

Attempts as Inquiry*

Sarah K. Paul

Here is a question many of us would like to know the answer to: “What could I (we) accomplish, if I (we) really tried?”

There are plenty of accomplishments that are of little interest to us, or that are relatively easy to know about. I know that if I tried right now, I could pick up my mug and drink some tea, in virtue of having done this countless times before. Not only this, but it’s the kind of thing that most everyone can do. I might have to try more than once, if I’m especially clumsy the first time. Or if there is some problem with the mug—it’s too hot to touch, or doesn’t actually contain tea—then I might have to do some maneuvering to put myself in a situation where I could successfully drink the tea. But this is also something I know from experience that I could do if I tried. I imagine I could use similar abilities to do something I have never done before—make and drink a mug full of mud, say—though I have no particular interest in doing that.

The question is of more significance when it comes to something we have never done before, and are interested in doing, but that we have reason to think will be challenging to do successfully. Consider certain kinds of long-term, difficult goals around which many of us structure our practical lives: things like career goals, athletic or artistic ambitions, long-term relationships, and political aspirations. These are often not limited to the sorts of things that most anyone can do if they try. In deciding where to devote our time and effort, we wonder: “Could I succeed at opening a popular restaurant, or publishing a novel, or completing the Ironman? Could I have a happy, well-functioning relationship with this person? Could we change US federal law by organizing around gun control?” The answer is salient because pursuing a difficult long-term goal will tend to involve effort, sacrifice, frustration, and heartbreak, as well as significant opportunity costs—things one could have done instead. One could organize around a different, state-level cause, or opt for the single life, or stick with an office job rather than becoming an entrepreneur. It therefore becomes pressing to know whether one’s efforts will ultimately pay off, or whether a particular endeavor is doomed to fail (Morton and Paul, forthcoming).

* The theme of this paper and many of the foundational ideas derive from co-authored work with Jen Morton, to whom I am perennially indebted and grateful. Thanks also to Adam Elga, Tom Kelly, Mike Titelbaum, Sara Purinton, Cristián Rettig, Wendy Salkin, Katia Vavova, and Gabby Zhang for helpful conversations, to Lara Buchak for comments, and to audiences at the 2026 NYUAD Normativity and Reasoning workshop, the Princeton University Center for Human Values, the University of Pennsylvania, Davidson College, and Mt. Holyoke.

So how do we come to know what kinds of difficult, long-term achievements would be within our reach if we really tried? We first need a better sense of what the question might be asking about. I propose that in the context of these kinds of goals, we are often interested in an underexplored sense of ability that I will call *achievement ability*. I will then argue that acquiring first-personal knowledge of this kind of ability is challenging in distinctive ways. To make significant progress on getting a well-founded answer to this question, there will often be no good substitute for simply trying to achieve the goal and seeing what happens. And on the other hand, "really trying" to succeed will involve inquiring into the conditions bearing on whether and how it will be possible to succeed. We end up with the idea that under some conditions, difficult attempts in the face of uncertainty just are, at the same time, instances of inquiry.

This means that there can be a single act-token to which both practical and theoretical standards and assessments apply. The final part of the paper explores the implications of this idea. I suggest that if there are epistemic norms on inquiry, cases of "inquisitive attempts" should incline us to think that such norms and assessments are not "pure"—they are not insulated from practical features of the context.

I. Varieties of ability

There are many different senses of ability we could be interested in when we ask questions of this form, so let us start by being more precise about what might be at issue.¹ First, theories of ability sometimes distinguish between *general* and *specific* abilities (Honoré 1964; Mele 2003). A general ability is something an agent can possess—often understood as a kind of skill or competence to reliably perform some act-type—even if she cannot exercise it in the present circumstances. A specific ability, in contrast, concerns something one is presently in a position to do because all the prerequisites have been met: "nothing stands between" the agent and the performance (Maier 2014). This is the difference between being able to swim while sitting in one's office and being able to swim while at the beach.

One might be wondering in any given case about either of these senses of ability. But when it comes to the difficult, long-term ambitions that are the focus here, I think that neither category captures what we often have in mind. On one hand, the question is often about a manifestation of agency that falls short of general ability. It would be nice to have the general ability to, say, repeatedly open profitable restaurants in various cities around the world. But when we are trying to accomplish something difficult for the first time, we will often be satisfied (for now) simply to be able to do it once, in the particular circumstances we are in. What is most practically relevant at this stage is whether we are in a position to realize this particular goal in the actual world, not whether we could reliably succeed in a

¹I hope to invoke the notion of ability here without committing to any particular view about what abilities are, or whether or how they explain performances. My real interest is primarily in true answers to a specified sense of the 'can' question, and the kinds of evidence that support such ability-ascriptions.

variety of circumstances that afforded the minimal conditions for a successful performance of that act-type.

On the other hand, it is not the case that nothing stands between the agent and the performance. Success will typically require the actual circumstances to be significantly changed over time—either through the agent’s intervention, or by other means. So although the agent is wondering whether it is *now* true that she can succeed, the question is not indexed to the circumstances that are now actual. The aspiring student of architecture does not now have the specific ability to win a major contract given that they lack a degree, a license, and a job with a prestigious firm. But their interest is in whether they will be able to bring those circumstances about, such that they will eventually have the specific ability in the future. Importantly, it will often be deeply unclear at the outset what those future circumstances will look like.

A second customary distinction concerns whether the ability is *simply* to do something, or to do it *intentionally*, or to *ensure* that it happens (Mele 2003). In the simple sense, an agent is able to do anything they in fact do—even those things they do entirely as a matter of luck, or have no idea that they’re doing. This very broad category includes slips, flukes, and accidents as well as genuine actions. Doing something intentionally is a narrower category, requiring more than simple ability, though it is a matter of controversy what more is required (see Paul 2020). Perhaps the deed must have been caused and guided by the right psychological states or mechanisms: those states and mechanisms, such as intentions, that are the source of our agency. Or perhaps the agent must exercise a certain degree of control over the outcome. And some views require that the agent have a special kind of “practical knowledge” of what she is doing, or even what she will have done.

Supposing that the requirements for being able to do something intentionally do not suffice for ensuring that one will do it, we can further speak of ensuring-level ability. For instance, I might have sufficient control to be able to give a good lecture intentionally—it would not be a fluke, if I did give a good lecture—but not enough control to ensure that it is good and sincerely promise you that it will be. If that gap exists (which some views of intentional action might deny), then the ability to ensure success will require even more than the ability to do something intentionally.

But again, though any of these notions might be at issue in a given instance, our interest in achieving difficult ambitions often does not fall cleanly into any of these categories. Normally, a mere fluke or lucky accident will not count as succeeding in achieving the accomplishment in question. The aspiring architect would not take herself to have accomplished her goal if she gets hired by a top firm only because of an identity mix-up. This would completely undermine the meaning of the “achievement.” I do not mean to deny the possibility of being interested primarily in outcomes rather than accomplishments or achievements, and orienting one’s practical reasoning toward positioning oneself to cause desired outcomes in fluky or accidental ways. We can aspire to win the lottery! But many of us do care about a stronger sense of doing something in which the outcome can be credited

to our agency. And for the purposes of deliberation, it will often be most relevant to focus on those things we have at least some control over.

On the other end of the spectrum, ensuring-level ability is generally not needed to make wise decisions about how to direct one's efforts (though again, it would be nice). In most such cases, we have not promised anyone else that we will succeed, and some amount of risk is part of the conception of what we are doing. Indeed, the difficulty and risk are often part of the attraction of certain ambitions. Being in a position to ensure success implies that the achievement is entirely up to you, and that's usually not the case in the examples we have in mind.

We might think the ability to succeed intentionally is the needed middle ground. But this may not be quite right either. It will depend, of course, on how we fill in the details of our theory of intentional action. Plausibly, however, the kinds of accomplishments we are interested in here are more compatible with something we might call "luck" playing a significant role than intentional actions are often thought to be. To succeed at the kinds of difficult long-term goals in question, many things will have to go right that are not directly in our control. We will need help from other people, relatively favorable conditions, and other unpredictable instances of good fortune. Opening one's first successful restaurant, publishing one's first novel, or winning one's first grass-roots campaign will often not be modally robust occurrences—it will not be the case that the agent would have succeeded in most nearby possible worlds. This might block the outcome from counting as a fully intentional action, but it does not necessarily prevent us from attributing it to the agent as her accomplishment if she does succeed. As long as the success is sufficiently explained by the agent's (or group's) planning, practical reasoning, and perseverance, they can legitimately claim to have done what they set out to do. But along the way, they will frequently lack practical knowledge of what they are doing under more ambitious descriptions, and it is not clear that they will have enough control to count as doing it intentionally all along (Morton and Paul, forthcoming).

Finally, a point famously made about the so-called "conditional analysis" of ability, in which an ability is understood in terms of what one could do if one tried, is that one might not be able to try (Lehrer 1968). That is, a talented student might have the ability to earn an advanced degree at a prestigious university in the US, but be unable to try because she cannot get a visa. Applied to the present context, the point helps us see that the agent who is striving to accomplish a difficult long-term goal might be interested not only in the question of whether they could succeed if they tried, but also in the question of whether they will be able to try. This is a slightly different question, however, and one I will largely set aside here.

It will be useful to have a label for the distinctive sense of ability I have pointed to here. Let us refer to it as *achievement ability*. Achievement ability is distinguished by the type of specificity involved, which includes the agent's actual circumstances considered over time rather than at a time, and with respect to its compatibility with relatively greater amounts of assistance and luck than more control-based, robust senses of ability. It is not the only sense

of ability we care about, but it is especially relevant to practical reasoning about whether to take up a difficult end and persevere in its pursuit.

II. The epistemology of achievement ability

Our focus is now on the specification of our initial question that is getting at achievement ability. We have gained a better grip on the truth-conditions governing such ascriptions. The next question is: how can we come to know whether we could accomplish something difficult, in the specified sense, if we tried?

I suggest that for many of the long-term, difficult ends that might appeal to us, there is no good substitute for simply attempting to do the thing. Of course, there will be some such goals we can rule out for ourselves on the basis of information that is available without the need to try them out. It might be clear that one is too tall to be an astronaut, or too blind to be a pilot, or too hopeless with languages to join the CIA, without needing to apply in order to find out. But in many other cases, the potential obstacles to succeeding will be less definitive or fixed. The available evidence will not exclude the possibility that one's current skills could massively improve, or that one's circumstances could be changed to become far more conducive to success than they currently are.

Even in these cases, we can get some insight into our ability to succeed through sources other than making the attempt (Vetter 2024). Most such accomplishments will depend on many other, more limited abilities that we know about from past experience. And we can use counterfactual reasoning and imagination to extend our existing ability knowledge to new situations. "I can consistently finish a ten-mile training run," one might think, "so if it were race day and I had more adrenaline, I could complete the full half-marathon at a similar pace." We can also draw on the testimony of others, as well as the evidence offered by the examples of those whom we take to be relevantly similar to ourselves (Morton 2022; Kikkert forthcoming). Perhaps many others who have followed this half-marathon training plan have succeeded in completing the race, leading me to think that I too will be able to succeed. Plausibly, in addition to past experience and certain kinds of reasoning, perception can also inform us of affordances in our environment. This can help us apply our knowledge of our general abilities to the specific circumstances.

All of this can contribute to informing an estimate of what would happen if we tried to accomplish a new, difficult, long-term goal. It will only get us so far, however. These sources of evidence will not bear decisively, at least in many cases, on first-personal questions about *me* (us) and *my* (our) actual circumstances.² On one hand, if the accomplishment is difficult, the evidence provided by the examples of others will often suggest that you will not succeed: many others who have attempted them have failed or

² This emphasis on knowledge with a first-personal character has resonances with Laurie Paul's work on the epistemology of transformative experience (2014). Some of my cases might involve this dimension, but the main points I make here apply even when nothing especially transformative or experiential is at issue.

given up (Marušić 2015). The evidential situation may be particularly hostile in conditions of prejudice, discrimination, or other forms of structural injustice operating against members of a social group with which the agent identifies. If few people in the poor community one grew up in have succeeded in escaping poverty, or if there are few examples of those who share one's race or gender accomplishing a similar goal, then reasoning from similarity will suggest that you too will be unable to succeed (this is what Jennifer Morton 2022 has termed a "pessimism trap").

Such evidence can make it rational to think it is too risky to even try for a difficult goal. That said, as is well known, the inference from statistics to the specific case is not straightforward: you might be the exception. And if the goal is important enough to you, then it might be reasonable to take seriously the possibility that you are the exception. In that case, specific evidence will be needed that you are *not* the exception before the first-personal question can be answered of whether you can achieve it.

Relatedly, there will generally be substantial uncertainty about the various parameters we need to fill in for counterfactual imagination and reasoning to tell us much about the actual world, especially in the far future. Success in these kinds of efforts will depend heavily on the specifics of the circumstances as they unfold over a significant period of time, and the relatively unpredictable judgments of others about whether some result ultimately counts as a success (whether a novel is well received, say). The amount of control one has to change the circumstances in a more favorable direction, or to ensure that others will have the sought-after reactions, is itself in question. And the agent will often be uncertain about where exactly the adversity will come from. Will it be due to the permanent lack of more fine-grained abilities that would be needed to, say, prosper as an opera singer? Or can she gain the needed abilities over time, with discipline and practice? Or are the obstacles primarily external to her—is she the target of discrimination or bias, or otherwise operating in circumstances that are hostile to her efforts in ways that she has little control over? She can imagine all sorts of possibilities, but since the kind of ability she is wondering about is so tightly tied to the circumstances and events unfolding in the actual world, her counterfactual thinking will be of limited use.

So while we can arrive at estimates of our chances of success before making any kind of attempt, informed by various sources of evidence, reasoning, imagination, and perception, our epistemic access is significantly constrained. Practically speaking, to improve our epistemic position with respect to the question of whether we could accomplish a difficult long-term goal if we tried, the best way forward may be to try to do it and seeing what happens.

This does not mean that the attempt will always or easily yield knowledge. Some conditions may be hostile to learning much from the attempt—if, for instance, one is not in a position to get useful feedback from others. Even if the attempt does succeed, the evidence may be undecisive as to whether too much luck was involved in order to count as an achievement of one's own. And frequently, the value of knowing the answer to the question will be outweighed by the disvalue of continuing too far in the pursuit of a risky endeavor

and paying the associated costs. All things considered, one should often quit before one learns definitively whether succeeding was possible in the relevant sense.

Further, insofar as the success conditions of the accomplishment are specified on a vague time-horizon, it will be difficult to get definitive evidence about what could happen in the far-out future if one kept trying. It will be relatively easy to find out whether you can succeed at an endeavor with a very specific deadline—either you made the Forbes “30 Under 30” list by your 30th birthday or you didn’t—but harder to find out whether you could achieve success as an entrepreneur at some unspecified future point. With all of this in mind, perhaps a better way to think of the theoretical goal here is in terms of improving one’s epistemic position with respect to the question. An attempt will often leave us short of knowledge, but it will sometimes be the best way of getting useful information that is sufficiently specific to the agent and her circumstances.

III. Attempts as inquiry

We do not always go into an attempt with this kind of theoretical goal or interest. But when we do, the upshot is that the attempt *just is*, at the same time, a kind of inquiry.³ We are wondering about something, and the attempt is the way to figure it out.⁴

We can drive at the same conclusion—that many difficult attempts are also inquiries—by thinking about the nature of action in the face of uncertainty. In order to succeed at the kinds of challenging ends at issue, in a way that is not simply a fluke or an accident, the realization of the goal needs to be sufficiently explained by the agent’s practical reasoning. She needs to figure out a way to realize her end, and her success must be partly attributable to the fact that she intentionally undertakes a set of actions conceived of as a means to succeeding at that end. As acknowledged earlier, the kind of agency at issue here is compatible with exercising less than full control over the outcome. The means chosen by the agent may not fully determine a successful result. But they do need to make a sufficient contribution, in the way that the agent thinks they might, for the result to count as an agential accomplishment.

It is typical of such cases, however, that the agent does not know in advance what a sufficient set of means might be to realize her end. She will generally begin with a schematic

³ I am minimally assuming here that inquiry must be guided by some theoretical interest—embodied in an interrogative attitude, perhaps (J. Friedman 2024) or in a question-directed intention (D. Friedman 2025)—but not that it must have some further constitutive aim beyond improving one’s epistemic position (Falbo 2023). Some accounts of inquiry may be more restrictive in the further conditions one has to meet in order to count as inquiring beyond being motivated and guided by some inquisitive attitude.

⁴ It is worth noting the resonances here with J. David Velleman’s view of rational agency, which also claims that agency is driven in part by theoretical aims. There are key differences, however. For Velleman, knowledge is acquired by first forming a belief about what you will do, in the expectation of a reliable connection between forming that belief and its coming true, and subsequently acting in the expected way. His account does not straightforwardly apply to exercises of agency in pursuit of goals one is uncertain one can achieve, such that antecedently forming the expectation of achieving them would be unwarranted.

idea of how it works to open a new restaurant, publish a novel, run a marathon, get a law degree and eventually make partner in a firm, and so forth. But this does not amount to knowledge of a specific series of actions that will suffice to make it happen. In order to arrive at these more specific plans over time, she will need a clearer understanding of the contours of her control within the circumstances as they unfold. To meet the demands of instrumental rationality, then, she must take steps to improve her epistemic situation. This can involve forms of inquiry other than making the attempt itself. It will be helpful to get testimony from others who have gone through a similar process, or who know more about the circumstances. But the best way to fully answer certain first-personal questions about the options she will have available, and how those might add up to a pathway to realizing the goal, will be to try out various ways of proceeding and see what happens. The inquiry is, at the same time, part and parcel of the attempt to succeed (Morton and Paul, forthcoming).

Let us call such actions “inquisitive attempts.” One might be inclined to balk at the metaphysics of claiming that the attempt to succeed at something difficult, in a context of uncertainty about the factors bearing on whether one can succeed, might at the same time be an instance of inquiry. Why not say instead that there are two things going on more or less concurrently—an attempt, which is practical in nature, and an inquiry, which is theoretical in nature? There is a respect in which this is the right thing to say. Inquiring is generally a type of doing: we gather information, think through it seeking insight and understanding, and otherwise act in order to improve our epistemic situation. And it is a familiar insight that our actions can come under a number of true descriptions, and may even be intentional under more than description. It follows that one and the same event can be an action, and even an intentional action, under multiple descriptions: for example, “trying to climb El Capitan” and “figuring out whether I can climb El Capitan, if I try” (Davidson 1963; Anscombe 1979). This is what is going on in the case of inquisitive attempts: there is a single act-token, which can be truly described both as “attempting to ϕ ” and “figuring out whether I can ϕ , if I try.”

These two descriptions impose different success conditions, such that one could succeed at one and fail at the other. And it is possible to abandon one aim—give up on inquiring, say—while continuing with the other (perhaps as at some point, you’re satisfied that you can or can’t do it, and just need to finish the attempt). In that sense, the inquiry and the attempt are two different things. But in acknowledging this, we should keep in view that as long as the agent does have both the practical and theoretical aims, directed at the ultimate end, there is a single action occurring that we are describing and assessing in different ways. Compare the simpler case of walking to the store both in order to get some exercise and to buy peanut butter. There is only one event of walking that comes under both of these descriptions, and that can be assessed in light of both aims. One might succeed in getting exercise while fail in getting peanut butter, or one might remember halfway that one already has peanut butter and still keep on going in order to get the

exercise. For the period in which both aims govern, however, they are each governing the same token occurrence of walking.

Inquisitive attempts have this structure as well, though they may be far more temporally extended and therefore far more “gappy”—there can be substantial interludes in which one is doing nothing to further the relevant goal(s). Such gaps are consistent with thinking of the act-token as a single, unified thing in virtue of being goal-directed (Thompson 2008). Some parts of the attempt might involve elements that strike us as purely inquisitive, as when the agent sits down and takes stock of everything she has learned so far. And other parts may look purely “effortful,” as it were. It does not follow that these aspects must be hived off from one another and considered parts of different overarching act-tokens. As observed above, figuring out how things are going and what one is up against *is a means* to achieving the practical goal, while trying to achieve the goal *is a way of inquiring* into the question of whether one can do it. These are not two different things one is doing at the same time, or even through the same motions; each activity is a constitutive part of the other. Sometimes the right advice will be of the form “Stop wondering whether you can succeed (at writing a dissertation, say) and just get it done!” My claim is that this is not equivalent to saying, “Stop inquiring and just attempt!” Rather, the inquisitive agent should hear the advice as saying, “Stop inquiring *like that*: the way to find out, as well as to succeed, is simply to work on the dissertation.”

So although we can meaningfully distinguish inquiry as an act-type from other types of actions like “writing a dissertation,” the claim here is that these different act-types can be instantiated by one and the same act-token. Inquisitive attempts are something the agent does, both in order to achieve the practical goal and in order to find out the answer to a question.

IV. Practical and theoretical, entangled

We might now wonder why the existence of difficult attempts that are also inquiries is especially significant. What do we learn by thinking about these cases in particular?

The first thing to say is that the typical examples we focus on in both philosophy of action and epistemology tend not to be hybrid in the way inquisitive attempts are. In the philosophy of action, the paradigm cases we theorize about tend to be things we already know how to do. We think of agency as tightly connected with the exercise of control over bringing about some outcome, and thus with knowing what we are doing when we act intentionally. And in the recent surge of interest in theorizing about inquiry, the focus tends to be on short-term rather than temporally-extended inquiries, in isolation from other practical things we might be doing at the same time. Moreover, the example questions into

which we are inquiring tend to be ones in which the answer does not depend on what we ourselves do.⁵

If the aim is to understand what is distinctive of action as such, or of inquiry as such, then it makes sense to focus on the pure, isolated cases of each. This lends support to the conviction that the domain of “the practical” can be cleanly disentangled from the domain of “the theoretical,” at least in theory. But if we place cases like inquisitive attempts at the center of our theorizing, a clean separation between practical and theoretical might not seem as obvious. And these are not fringe cases with an artificially contrived puzzle at the center. Striving to do difficult things, and to find out whether we can do them, is central and important to our lives.

One way to draw out the entanglement of practical and theoretical in these cases is in terms of the common (though not uncontested) distinction between the different directions of fit an attitude or endeavor can be said to have (Anscombe 2000; Searle 1983; see Frost 2014 for critique). This distinction is meant to help us get a grip on a deep and important difference between two kinds of thought and the normative standards appropriate to each. Roughly and metaphorically, theoretical thought aims at reflecting the world as it is, while practical thought aims at getting things done and thereby altering the world. But which direction of fit do inquisitive attempts have—is the agent in these cases trying to figure something out, or is she trying to act on and alter the world? Yes. It is not only that her activity has both directions of fit, but the two are intertwined with one another. Her inquiry will help determine the fact of the matter about what she ends up doing, and what she ends up doing will help determine the answer to the question she is pursuing.

One implication is that the inquiring agent’s attempt should be shaped by her epistemic as well as practical goals. Given that she wants to acquire relevant information about her current abilities, circumstances, affective proclivities, behavioral dispositions, and other factors that will bear on whether she can succeed, there will be instrumental pressure in favor of selecting means that will be fruitful for gathering evidence. Which parts of a large project she focuses on first, for instance, might be selected partly with such considerations in mind. It may even be rational to choose a means that she thinks is unlikely to be causally effective in furthering the practical goal, but that is highly likely to yield useful information. This fact will go completely unnoticed if we think of the quality of a means as measured entirely by its causal effectiveness.⁶ Or in the context of collective action, the commitment and motivation of other group members might be a particularly potent source of doubt, giving the group reasons to take actions that will not only firm up the group’s *esprit*

⁵ In fact, Gabby Zhang (forthcoming) argues that intentionally causing an answer to be true so as to know the answer to the question is *not a way* of inquiring. However, the cases she takes to support this principle all involve short-term exercises of agency with complete control over the outcome—locking someone in his office in order to know that he is there, for instance. My temporally-extended cases in which the agent lacks complete control, or has insufficient reason to believe that he has it, challenge the general applicability of Zhang’s principle.

⁶ As Jason Bridges pointed out in a symposium on “Doing Difficult Things” at the 2026 Central APA.

de corps—thereby promoting the likelihood of achieving the goal—but that will thereby increase each member’s confidence in the commitment of the others. In general, the practical reasons bearing on how best to go about achieving the goal might point in one direction, while the reasons bearing on how to arrive more quickly at an answer to the question “will I be able to do this?” might point in another. It may make more sense to focus on the hardest part first, or to proceed in a way that will help to identify the biggest obstacles one will encounter, even if these approaches are more costly, uncomfortable, or otherwise practically worse.

More controversially, the practical interests at stake might favor giving the agent’s inquiry a shape that it otherwise would not have, in ways that depart from how we ordinarily think of the norms of good inquiry. For one thing, the agent will usually be highly motivated to reach a positive conclusion—to give herself reason to believe she can succeed. It does not follow that she must be uninterested in reaching a negative conclusion, if it is true; we often just want to know either way. But in an inquisitive attempt, the agent is actively working toward success at achieving the goal, which would entail a positive answer to the question they are inquiring into. And the reasons for which they want to succeed will largely be practical in nature; they are in it for the rewards, whether those be intrinsic to succeeding or the extrinsic consequences of success. It follows that such agents are far from disinterested, as ideal inquirers are sometimes held to be.⁷

This does not mean that we have good practical reason to conduct such inquiries in ways that will significantly distort their accuracy. As a descriptive matter, it is no doubt common to engage in highly biased, motivated reasoning that involves seeking out only evidence that confirms our desired conclusion, ignoring evidence that is contrary to it, and subjecting the contrary evidence we have to more skeptical scrutiny than our confirming evidence. But inquiry that is significantly biased in this way will tend to be a hindrance to the agent in achieving her practical goals. If she is to succeed, she needs an accurate understanding of her own limitations and the challenges she will face, and so should not avoid evidence that suggests she may fail. The pressure toward accuracy in inquiry is further supported by the agent’s practical interest in giving up on the attempt if it is unlikely to succeed, rather than continue to devote resources to a failing endeavor and pay potentially high opportunity costs. If she is likely to fail, she wants to be able to recognize that this is so. In this respect, the practical reasons at stake will not tend to favor inquiry that is significantly biased toward gathering or contemplating only favorable information.

That said, the agent does have an interest in conducting her inquiry, and sustaining provisional attitudes about how it will turn out, in ways that will support rather than detract from her efforts to succeed. After all, her assessments at any given point of whether and how she will be able to succeed will interact with her exertion of effort, her motivation, and

⁷ As is encoded in the Mertonian norms of science, for instance. See Kelly (2022), chapter six for a very useful discussion of norms of objectivity, including a defense of the Socratic norm of “following the argument where it leads” (though Kelly does not take this norm to be overriding).

her decision-making about whether to persevere or give up in the face of difficulty.⁸ As Morton and I have argued elsewhere, one common response to encountering difficulty is to despair that you will not be able to succeed even if you continue to try (Morton and Paul 2019; forthcoming). The agent might erroneously think she has sufficient evidence to conclude that she can't do it, even if things would improve if she kept going. And such despair will be self-fulfilling, leading her to give up. This means that maintaining a sufficiently optimistic assessment of one's chances of success can support perseverance and thereby contribute to one's ability to succeed. Again, wild optimism that ignores important sources of evidence will not generally be beneficial overall; it may support perseverance, but at the cost of our effectiveness in figuring out how to succeed, and of being able to recognize when the endeavor is a lost cause. But a more limited version of "positive thinking" may well be beneficial in supporting agency, and it is possible for our inquiry to be structured so as to guard against premature losses of confidence. Alternatively, some agents may actually be motivated by difficulty: their interest in succeeding is piqued by the thought that the task will be very challenging. This kind of agent will have an interest in avoiding the premature conclusion that the task is too easy.

Either way, one's interest in sustaining motivating and confidence may factor into the question raised earlier, of when to engage in those parts of a project that will give us certain kinds of information. The agent might want to take care to structure his inquiry so that he learns things in a particular order, or at the right time, so as to avoid the inclination to quit prematurely. Getting unfavorable information too soon, or too much at once, could undermine her will to continue and thereby frustrate the attempt. In order to succeed qua attempt, then, the agent might need to delay getting new information that would otherwise be relevant to the endeavor qua inquiry. Certain kinds of critical feedback might be too devastating for the graduate student just setting out to write a dissertation, in the face of significant doubt that he can pull off this daunting task. He might therefore intentionally avoid showing his early drafts to the faculty members whose practice it is to immediately point out all the powerful objections to a view—even if that information would be very valuable to have, both for planning and for prediction. Gathering favorable evidence first and avoiding negative evidence will help sustain this kind of agent's pursuit when he later gets the latter kind of evidence he also needs, both in order to ensure that his inquiry is truth-conducive and in order to know how to improve.

⁸ Lara Buchak points out to me that some of these features may be exhibited even in "pure" cases of inquiry, in which the answer does not depend on the agent, in virtue of the agency involved in trying to find out the answer: the agent may be uncertain about whether she'll be able to find out the answer, and may have reason to inquire in ways that support perseverance or otherwise improve my chances of succeeding in getting an answer. Still, it is distinctive of inquisitive attempts (and maybe other hybrid cases) that the practical reasons derive not merely from an interest in succeeding in inquiry one way or the other, but from a desire to reach a *particular* answer. And that answer will depend in part on how the agent is thinking about her chances, which will in turn depend on how her inquiry unfolds, in combination with the contextual and psychological factors that may incline her to overreact to the evidence she has (by quitting too soon, for instance).

But this looks a lot like confirmation bias. Moreover, the agent will thereby end up with a different body of evidence than he would have had if his inquiry had led him to encounter the information in a different order. As Tom Kelly has pointed out, the evidence acquired at an early stage of inquiry can causally shape the evidence one later comes to possess, insofar as it leads one to consider certain epistemic possibilities as relevant and fruitful while others are given relatively less attention (2008, 628). In a way, this is the point for the striving agent; he wants to be in an evidential position to give less attention to the demoralizing possibility that he will fail, and more attention to fleshing out the possibilities in which he succeeds. Once he is in this position, however, he has good theoretical reasons to continue to focus his inquiry on gathering information relevant to those possibilities, and may as a result be epistemically justified in believing different things, or having different credences, than a counterpart who inquired into the question purely out of curiosity.

Further, the inquiring agent will tend to be more successful, *qua* agent, if she incorporates and updates on new information in ways that support the stability of her motivation—often by sustaining optimism.⁹ One’s ability to persevere at difficult tasks can be enhanced by treating certain assumptions or models of the situation as the default in inquiry, even if this conflicts with a completely open-minded method of following the evidence where it leads. For example, thinking of ourselves as efficacious rather than helpless, until demonstrated otherwise—as being in a position to control or change many relevant features of our environment—can be instrumental in sustaining agency (Bandura 1997; Seligman 2011; Bortolotti 2018). And we might start out by treating the possibility that we will eventually figure out how to succeed as the default, to be proven otherwise, rather than treating failure as the possibility to be disproven. As long as it is possible over time to correct course and end up with a balanced and representative body of evidence, these practices need not signify that the agent is engaged in something other than truth-directed inquiry. Like Augustine and chastity, the inquisitive agent wants the truth, but not yet.

These may appear to be merely practical reasons to avoid information, in the same way that our desire to enjoy a movie gives us reason not to want to find out the ending before we finish watching it. But I think it is not that simple, since—I want to suggest—engaging in somewhat biased, motivated, overly positive inquiry can in many cases be needed in order to figure out the answer to the question “can I do this, if I try?” Imagine a completely disinterested counterpart to our invested agent, led entirely by the evidence without concern for the impact this evidence and her resulting beliefs will have on her motivation to persevere. There is a sense in which this person, I submit, is *not really trying to succeed*. Insofar as her chances of succeeding would be enhanced by engaging in inquiry that

⁹ Morton and Paul (2019 and forthcoming) argue that this might involve being relatively slower to revise one’s assessment of one’s chances downward in response to discouraging feedback, by employing a higher evidential threshold for belief change in this context than a neutral observer might. The argument here regarding inquiry is meant to be complementary to that picture of “epistemic resilience.”

is structured so as to support optimism about reaching the desired conclusion, then omitting to do so would be a matter of not doing all one reasonably can to succeed.

Now, there are multiple legitimate readings of the notion of trying. In some contexts, an agent counts as having tried to A simply in virtue of willing herself to A, or performing some basic action in order to A, even if no progress is made toward achieving A and she immediately gives up. “I tried,” the child might say after making a half-hearted gesture at cleaning up an enormous mess they have made, and there is a sense in which this is true. But in another sense more closely tied to accountability, the child hasn’t really tried. It is difficult to spell out what else must be involved, beyond minimal trying, in order to count as having tried in the more robust sense. It is sometimes glossed as “doing what one can/what is in one’s power,” or “doing one’s best under the circumstances” (Hornsby 2010; Tenenbaum 2020). These glosses cannot be taken literally, however, on pain of requiring that one commit all sorts of unsavory but available acts in order to count as having tried. There are lots of things we could do in order to further our goals that we ought not to do—lie, steal, cheat, kill—and such options do not count against us having genuinely tried.

This is a vexing issue, but we can largely sidestep it for present purposes. Whatever is involved in positive thinking is unlikely to fall into the category of morally problematic or deeply imprudent means that must be excluded by the requirements for having really tried.¹⁰ Or rather, positive thinking that is deeply imprudent, in virtue of being delusional, will be ruled out—but I take it that some milder forms will escape these constraints. There may be some practical costs or downsides to engaging in positive thinking, but this is not a decisive objection as such; nearly all means we might take to our ends have some practical cost, if only in virtue of the agential resources they absorb that could have been used in a different way.

That said, the costs do matter insofar as engaging in motivated, biased inquiry is not strictly a necessary means to succeed, or even to sustain sufficient optimism about success to keep going. It may well be possible to maintain optimism in a way that floats free to some extent from one’s take on the evidence at any given point, especially if the evidence permits a range of doxastic responses (Morton and Paul 2019). Or one might be so fortunate as to have sufficient evidence to sustain optimism about success without engaging in non-neutral inquiry. The claim is not that that inquiry *must* be motivated, as a matter of necessity, in order to give oneself the best shot of prevailing. But as long as the agent’s assessment of the feasibility of the enterprise is in fact sensitive to and constrained by the evidence she possesses—as it should be—then the way her inquiry unfolds will have a (potentially significant) effect on such assessments even if they are already as optimistic as the evidence permits. And we cannot generally count on being fortunate in the evidence we happen to get. So, trying one’s best will often, though not always, include non-neutral inquiry.¹¹

¹⁰ Unless, I suppose, one thinks we have a compelling moral duty to disinterestedly follow the evidence where it leads at each point in time.

¹¹ We might worry that “doing what one can” or “doing one’s best” will not include alterations in our doxastic states, since we lack the required control—it may not be clear that we *can* have optimistic beliefs

All of this is contingent on the specific features of a given case; I do not mean to be defending a perfectly general claim that has no counterexamples. But the suggestion is that in many cases, the structure of our inquiry will matter to giving ourselves the best chance of succeeding. And so, to count as having really tried—as having done the permissible things that were in our power to do—we must inquire in a way that is driven by factors other than simply following all of the available evidence where it leads. It does not follow that we always or even mostly have sufficient reason to do this, all things considered. We may have good reason not to go as far as we can in attempting something difficult, perhaps because the costs of doing so are higher than we are willing to pay. The conclusion is rather that *if* we want to know what we could do if we really tried, and can only find out by way of really trying, then there will often be a doxastic component involved in what it is to try. In order to give the attempt all that we reasonably can, we must think as well as act accordingly.

If that's right, then the reasons favoring this kind of positive thinking are not merely practical; they are reasons of inquiry. Jane Friedman has articulated and defended what she calls the Zetetic Instrumental Principle: "If one wants to figure out Q ?, then one ought to take the necessary means to figuring out Q ?" (where Q ? designates a question) (2020). This is not the most plausible formulation of an instrumental norm governing inquiry, as others have pointed out; merely wanting to find out whether Q cannot bootstrap us into an obligation to take immoral or highly imprudent means to doing so. But we can set these details aside here. If something in the neighborhood of ZIP is right, then when we input questions that in some way call for a kind of inquiry that normally looks improper or defective, we get reasons *of inquiry* to inquire in that way.

V. *Implications for normativity*

Does this tell us anything interesting about the relationship between practical and epistemic normativity? There are a number of views about epistemic normativity that will see no clash here. One fairly traditional view, which we may call *doxasticism* (following Friedman 2025), takes epistemic norms and assessments to apply only to states like belief, credence, and the suspension of judgment and not to activities like gathering evidence. The norms of inquiry should all be understood as instances of prudential and moral norms—the norms that apply to doing things, not those that apply to believing things. For this kind of view, it comes as no surprise that our practical interests can shape how we ought to inquire.

about our own efficacy just because they would be useful. But what is needed here is not the ability to voluntarily will a belief directly into being, on the basis of the practical reasons that favor it. As long as there are aspects of our cognitive lives that are subject to our control like patterns of attention, habits of thought, ways of framing the challenges one faces and the evidence bearing on whether we can overcome them, "manifesting" the confidence of a successful future self, and so forth, this will suffice. The agent merely needs to be in a position to cultivate a positive outlook over time, and this is something we seem to be able to do.

Two other approaches seek to reduce one domain of normativity to the other. According to *cognitivism*, practical reasons for action (or intention, if there are such things) can be reduced to reasons for belief (Harman 1976; Velleman 1989). And according to *epistemic instrumentalism*, reasons for belief can be reduced to instrumental reasons to promote our goals and interests, epistemic and otherwise. Either way, there is only one kind of normativity here and not two.

Finally, there is the fairly radical view that epistemology should be identified with the study of inquiry (*zeteticism*, again following Friedman). This kind of view will need something to say about inquisitive attempts and other cases in which an instance of inquiry is also an instance of pursuing some other practical goal. But again, I think it will come as no surprise that such an expansive view of epistemology will end up including much of what we might have thought was practical normativity.

Cases like inquisitive attempts are potentially significant, however, for views of epistemic normativity that are either less narrow or less radically expansive. In particular, I think they put pressure on a combination of views that one might want to hold: (1) that the domain of epistemic normativity is distinct from practical normativity, and does not encompass all aspects of inquiry; (2) that there are nevertheless some norms of inquiry that should be understood as epistemic norms; and (3) that epistemic norms are categorical and insulated from contingent practical interests; they cannot be “encroached upon” by practical considerations.

I assume that many of us find (1) appealing. And the kind of “purism” expressed by (3) is also widely held. Many philosophers endorse something like the idea that only alethic factors can legitimately bear on epistemic forms of assessment, and that epistemic reasons are categorical rather than contingent on the thinker’s practical goals and interests.

In favor of (2), many of us have the intuition that it is bizarre for epistemology be silent on the topic of how we should go about figuring out what to believe. And more substantively, as Carolina Flores and Elise Woodard have argued, certain evidence-gathering practices do seem legitimately subject to distinctively epistemic forms of criticism, and to demand epistemic forms of repair (2023). This suggests that the norms that are violated in these cases are themselves epistemic. One of their main examples concerns precisely the kinds of motivated evidence-gathering at issue here, in which wanted evidence is selected, or even manufactured, while unwanted evidence is avoided (see also Avnir & Scott-Kakures 2015; Hughes 2023). Another example concerns the practice of subjecting undesirable hypotheses to more skeptical scrutiny, while more easily accepting evidence that confirms our favored hypothesis.

Flores and Woodard argue that such violations lead us to reduce epistemic trust in a person, or otherwise downgrade our relationship with them along epistemic dimensions, when they lack excuse (2023, section 3.2). Other things equal, we are less inclined to rely on them epistemically, in a way that is different from the kind of damage done by the violation of a prudential or moral norm. And there seems to be nothing illegitimate about such practices in general. Not all defects in inquiry will end up being epistemic on Flores and

Woodard's view. Some norms governing inquiry may be merely instrumental norms on how to act with an eye to achieving our epistemic goals, such as eating sandwiches when we are too hungry to think straight. To disentangle the one sort from the other, we look to the practices of reaction, accountability, and criticism that seem to be appropriate.

These ideas are *prima facie* compelling. Entirely ceding the territory of evidence-gathering to the sphere of practical and moral normativity fails to be responsive to the forms of distinctively epistemic criticism and repair that seem appropriate there. And evidence-gathering processes that are motivated toward reaching a desired conclusion and selectively biased are especially good candidates for violations of epistemic norms. If we take this idea on board, however, what should we make of inquisitive attempts? What these cases help us see is that sometimes, the practical considerations at work make it the case that the best way to inquire involves practices that would normally be epistemically criticizable. The question is whether and how such practical considerations alter such epistemic assessments.

One possibility is to hold that the inquisitive agent *is* epistemically criticizable for evidence-gathering that is motivated and biased, in just the way he would be if he were not at the same time attempting to achieve the goal—though perhaps he has sufficient *practical* (and maybe all-things-considered) justification for proceeding as he does. There will certainly be numerous cases in which this is intuitively the right verdict. Imagine a person who is inquiring into whether *P*, but if it turns out that not-*P*, his heart will break. He may have good reasons deriving from the need to protect his own mental and emotional well-being to inquire in a biased, motivated way that will support the hypothesis that *P* as long as possible. In that kind of example, we can separate out the practical pressures from the epistemic ones by imagining a counterpart who is inquiring into the same question without those same practical interests. The counterpart's inquiry should be neutral, balanced, and open-minded, simply following the evidence where it leads. The emotionally invested agent has a good practical excuse for not investigating in this way. But it does not serve as an epistemic justification; we would still downgrade his credibility in this domain.

Things are different in the case of inquisitive attempts, however: there *is* no relevant counterpart who could be conducting the very same inquiry without similar practical interests. True, we could imagine a quite different version of the agent whose motivations did not depend at all on how he thinks his endeavor will turn out at a given moment. This is not impossible, but it is too dissimilar to the ordinary case of agents like ourselves to bear on how we should be assessed. For agents like us, I have argued that engaging in what would ordinarily be criticizable inquiry can be a means to succeeding, and therefore part of really trying to succeed, and therefore necessary in order to answer the question at hand. We cannot get a grip on an epistemically more ideal counterpart, engaged in the same first-personal action-dependent inquiry, in comparison with which the agent is falling short.

This is a reason to think that the inquisitive agent is justified in proceeding as she does, in a way that deflects criticism for what would ordinarily be defective evidence-gathering practices. And this justification is partly grounded in practical considerations; it is because of what she is trying to do, and her interest in succeeding, that she is justified in

inquiring as she does. If that's right, then there can be practical justifications that serve to deflect epistemic criticism.¹² It is worth emphasizing that this is not simply an instance of what we might call "the pragmatics of inquiry and belief management" (Titelbaum 2015). The point is familiar that we are cognitively and temporally limited creatures with bounded capacities to put toward figuring out what to believe, such that even if we had no ability to act, we would still need to avoid cluttering our minds with trivialities and employ certain heuristics and shortcuts that are not ideally reliable, but that cut down on time and effort (Harman 1986). The practical considerations at issue here, however, go beyond our merely cognitive or intellectual limitations. The suggestion is that our epistemic assessments—of inquiry, at least—can depend on what else the thinker in question is trying to do.

This puts pressure on the idea that the epistemic norms of evidence-gathering apply categorically to all similar thinkers in a way that is isolated from differences in practical interest. Rather, if there are epistemic norms on inquiry, they seem to be able to interact with practical considerations. The extension of the epistemic to the zetetic, even in its more modest form, goes hand in hand with embracing pragmatic encroachment.

This does not yet entail that there is pragmatic encroachment on the epistemic norms that govern our doxastic *states*. One could try to remain a purist about these norms. The question will rearise, however, when we try to assess the interim doxastic states formed by the inquiring agent in light of the evidence he has gathered. Supposing these states conform to what his evidence supports, should we take him to be justified in holding them, or not? We might want the ability to say that if the etiology of a belief includes epistemically criticizable evidence-gathering practices, then this can negatively affect the thinker's justification for holding the belief (Avner and Scott-Kakures 2015; Hughes 2023). When motivated reasoning, confirmation bias, and the manufacturing of evidence lead one to possess a body of misleading evidence that supports racist beliefs, we might want to hold that the beliefs are problematic even if they are a rational response to that evidence. But if the criticizability of one's evidence acquisition can matter to doxastic justification, and if practical considerations can matter to the criticizability of one's evidence acquisition, then practical considerations can matter to doxastic justification. This would amount to an inquiry-based argument for pragmatic encroachment on the justification of belief.

This train of thought would need much more slow and careful argument to be fully convincing. There are many places to get off the boat before reaching this conclusion. My more modest contention is that thinking about cases like inquisitive attempts is fertile ground for trying to understand the interaction between practical and epistemic normativity, and that these two things seem to be intertwined here in a way that poses an

¹² It's true that in virtue of the practical context, the positive thinking supported by the kind of inquiry I have described will be partly truth-conducive, making it more likely that the agent's positive expectations will be borne out than if she did not have those expectations. But importantly, they are not self-fulfilling. And if we bracket the agent's practical interest in succeeding, purely truth-conducive considerations would often point in the direction of negative thinking, which is much more likely to be self-fulfilling by leading to a negative answer. This is why the practical considerations make a non-alethic difference here.

apparent challenge for those who wish to defend certain kinds of “purism” about epistemic normativity.

References

- Anscombe, G.E.M. (1979). Under a description. *Noûs*, 13(2), 219–233.
- Anscombe, G.E.M. (2000). *Intention*, 2nd edition. Harvard University Press.
- Avnir, Y. & Scott-Kakures, D. (2015). How irrelevant influences bias belief. *Philosophical Perspectives*, 29, 7-39.
- Bandura, A. (1997). *Self-Efficacy: The Exercise of Control*. Worth Publishers.
- Bortolotti, L. (2018). “Optimism, agency, and success.” *Ethical Theory and Moral Practice* 21, 521-35
- Davidson, D. (1963). Actions, reasons, and causes. Reprinted in *Essays on Actions and Events*, 2nd Edition, 2001. Oxford University Press.
- Falbo, A. (2023). Inquiring minds want to improve. *Australasian Journal of Philosophy* ,101(2), 298–312
- Flores, C., & Woodard, E. (2023). Epistemic norms on evidence-gathering. *Philosophical Studies*, 180(9), 2547-71.
- Friedman, D. (2025). Intentions and inquiry. *Mind* 134(533), 85-106.
- Friedman, J. (2020). The epistemic and the zetetic. *Philosophical Review*, 129(4), 501-536.
- Friedman, J. (2025). Inquiry in epistemology: doxastic attitudes as zetetic attitudes. *A Companion to Epistemology*, 439-49.
- Frost, K. (2014). On the very idea of direction of fit. *Philosophical Review*, 123(4), 429-84.
- Harman, G. (1976). Practical reasoning. *The Review of Metaphysics*, 431-463.
- Harman, G. (1986). *Change in View: Principles of Reasoning*. The MIT Press.
- Honoré, A.M. (1964). Can and can't. *Mind*, 73, 463-479.
- Hornsby J. (201). Trying to act. In *A Companion to the Philosophy of Action*. Wiley-Blackwell, 18–25.

Hughes, N. (2023). Epistemic feedback loops (or: how not to get evidence). *Philosophy and Phenomenological Research*, 106(2), 368-393.

Kelly, T. (2022). *Bias: A Philosophical Study*. Oxford University Press.

Kelly, T. (2008). Disagreement, dogmatism, and belief polarization. *The Journal of Philosophy*, 105(10), 611-33.

Kikkert, S. (Forthcoming). Ability knowledge and epistemic disadvantage. In *The Epistemology of Ability*, B. Vetter & T. Schoonen, eds. Oxford University Press

Lehrer, K. (1968). Cans without ifs. *Analysis*, 29, 29–32.

Maier, J., and S. Kikkert. Abilities. *The Stanford Encyclopedia of Philosophy* (Summer 2025 Edition), Edward N. Zalta & Uri Nodelman (eds.).

Marušić, B. (2015). *Evidence and Agency: Norms of Belief for Promising and Resolving*. Oxford University Press.

Mele, A. (2003). Agents' abilities. *Noûs*, 37, 447–470.

Morton, J. M. (2022). Resisting pessimism traps: the limits of believing in oneself. *Philosophy and Phenomenological Research*, 104(3), 728-46.

Morton, J. M., & Paul, S. K. (Forthcoming). *Striving: Uncertain Agency*. Oxford University Press.

Morton, J. M., & Paul, S. K. (2019). Grit. *Ethics*, 129(2), 175-203.

Paul, L. A. (2014). *Transformative Experience*. Oxford University Press.

Paul, S. (2020). *The Philosophy of Action: A Contemporary Introduction*. Routledge Press.

Searle, J. (1983). *Intentionality*. Cambridge University Press.

Seligman, M. (2011). *Learned Optimism: How to Change Your Mind and Your Life*. Knopf Doubleday Publishing Group.

Tenenbaum, S. (2020). *Rational Powers in Action: Instrumental Rationality and Extended Agency*. Oxford University Press.

Thompson, M. (2008). *Life and Action: Elementary Structures of Practice and Practical Thought*. Harvard University Press.

Titelbaum, M. G. (2015). Continuing on. *Canadian Journal of Philosophy*, 45(5-6), 670-691.

Velleman, J. D. *Practical Reflection*. Princeton University Press, 1989.

Vetter, B. (2024). Abilities and the epistemology of ordinary modality. *Mind*, 133(532), 1001-27.

Zhang, Gabby. (Forthcoming). Inquiry and the inquiring way. *Philosophers' Imprint*.