close of §3 and that was a guiding thought behind our discussion there and in §4. To recall, the basic idea was that if trying requires believing that one will succeed, or wanting or intending to succeed, then we could imagine someone having these beliefs, wants, or intentions toward their trying to ϕ while lacking them toward their ϕ -ing, thereby generating a counterexample to COLLAPSE. In our view, this style of argument is extremely important, and that’s why we took some time with it above. But we argued that each of the relevant principles fails, and so we take the objection to have been defused in advance.

We’ve presented and responded to a number of objections to IDENTITY. But perhaps more important than the specifics of our responses to these objections is the following general point. Once one takes care with the relevant readings of ‘try’, the objections above don’t seem to undermine the motivations for COLLAPSE, ITERATION, or even IDENTITY. In this sense, they’re quite different from the objections we raised in Sects. 3 and 4 against various attitudinal constraints on trying. A moral we drew

from those earlier cases is that one can have reasons for trying to ϕ that have little or nothing to do with one's reasons for successfully ϕ -ing. With this moral in view, one loses the sense that there *should* be relevant necessary connections between one's trying to ϕ and one's beliefs, wants, and intentions about one's success in ϕ -ing. By contrast, the putative counterexamples to COLLAPSE, ITERATION, and IDENTITY we've just discussed do not seem to be underwritten by a similarly illuminating moral. We thus take the case in favor of this strong, simple theory of trying to be in good standing.

7 Agential perfection in action

We now turn to some consequences of IDENTITY, with the aim of illustrating the broader philosophical significance of this thesis. Our main focus will be on its consequences for debates about basic action and the logic of intentional action. But we will also briefly discuss consequences concerning the nature of options in decision theory.

7.1 Acting by trying

As a prelude to our discussion of basic and intentional action, we begin by introducing a key connection between intentional action and trying.

Reflecting on paradigm cases of intentional action, it is natural to think that acting intentionally requires trying, that is:²¹

INTENTIONAL \rightarrow TRYING If S ϕ s intentionally, then S tries to ϕ .

But there are in fact strong reasons to reject this principle, stemming from a style of example discussed by Harman (1976), Bratman (1984), and Knobe (2003, 2006), among others. To take Bratman's case, suppose you plan to run a marathon. You realize that by running it, you'll almost surely end up destroying your shoes, which happen to be a family heirloom. You don't want to destroy your shoes—they're a family heirloom!—but you love running and your family is long dead, so you treat the likely destruction of the shoes as a necessary evil. The time of the marathon comes. You run it and, lo and behold, your shoes are destroyed. Plausibly, you did not try to destroy your shoes, but you nonetheless destroyed them intentionally.

At least for the sake of argument, we will accept that these cases are counterexamples to INTENTIONAL \rightarrow TRYING. Still, we think there is an intuitive distinction between cases like these and more paradigmatic cases of intentional action like those that motivate INTENTIONAL \rightarrow TRYING. We'll call the intentional action involved in cases like Bratman's *intentional in the secondary way*, and the rest we'll call *intentional*

²¹ Proponents of this claim include, e.g., O'Shaughnessy (1973, 1980), McCann (1975), Gorr (1979), Hornsby (1980, 2010).

in the primary way. Effectively by definition, then, your running the marathon was intentional in the primary way; your destroying your shoes was intentional in the secondary way. Likewise, your arriving at the marathon's starting location was presumably intentional in the primary way; your missing your lecture to run the marathon was presumably intentional in the secondary way.

With this distinction in hand, we offer the following as a conservative modification to INTENTIONAL \rightarrow TRYING:²²

PRIMARY INTENTIONAL \rightarrow TRYING If S ϕ s intentionally in the primary way, then S tries to ϕ .

In paradigmatic cases of intentional action—that is, cases of acting intentionally in the primary way—if one ϕ s intentionally, one tries to ϕ . But crucially it is not that one tries to ϕ and just coincidentally also ϕ s; instead, in these cases one ϕ s intentionally *by* trying to ϕ . Indeed, PRIMARY INTENTIONAL \rightarrow TRYING seems to be in part motivated by this stronger idea:

BY TRYING If S ϕ s intentionally in the primary way, then S ϕ s by trying to ϕ .²³

We won't say much in defense of this principle, but we think the case for it becomes clear when one reflects on perhaps the most obvious class of putative counterexamples: simple bodily movements. Some might want to reject BY TRYING because they judge that we can (for instance) intentionally move our fingers without trying to. But we are skeptical of this judgment. It's true that simple bodily movements are often effortless. But as we argued in §2, the mere fact that we (stereotypically) perform simple bodily movements without effort is not good evidence that we can perform them without trying to (cf. Ludwig (2021, §6)). On the contrary, for many cases of intentionally moving one's body, it seems perfectly reasonable to assert conditionals like 'Had I not tried to move that way, I wouldn't have moved that way'. If this sentence is true, as it seems to be, then one must have tried (and succeeded) to move one's body. Moreover, in a poll of the room, one does not expect to get a larger number in answer to the question 'How many people here moved their fingers intentionally?', than in answer to 'How many people here moved their fingers by trying to move them?'. Finally, it seems typical for one's knowledge that one has moved one's fingers intentionally rather than accidentally (or because of an external force) to be at least partly grounded in one's knowledge that one moved one's fingers by trying to. Outside of the special secondary cases, then, it seems that what we do intentionally we do by trying.

²² Those who deny that examples like Bratman's are counterexamples to INTENTIONAL \rightarrow TRYING can simply delete 'in the primary way' in its key occurrences in what follows; nothing essential will turn on whether any action is intentional in a non-primary way.

²³ For sympathetic discussion of this principle, see, e.g., Hornsby (1980), O'Shaughnessy (1980), Pietroski (2000).

7.2 Basic action

Having made a tentative case for BY TRYING, we now turn to basic action. The correct characterization of the technical notion of “basic action” is hotly contested, and we cannot begin to engage seriously with this controversy here.²⁴ Instead, our goal will be to show how one reasonable conception of basic action is supported and illuminated by IDENTITY.

The core idea of the conception we’ll focus on is that to perform a basic action is to do something intentionally and not by doing anything else. More exactly, we will say that S ϕ s as a basic action iff: (i) S ϕ s intentionally in the primary way; and (ii) for all ψ , if S ϕ s by ψ -ing, then for S to ϕ just is for S to ψ .²⁵

A striking consequence of this definition, together with BY TRYING, is that if one ϕ s as a basic action, then for one to ϕ just is for one to try to ϕ . In simplified slogan form: all basic actions are tryings. To see this, note that if one ϕ s as a basic action, then (by (i)) one ϕ s intentionally in the primary way. BY TRYING then implies that one ϕ s by trying to ϕ . But now (by (ii)), for every ψ such that one ϕ s by ψ -ing, for one to ϕ just is for one to ψ . So if one ϕ s as a basic action, it must be that for one to ϕ just is for one to try to ϕ .

So BY TRYING and the conception of basic action we’ve focused on together impose tight constraints on basic action. But are those constraints too tight, ruling out the very possibility of basic action? We’ll now show how IDENTITY plays a key part in an argument that they do not.

We’ve established that if a person ϕ s as a basic action, then for them to ϕ is for them to try to ϕ . So if they ϕ as a basic action, they must also try to ϕ as a basic action. Our question now is whether one ever tries to ϕ as a basic action. By definition, a person can try to ϕ as a basic action if and only if: (i) they intentionally try to ϕ in the primary way and (ii) for any ψ , if they try to ϕ by ψ -ing, then for them to try to ϕ is for them to ψ . We think it’s clear that (i) is often satisfied. Often when people try to do things, they intentionally try to do them (in the primary way): people

²⁴ For discussion see, e.g., Danto (1965), Goldman (1970), Davidson (1971), Hornsby (1980, Ch. 5–6), Thompson (2008), Lavin (2013), Amaya (2017), Ludwig (2016, Ch. 6.1, 6.6), Kelley (2024a).

²⁵ This characterization is a standard “teleological” conception of basic action, with one important caveat: it treats *basicness* adverbially—as a way of doing things—rather than as a property of *token doings* (i.e., particular doings). The adverbial characterization is consonant with our practice, throughout the paper, of discussing what is expressed by finite-verb phrases, and our corresponding uses of property identity, as in IDENTITY. This shift away from a focus on token actions allows us to set aside notoriously tricky questions about how “fine-grained” token events are (cf. Davidson (1971), Hornsby (1980), and Payton (2021b, pp. 14–18)). It also enables us to describe a trying-theoretic account of basic action while continuing to be neutral (as we have been throughout) on the question of whether there are such things as *particular tryings* (see Ruben (2013, 2015, 2016, 2018, 2022) for the claim that there are not, with critical discussion by Payton (2021a)). Our adverbial characterization of ϕ -ing as a basic action would entail a corresponding characterization of token basic actions, if one made the natural further assumption that whenever someone ϕ s as a basic action, there is a particular event of their ϕ -ing that is a basic action. But we will not make this assumption here. Sometimes we will speak below, inexactly, of “performing basic actions” and the like, but we ask the reader to understand these statements as abbreviations for the corresponding (clunkier, but official) adverbial claims. Thanks to two anonymous referees for extremely helpful discussion here.

intentionally try to pay their taxes, intentionally try to pass tests, and intentionally try to go to space. So let's assume for now that (i) is met. (Later, we'll argue that it's *always* met.)

Given (i), the focus falls on (ii). And it's here that IDENTITY plays a starring role. BY TRYING says that if one intentionally tries to ϕ in the primary way, then one tries to ϕ by trying to try to ϕ . On the face of it, this would appear to lead to a problematic regress: if for every ϕ , *trying to ϕ* were distinct from *trying to try to ϕ* , then no one could ever try to ϕ as a basic action, since one would always try to ϕ by doing something else. But in the presence of IDENTITY there can be no such regress: since for every ϕ , to try to ϕ *just is* to try to try to ϕ , the fact that one tries to ϕ by trying to try to ϕ does not mean that one tries to ϕ by doing something else. Given IDENTITY, 'to try to ϕ ' and 'to try to try to ϕ ' are just two descriptions of the very same thing.

This is a step in the right direction. Still, IDENTITY doesn't quite get us all the way to the claim that sometimes, we try to ϕ as a basic action. For while we've seen that there's no problem with trying to ϕ by trying to try to ϕ (and trying to try to try to ϕ , and so on), it might still be that one always tries to ϕ in part by doing something else altogether. Suppose, for instance, that someone tries to get rich by investing in plastics. If so, then even granting that they also try to get rich by trying to try to get rich (and so on), they will still not have tried to get rich as a basic action, since they tried to get rich by *investing in plastics*, and for them to try to get rich is not the very same thing as for them to invest in plastics.

So it is not always true that when a person tries to do something, they try to do it as a basic action. But it remains plausible that at least *sometimes* people try to do things as a basic action. The person who tries to lift a rock with their mind strikes us as a clear case: plausibly the person intentionally tries to lift the rock with their mind (in the primary way), and plausibly there is nothing else they do by which they try to lift the rock with their mind. So, it seems, they try to lift the rock with their mind as a basic action.

If that's right, our conception of basic action is non-vacuous. But this still doesn't yet show that it does all the work it's meant to do. A popular view, which we're inclined to endorse, is that all (primary) intentional action eventually "bottoms out" in basic action—i.e., if a person ϕ s intentionally in the primary way, then, for some ψ , they ψ as a basic action.²⁶ Does this claim also hold? We don't have space to give a full dress argument for it here. But we think our conception of trying makes it plausible. Briefly: we saw above that one might try to get rich, but not as a basic action (as when one tries to get rich by investing in plastics). But even so, one might still try to invest in plastics as a basic action (if one tries to invest in them only by trying to do so). And even in the normal case where that isn't true—say because one tries to invest in plastics by pressing a button—one might still try to press the button as a basic action (because one tries to do so, but not by doing anything else). Or perhaps the process continues on a bit longer. But even if it does, we see little

²⁶ This is the view of Danto (1965, 1979), and has been influentially endorsed by Davidson (1973), and Hornsby (1980, chs. 5–6). But it is not without its detractors—see, e.g., Sneddon (2001), Thompson (2008), and Lavin (2013).

reason to think there are any cases where it goes on forever. So long as it always ends somewhere—with one trying to do something only by trying to try to do it, and so forth—the claim that all (primary) intentional action eventually “bottoms out” in basic action can be vindicated.

So our view of trying motivates not only the claim that to do something as a basic action is to try to do something—i.e., that all basic actions are tryings—but also the claim that some of what we try to do we try to do as a basic action. More than this: it leads to a natural story about how it could be that intentional action always “bottoms out” in basic action. Key to this overall picture, which puts together BY TRYING with a fairly standard conception of basic action, is IDENTITY, which guarantees that the claim that one tries by *trying to try* is fully compatible with the claim that one tries but not by doing anything *else*.

7.3 The logic of intentional action

We now turn to consequences of IDENTITY for structural laws governing intentional action. We’ll begin by arguing for a claim that surfaced a moment ago: that whenever a person tries to do something, they intentionally try to do it (in the primary way)—i.e., that trying is essentially intentional in Anscombe’s sense.²⁷ Since it’s clear that if a person intentionally ϕ s (in the primary way), they must in fact ϕ , any argument for the claim that trying is essentially intentional is an argument for:

TRYING IS INTENTIONAL Necessarily, S tries to ϕ if and only if S intentionally tries to ϕ in the primary way.

And just as we suggested in §6 that the agential perfection of trying stands and falls with IDENTITY, here too, we think that TRYING IS INTENTIONAL stands and falls with a parallel identity:

INTENTIONAL IDENTITY To try to ϕ just is to intentionally try to ϕ in the primary way.

These two principles strike us as intuitively plausible—at least to the extent that we have intuitive judgments about the intentionality of tryings at all. The idea that one might unintentionally try to do something seems about as bizarre as the idea that one might unintentionally choose to do something. It just sounds strange to actively deny that one’s trying was intentional. If we were to ask you ‘Why did you try to ϕ ?’, we would be quite surprised to hear that you didn’t mean to, or that your trying was a mere accident. Similarly, it’s hard to see how a person could try to do something intentionally but not in the primary way. Under what conditions might you try to ϕ not as a means toward some end, say, but instead as a merely foreseen consequence of something else you do intentionally?

²⁷ For further discussion of the intentionality of trying, see also McCann (1974, 1975, 1986), McGinn (1982, ch. 8), Ginet (1990), Adams (1995, 1997). For recent discussion of essentially intentional doings, see, e.g., Beddor and Pavese (2021).

Beyond these intuitive judgments—which are at best attenuated—there is also a more direct argument for TRYING IS INTENTIONAL, using a natural further assumption about the conditions under which one can do something *without* doing it intentionally. The assumption is that if it's true that a person ϕ -d but not true that they intentionally ϕ -d, then either (i) they didn't try to ϕ —as in the case of accidental action—or (ii) they tried to ϕ , but succeeded in a way that was relevantly abnormal, deviant, or lucky.²⁸ This assumption, together with the claim that trying is agentially perfect, implies the controversial direction of TRYING IS INTENTIONAL (that if you try, you intentionally try). For, if trying is agentially perfect, neither (i) nor (ii) can hold when ϕ is a trying. If a person tries to ϕ , it can't be that they don't try to try to ϕ (by necessity, whenever they try, they try to try), and it can't be that the connection between their trying to try to ϕ and their trying to ϕ was in any sense abnormal, deviant, or lucky (by necessity, whenever they try to try, they try). So a person who tries to ϕ must intentionally try to ϕ in the primary way.

As we said, any case for TRYING IS INTENTIONAL seems to us a case for INTENTIONAL IDENTITY. And the latter principle has an important consequence. It immediately implies that to try to ϕ just is to intentionally...intentionally try to ϕ , where the ellipsis can be filled in with an arbitrary number of repetitions of 'intentionally'. In itself, this is a striking conclusion. There are a wide array of important notions in philosophy for which there are no true cases of indefinite iteration on "non-trivial" conditions. Perhaps most famously, those who follow Williamson (2000) hold that that no one can have arbitrarily many iterations of knowledge of non-trivial truths. Given INTENTIONAL IDENTITY, intentional action would differ from knowledge (so understood): the former, unlike the latter, would allow arbitrary iterations on some non-trivial conditions.²⁹

²⁸ Here we have in mind the kind of cases of successful action "by fluke" discussed in, e.g., Chisholm (1971), Davidson (1973), Armstrong (1973), and Harman (1976). A person who defuses a bomb by (correctly) guessing the code, and a person who tries to make their friend laugh by sliding Kramer-style into a room, only to trip and fall, earning a round of laughs, both in different ways try to ϕ and succeed because they tried to. But intuitively neither ϕ s intentionally.

²⁹ It is perhaps illuminating to consider how the agential perfection of trying undermines an argument, broadly parallel to Williamson's argument about knowledge, against the possibility of there being things we intentionally...intentionally do, for arbitrary iterations of 'intentionally'. The argument turns on two main premises: first, that if one ϕ s intentionally, then one could not have easily tried but failed to ϕ ; and, second, that for any ϕ , there is some number of repetitions of 'could have easily' such that one could have easily been such that...one could have easily been such that one tried but failed to ϕ . Taken together, these claims imply that for any ϕ , there is a greatest number of iterations of 'intentionally' for which one can intentionally...intentionally ϕ —contradicting INTENTIONAL IDENTITY. Shepherd and Carter (2021) and Carter and Shepherd (2022) develop counterexamples to the first premise of this argument, and we are sympathetic to their cases (see also Ludwig (2016, Ch. 7.4)). Roughly the idea is that, while true beliefs that could have easily been false do not give rise to knowledge, successful attempts that could have easily failed do often give rise to intentional action. In this sense, knowledge is fragile, where intentional action is robust. But even if this first premise of the argument were true, the agential perfection of trying would undermine the case for the second. For, if trying is agentially perfect, it is impossible that one tries to try to ϕ but fails to try to ϕ . So if one intentionally tries to ϕ , then for any number of repetitions of 'could have easily', it's just not true that one could have easily been such that...one could have been easily such that one tries but fails to try to ϕ .

This abstract contrast has downstream consequences too. Since the fact that one is trying to ϕ is typically non-trivial, followers of Williamson will say that one can't have arbitrarily many iterations of knowledge that one is trying to ϕ . But, given INTENTIONAL IDENTITY, it is always true that for arbitrary many iterations of 'intentionally', one intentionally...intentionally tries to ϕ . So plausibly it can't be that one intentionally ϕ s only if one knows that one ϕ s, since that would seem to require cases of the forbidden arbitrary iterations of 'know' over (non-trivial) 'try'. So INTENTIONAL IDENTITY provides a new argument for a negative answer to the much discussed question whether, if one ϕ s intentionally, one knows that one ϕ s.³⁰

7.4 Options in decision theory

Finally, we turn to some consequences of the agential perfection of trying for normative decision theory. Decision theorists typically appeal to a set of *options* as the basis for their verdicts about what a person ought (subjectively) to do. But what are these options?

A number of constraints have been proposed. Among them is the idea that options must be essentially successful—see, e.g., Pollock (2002).³¹ The intuitive motivation for this idea comes from the thought that a person's options should be things they cannot fail to do, where to fail to ϕ in the relevant sense is to try to ϕ but not ϕ . In other words: if ϕ -ing is an option, then one should be guaranteed to succeed in ϕ -ing if one tries to ϕ , i.e. ϕ -ing should be essentially successful. Supposing this thought is right, then since essential success is a very demanding condition, trying would be at least *prima facie* important for decision theory. For, if trying is agentially perfect, it is essentially successful, so tryings have already passed an important and restrictive test for being options.

More than this, though, the overall picture we've developed—on which trying is essentially successful—strengthens the idea that the property of being essentially successful is important for decision theory in the first place. To see this, suppose ϕ -ing was essentially successful but that trying to ϕ was not. Then even if one could not fail to ϕ (in the sense that whenever one tries, one succeeds), it might be that one could still fail (in this sense) to *try to* ϕ . And if one could fail to try to ϕ , then in a case where one's only means of ϕ -ing was trying to ϕ , one might try to try to ϕ , but still fail to try to ϕ , and thus not actually ϕ . Such cases wouldn't be ones where the person *tried and failed* to ϕ , but they would be cases where the person did not ϕ , in spite of trying to try to ϕ . So they would raise the question of why we should care about the fact that ϕ -ing is essentially successful. By contrast, if trying itself is essentially successful, the problem does not arise. For in this case the essentially successful doings are the things we do without fail, not just because trying to do them suffices for doing them, but also because trying to try to do them

³⁰ See, e.g., Davidson (1971), Setiya (2008), Paul (2009), Piñeros Glasscock (2019), and Kelley (2024b) (and for a survey, see Paul (2020, Ch. 6)). For responses, see, e.g., Anscombe (1957), Thompson (2011), Pavese (2020, 2022), and Beddor and Pavese (2021).

³¹ For related discussion, see Hedden (2012, 2015), Koon (2020), Schwarz (2021).

suffices for doing them, and so on. If trying is essentially successful, then that which is essentially successful is that which we can do without fail, in the fullest sense of the phrase.

8 Conclusion

We have argued that trying is agentially perfect, and drawn out several consequences of this claim. The agential perfection of trying opens a new path to a trying-theoretic conception of basic action, yields important results in the logic of intentional action, and strengthens the case for a trying-theoretic conception of options in decision theory.

But many are skeptical that the notion of trying is of any interest in its own right (cf. Ludwig (2021)). In closing, we want to respond to these skeptics with one last plea for our topic. We'll do so by developing a broad parallel between epistemology and the philosophy of action—a parallel we hope helps to illuminate the case for the study of trying on its own terms.

There is a natural analogy between belief and knowledge, on the one hand, and trying and intentional action, on the other.³² Trying is plausibly necessary for (primary) intentional action in the way that belief is plausibly necessary for knowledge. Likewise, one's environment imposes fewer constraints on what one can try to do than on what one can do intentionally, in much the same way that one's environment imposes fewer constraints on what one can believe than it does on what one can know. If one is unknowingly paralyzed, many of one's attempts will fail to manifest themselves in intentional action; if one is unknowingly envatted, many of one's beliefs will fail to constitute knowledge.

The analogy between these notions also extends to questions about conceptual priority. Is knowledge to be understood in terms of belief, or is belief to be understood in terms of knowledge (or is neither reducible to the other)? Is intentional action to be understood in terms of trying, or is trying to be understood in terms of intentional action (or neither)?

The similarities here are as important as the differences. Post-Gettier epistemology sees few who take seriously the prospects of offering a non-trivial analysis of knowledge in terms of belief (or vice-versa, for that matter). But a good deal of work in the philosophy of action takes seriously the prospects of providing such non-trivial analyses. Some take trying to be prior, and claim that intentional actions are events that are caused by tryings in the right way (Armstrong 1973), or that what makes an event an action is simply that it *is* a trying (Hornsby 1980, Pietroski 2000). Others take action to be prior, and claim that trying is doing what one takes to be

³² Williamson (2017) argues that *intention* (not trying) stands to acting as belief stands to knowledge. But intentions can fail to lead to action because the person changes their mind before it is time to act (or just because they never try), while beliefs have no analogous failure mode. Moreover, few hold that successful intentions are identical to intentional actions. But just as it is a popular view that "successful" beliefs are identical to states of knowledge, it is a popular view that some successful tryings are identical to intentional actions.

necessary to perform a certain action (O’Shaughnessy 1973), or that to try to ϕ is to be such that if one were in the right circumstances, one would succeed in ϕ -ing (Ruben 2016, 2018, 2022), or that to try to ϕ is to do something with the intention of bringing it about that one ϕ s (Ludwig 2021).

We have not entered into this debate. We have argued that trying is necessary for (primary) intentional action, and that it plays an important role in the characterization of basic action. But we’ve remained silent on the question of whether this means that intentional action itself can be analyzed in terms of trying, or vice-versa. For all we’ve said, one could endorse an “action-first” theory of trying, or, for that matter, a “trying-first”, *peirastic* theory of action. However, our own suspicion is that neither is right: the notions of action and trying are both conceptual primitives, and there is no non-trivial analysis of intentional action in terms of trying, or of trying in terms of intentional action.³³

But in epistemology, even those who think of belief as a kind of mere “defective knowledge” admit that it may be subject to rich structural laws of its own—that many substantive questions about its nature and norms are unobvious and of great theoretical interest. We have made the case for a similar view of trying. Even those who see trying as mere “defective intentional action” should accept that trying is subject to rich structural laws of its own, and that questions about its nature and norms are worthy of study in their own right.

Acknowledgements The authors contributed equally to the ideas and writing of the paper. Thanks to David Barnett, Bob Beddor, Kyle Blumberg, Mike Caie, Sam Carter, Imogen Dickie, Adam Elga, Davide Fassio, Anton Ford, Jane Friedman, Eric Funkhouser, Jessica Gelber, Simon Goldstein, Caspar Hare, John Hawthorne, Nico Kirk-Giannini, Hongwoo Kwong, Matt Mandelkern, Jack Mullaney, Jennifer Nagel, Carlotta Pavese, Antonia Peacocke, Ian Phillips, Bryan Reece, Gideon Rosen, Kieran Setiya, Michael Smith, Robert Stalnaker, Eric Swanson, Sergio Tenenbaum, and Timothy Williamson for conversations, correspondence and written comments. Friedman, Goldstein, Phillips, Rosen, and Williamson especially went above and beyond. Thanks also to audiences at Agency and Intentions in Language 3, the University of Arkansas, the ACU Workshop on Action, Intention and Language, the Centre for Human Abilities (Berlin), CUNY, JHU, UMass Amherst, MIT, Princeton, the Ranch Metaphysics Conference, UT Austin, and the University of Toronto. Thanks finally to three anonymous referees for helpful comments.

References

- Adams, F. (1986). Intention and intentional action: The simple view. *Mind and Language*, 1(4), 281–301.
- Adams, F. (1991). He doesn’t really want to try. *Analysis*, 51(2), 109–112.
- Adams, F. (1994a). Of epicycles and elegance. *Canadian Journal of Philosophy*, 24(4), 637–641.
- Adams, F. (1994b). Trying, desire, and desiring to try. *Canadian Journal of Philosophy*, 24(4), 613–626.
- Adams, F. (1995). Trying: You’ve got to believe. *Journal of Philosophical Research*, 20, 549–561.
- Adams, F. (1997). Cognitive trying. In G. Holmström-Hintikka & R. Tuomela (Eds.), *Contemporary action theory volume 1: Individual action* (pp. 287–314). Springer Netherlands.
- Adams, F. (2007). Trying with the hope. In M. Timmons, J. Greco, & A. R. Mele (Eds.), *Rationality and the good: Critical essays on the ethics and epistemology of robert audi*. Oxford University Press.

³³ Here we see ourselves as aligned in different ways with Anscombe (1957), Thompson (2008), Levy (2013), and Ford (2017).

- Adams, F., & Mele, A. R. (1992). The intention/volition debate. *Canadian Journal of Philosophy*, 22(3), 323–337.
- Albritton, R. (1985). Freedom of the will and freedom of action. *Proceedings and Addresses of the American Philosophical Association*, 59(2), 239–51.
- Amaya, S. (2017). Basic actions reloaded. *Philosophy Compass*, 12(9), e12435.
- Anscombe, G. E. M. (1957). *Intention*. Harvard University Press.
- Armstrong, D. M. (1973). Acting and trying. *Philosophical Papers*, 2(1), 1–15.
- Beddor, B., & Pavese, C. (2021). Practical knowledge without luminosity. *Mind*, 131(523), 919–936.
- Bratman, M. (1984). Two faces of intention. *Philosophical Review*, 93(3), 375–405.
- Bratman, M. (1987). *Intention, plans, and practical reason*. Harvard University Press.
- Buckwalter, W., Rose, D., & Turri, J. (2021). Impossible intentions. *American Philosophical Quarterly*, 58(4), 319–332.
- Carter, J. A., & Shepherd, J. (2022). Intentional action and knowledge-centred theories of control. *Philosophical Studies*, 180(3), 957–977.
- Chisholm, R. (1971). On the logic of intentional action. In R. W. Binkley, R. N. Bronaugh, & A. Marras (Eds.), *Agent, action, and reason* (pp. 38–80). University of Toronto Press.
- Danto, A. C. (1965). Basic actions. *American Philosophical Quarterly*, 2(2), 141–148.
- Danto, A. C. (1979). Basic actions and basic concepts. *Review of Metaphysics*, 32(3), 471–485.
- Davidson, D. (1971). Agency. In A. Marras, R. N. Bronaugh, & R. W. Binkley (Eds.), *Agent, action, and reason* (pp. 1–37). University of Toronto Press.
- Davidson, D. (1973). Freedom to act. In *Essays on Actions and Events* (pp. 63–81).
- Descartes, R. (1985–1992). *The Philosophical Writings of Descartes*. Cambridge University Press. trans. John Cottingham, Robert Stoothoff, Dugald Murdoch, and Anthony Kenny.
- Descartes, R. (1996). *Oeuvres de Descartes*. J. Vrin. ed. Charles Adam and Paul Tannery.
- Ford, A. (2017). The representation of action. *Royal Institute of Philosophy Supplement*, 80, 217–233.
- Frankfurt, H. (1969). Alternative possibilities and moral responsibility. *Journal of Philosophy*, 66(23), 829.
- Ginet, C. (1990). *On action*. Cambridge University Press.
- Goldman, A. (1970). *A theory of human action*. Princeton University Press.
- Gorr, M. (1979). Willing, trying and doing. *Australasian Journal of Philosophy*, 57(3), 237–250.
- Grano, T. (2011). Mental action and event structure in the semantics of ‘try’. *Semantics and Linguistic Theory*, 21, 426–443.
- Grano, T. (2017). Control, temporal orientation, and the cross-linguistic grammar of trying. *Glossa*, 2(1), 94.
- Grice, P. (1989). *Studies in the way of words*. Harvard University Press.
- Hampshire, S. (1959). *Thought and action*. Chatto Windus.
- Harman, G. (1976). Practical reasoning. *The Review of Metaphysics*, 29(3), 431–463.
- Harman, G. (1986). Willing and intending. In R. Grandy & R. Warner (Eds.), *Philosophical grounds of rationality: Intentions, categories* (pp. 363–380). Clarendon Press.
- Heath, P. (1971). Trying and attempting. *Aristotelian Society Supplementary*, 45(1), 193–227.
- Hedden, B. (2012). Options and the subjective ought. *Philosophical Studies*, 158(2), 343–360.
- Hedden, B. (2015). Options and diachronic tragedy. *Philosophy and Phenomenological Research*, 90(2), 423–451.
- Hedman, C. G. (1970). Intending the impossible. *Philosophy*, 45(171), 33–38.
- Hermes, C. (2006). Does attempting to try to a imply trying to a? *Southwest Philosophy Review*, 22(2), 63–70.
- Hornsby, J. (1980). *Actions*. Routledge and Kegan Paul.
- Hornsby, J. (1995). Reasons for trying. *Journal of Philosophical Research*, 20, 525–539.
- Hornsby, J. (2010). Trying to act. In T. O’Connor & C. Sandis (Eds.), *A companion to the philosophy of action. Blackwell companions to philosophy series* (pp. 18–25). Wiley.
- Hunter, J. F. M. (1987). Trying. *Philosophical Quarterly*, 37(149), 392–401.
- Jones, O. R. (1983). Trying. *Mind*, 92(367), 368–385.
- Kelley, M. (2024a). How to perform a nonbasic action. *Noûs*, 58(1), 106–125.
- Kelley, M. (2024b). Separating action and knowledge. *The Philosophical Quarterly*. <https://doi.org/10.1093/pq/pqae050>
- Knobe, J. (2003). Intentional action in folk psychology: An experimental investigation. *Philosophical Psychology*, 16(2), 309–325.
- Knobe, J. (2006). The concept of intentional action: A case study in the uses of folk psychology. *Philosophical Studies*, 130(2), 203–231.

- Koon, J. (2020). Options must be external. *Philosophical Studies*, 177(5), 1175–1189.
- Lavin, D. (2013). Must there be basic action? *Noûs*, 47(2), 273–301.
- Levy, Y. (2013). Intentional action first. *Australasian Journal of Philosophy*, 91(4), 705–718.
- Ludwig, K. (1992). Impossible doings. *Philosophical Studies*, 65(3), 257–281.
- Ludwig, K. (1995). Trying the impossible: Reply to Adams. *Journal of Philosophical Research*, 20, 563–570.
- Ludwig, K. (2016). *From individual to plural agency: Collective action: Volume 1* (Vol. 1). Oxford University Press.
- Ludwig, K. (2021). Let me go and try. *Philosophical Explorations*, 24(3), 340–358.
- Mandelkern, M., Schultheis, G., & Boylan, D. (2017). Agentive modals. *Philosophical Review*, 126(3), 301–343.
- McCann, H. (1974). Volition and basic action. *Philosophical Review*, 83(4), 451–473.
- McCann, H. (1975). Trying, paralysis, and volition. *Review of Metaphysics*, 28(3), 423–442.
- McCann, H. J. (1986). Rationality and the range of intention. *Midwest Studies in Philosophy*, 10(1), 191–211.
- McCann, H. J. (1989). Intending and planning: A reply to Mele. *Philosophical Studies*, 55(1), 107–110.
- McCann, H. J. (1991). Settled objectives and rational constraints. *American Philosophical Quarterly*, 28(1), 25–36.
- McCormick, S., & Thalberg, I. (1967). Trying. *Dialogue*, 6(1), 29–46.
- McGinn, C. (1982). *The character of mind*. Oxford University Press.
- Mele, A. R. (1990). He wants to try. *Analysis*, 50(4), 251–253.
- Mele, A. R. (1991). He wants to try again: A rejoinder. *Analysis*, 51(4), 225–228.
- Mele, A. R. (1992). *Springs of action: Understanding intentional behavior*. Oxford University Press.
- Mele, A. R. (1994). Desiring to try: Reply to Adams. *Canadian Journal of Philosophy*, 24(4), 627–636.
- O’Shaughnessy, B. (1973). Trying (as the mental “pineal gland”). *Journal of Philosophy*, 70(13), 365–386.
- O’Shaughnessy, B. (1980). *The will: A dual aspect theory*. Cambridge University Press.
- Paul, S. K. (2009). Intention, belief, and wishful thinking: Setiya on ‘practical knowledge’. *Ethics*, 119(3), 546–557.
- Paul, S. (2020). *Philosophy of action: A Contemporary introduction*. Routledge.
- Pavese, C. (2020). Probabilistic knowledge in action. *Analysis*, 80(2), 342–356.
- Pavese, C. (2022). Practical knowledge first. *Synthese*, 200(5), 1–18.
- Payton, J. D. (2021a). Attempts. *Philosophical Studies*, 179(2), 363–382.
- Payton, J. D. (2021b). *Negative actions: Events, absences, and the metaphysics of agency*. Cambridge University Press.
- Pietroski, P. M. (2000). *Causing actions*. Oxford University Press.
- Piñeros Glasscock, J. S. (2019). Practical Knowledge and Luminosity. *Mind*, 129(516), 1237–1267.
- Pollock, J. L. (2002). Rational choice and action omnipotence. *Philosophical Review*, 111(1), 1–23.
- Ruben, D.-H. (2013). Trying in some way. *Australasian Journal of Philosophy*, 91(4), 719–733.
- Ruben, D.-H. (2015). The physical action theory of trying. *Methode*, 4(6), 1–19.
- Ruben, D.-H. (2016). A conditional theory of trying. *Philosophical Studies*, 173(1), 271–287.
- Ruben, D.-H. (2018). *The metaphysics of action: Trying, doing, causing*. Palgrave Macmillan.
- Ruben, D.-H. (2022). Reply to “attempts”: A non-davidsonian account of trying sentences. *Philosophical Studies*, 179(12), 3817–3830.
- Schroeder, S. (2001). The concept of trying. *Philosophical Investigations*, 24(3), 213–227.
- Schwarz, W. (2021). Objects of choice. *Mind*, 130(517), 165–197.
- Setiya, K. (2008). Practical knowledge. *Ethics*, 118(3), 388–409.
- Shepherd, J. (2016). Conscious action/zombie action. *Noûs*, 50(2), 419–444.
- Shepherd, J., & Carter, J. A. (2021). Knowledge, practical knowledge, and intentional action. *Ergo An Open Access Journal of Philosophy*, 9, 556.
- Sneddon, A. (2001). Does philosophy of action rest on a mistake? *Metaphilosophy*, 32(5), 502–522.
- Storrs-Fox, D. (2019). *Non-Cartesian agency and normative theory*. Unpublished manuscript.
- Thalberg, I. (1962). Intending the impossible. *Australasian Journal of Philosophy*, 40(1), 49–56.
- Thompson, M. (2008). Naive action theory. In M. Thompson (Ed.), *Life and action*. Harvard University Press.
- Thompson, M. (2011). Anscombe’s intention and practical knowledge. In A. Ford, J. Hornsby, & F. Stoutland (Eds.), *Essays on Anscombe’s intention*. Harvard University Press.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford University Press.

- Williamson, T. (2017). Acting on knowledge. In J. A. Carter, E. C. Gordon, & B. W. Jarvis (Eds.), *Knowledge first: Approaches in epistemology and mind* (pp. 163–181). Oxford University Press.
- Wittgenstein, L. (1950). *Philosophical investigations*. Macmillan.
- Yaffe, G. (2004). Trying, intending, and attempted crimes. *Philosophical Topics*, 32(1/2), 505–531.
- Yaffe, G. (2010). *Attempts: In the philosophy of action and the criminal law*. Oxford University Press.

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

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.