

Penultimate draft, accepted for publication at *Philosophical Studies*. Please cite published version once available.

Alienation or Regress:

On the Non-Inferential Character of Agential Knowledge

Abstract: A central debate in philosophy of action concerns whether agential knowledge, the knowledge agents characteristically have of their own actions, is inferential. While inferentialists like Paul (2009a) hold that it is inferential, others like O'Brien (2007) and Setiya (2007, 2009, 2008) argue that it is not. In this paper, I offer a novel argument for the view that agential knowledge is non-inferential, by posing a dilemma for inferentialists: on the first horn, inferentialism is committed to holding that agents have only alienated knowledge of their own actions; on the second horn, inferentialism is caught in a vicious regress. Neither option is attractive, so inferentialism should be rejected.

Keywords: agential knowledge, practical knowledge, inferentialism, cognitivism, alienation, regress.

Alienation or Regress:

On the Non-Inferential Character of Agential Knowledge

Word Count: 4,702 (6,017 with footnotes)

July 10, 2020

Since Anscombe ([1957] 1963) introduced the notion, philosophers of action have sought to understand the nature of agential knowledge, the knowledge agents characteristically have of their own actions.¹ A central debate in the literature concerns whether agential knowledge is inferential. While inferentialists like Sarah Paul (2009a) hold that it is, others like Lucy O'Brien (2007) and Kieran Setiya (2007, 2008, 2009) contend that it is not. In this paper, I offer a novel argument for the view that agential knowledge is non-inferential, by posing a dilemma for inferentialists: on the first horn, inferentialism is committed to holding that agents have only alienated knowledge of their actions; on the second horn, inferentialism is caught in a vicious regress. Neither option is attractive, so inferentialism should be rejected.

Let me begin by explaining in more detail what agential knowledge is. I shall assume that agential knowledge is, in the first instance, knowledge of ongoing action (what I am doing) and action in the future (what I will do): the ordinary ways in which I know that I am raising three fingers or that I will go for a walk will thus be paradigm examples of agential knowledge. As Anscombe noted ([1957] 1963), there seems to be an important sense in which this knowledge is non-observational (pp.13-15):² ordinarily, I do not find out how many fingers I am raising on the

¹ Better: "reintroduced", as the notion appears already in the writings of Avicenna, Maimonides, and Aquinas. The literature on agential knowledge is extensive. See Schwenkler (2012) for a helpful review. Unless otherwise noted, when I speak of 'actions' I mean 'intentional actions'. I use the term 'agential knowledge' where many others use 'practical knowledge', because I think there are forms of practical knowledge that do not concern an agent's actions (Piñeros Glasscock ms.). However, I shall not argue for this point in the paper, and the reader is free to substitute their preferred label.

² Anscombe goes on to argue that agential knowledge is non-observational because it is "practical", insofar as it is "the cause of what it understands" (p.87). See Moran (2004) on the relation between the two.

basis of looking at my hand, nor do I ordinarily find out that I am going for a walk after lunch by making a prediction based on past occurrences. Rather, agents ordinarily seem to know what they are doing and what they are going to do *directly*. How to understand this characteristic directness, however, is among the disputed questions between inferentialists and non-inferentialists.

Another characteristic of agential knowledge is what we might call ‘spontaneity’. Agential knowledge is available from the practical point of view from which the action is regarded by the agent as up to her. I know that I am walking to the store, or that I will do so in 15 minutes, because *I settle* to do so. By contrast, ‘alienated’ knowledge, knowledge had from a standpoint from which the action is regarded as a settled matter, is not agential knowledge.³ The addict’s reply, “I am going to smoke a cigarette today, I just know it—you see, I am addicted” might thus express knowledge, but not of the agential kind. An indication of this is that such a statement has the character of a mere prediction, rather than an expression of intention.⁴ The claim that agential knowledge is spontaneous will play a crucial role in my argument, and, as such, will be explored further below.

Determining whether agential knowledge is inferential will give us insight into the nature of this important phenomenon. In addition to its intrinsic interest, however, a determination on this question would have repercussions for further debates in the philosophy of action. One pertains to the nature of intention. According to cognitivism, intentions are, or are constituted by beliefs, precisely the kinds of beliefs capable of constituting agential knowledge.⁵ However, inferentialism

³ I am indebted here to Moran’s (2001) and Marušić’s (2015) work on alienation and the agential stance.

⁴ On the difference, see Anscombe ([1957] 1963).

⁵ Cognitivists include Velleman (2001), Marušić and Schwenkler (2018), and Setiya (2007, 2008), though in his latest work Setiya states the view in terms of increases in epistemic confidence when one acts intentionally. Weaker versions of cognitivism hold that having an intention entails having a belief, but take intentions to be a separate state (see e.g. Clark (2020) for a recent defense of such a view). Setiya’s argument against inferentialism below (pp.6-7) shows why even these weaker views sit uneasy with inferentialism.

sits uneasy with such a cognitivist position because, as we shall see, on the best inferential models intention figures as a factor independent of the beliefs that constitute agential knowledge.⁶ Relatedly, the determination bears on the thesis that Kieran Setiya (2009) identifies as the central point at issue between Anscombeans and their opponents. This is the thesis that, like our perceptual faculties, the will is “a capacity for practical knowledge” (p.131). It would go a long way towards vindicating this view if we could show that the will, like perception, gives us knowledge directly without the need of an inference.

With this in mind, consider a case where I will go for a walk after lunch and I know this in the way characteristic of agential knowledge. According to inferentialism, I know that I will go for a walk on the basis of an inference. What distinguishes agential knowledge from other sorts of inferential knowledge is its content (it concerns actions) and its bases (the inference relies on knowledge of mental states—intentions or decisions—of the agent). I shall concentrate on what I consider the best inferentialist account in the literature, defended by Paul (2009a), though the criticism of inferentialism I shall raise should generalize to any version of the view.⁷

On Paul’s account, agential knowledge is acquired in two stages. First, an agent takes advantage of her privileged access to her mental states to come to know that she has an intention to Φ .⁸ At the second stage, the agent makes an inference on the basis of this knowledge. Paul argues that generally and in the right circumstances, intending to Φ suffices for Φ ing.⁹ Similarly,

⁶ Here, I am in agreement with Paul (2009a), who argues that Grice’s (1971) cognitivist inferentialism is an unstable position, lacking the resources to explain how we could ever have practical *knowledge*, as our beliefs about our actions “would be neither based on evidence nor self-fulfilling” (p.20).

⁷ Other prominent inferentialists include Grice (1971), on whose work Paul is building, and O’Shaughnessy (1980, p.515; 2003, p.349). On the other hand, a view like Marušić and Schwenkler’s (2018) (and arguably Anscombe’s), on which agential knowledge is gained on the basis of a *practical* inference is not called into question by my argument.

⁸ In her (2012), Paul presents an account of the mechanism by which we come to know our intentions on the basis of our decisions. This view might explain why knowledge of our intentions is spontaneous. Even if it does, I shall argue that the spontaneity is not transferred to knowledge of our actions on an inferential account.

⁹ At least if Φ is of the right type (e.g. if you’re in the middle of a street, and you form the intention to walk, you will be thereby engaged in the process of walking (note the progressive aspect on the formulation)). As Paul notes (p.16)

for action in the future: in the right circumstances, an agent who forms an intention to Φ (at some future time), will generally Φ in the future.

Given this general and reliable connection between intention and action, then, an agent can reason in accordance with the following schema:

Inferentialist Schema (IS)

- (1) I intend to Φ [in circumstances C_1, C_2, \dots, C_n].
- (2) If I intend to Φ [in circumstances C_1, C_2, \dots, C_n], then I am Φ ing/I will Φ .
- (3) Therefore, I am Φ ing/I will Φ .¹⁰

A few remarks about IS are in order. First, IS represents a schema of the contents that certain beliefs must have to warrant drawing a certain conclusion. A belief with content (3) will be inferentially warranted by beliefs with content (1) and (2) only if the subject is also warranted in believing (1) and (2). However, to acquire such warrant further conditions are needed. For instance, the subject must *base* her belief in (3) on her beliefs in (1) and (2);¹¹ and her warrant for believing (1) and (2) should not be based on (3), at pains of circularity. Second, the qualification about circumstances is in square brackets because there are more or less demanding ways of understanding this qualification. On the more demanding way of understanding it, the qualifications are part of the content of the beliefs that the agent employs in the inference. On a less demanding account, it is sufficient if the relevant circumstances obtain at the time at which the agent draws the inference. Although Paul endorses the more demanding view, I want to remain

further complications arise when considering telic action descriptions, but I ignore these complications since they are irrelevant to the issue at hand.

¹⁰ For simplicity, I take agential knowledge to be propositional, though see Thompson (2008) and Hornsby (2016) for arguments that suggest otherwise. However, the question is orthogonal to the present argument, which can be formulated *pari passu* in terms of self-ascriptions.

¹¹ This is a standard requirement on inferential warrant, though it is a difficult question what it takes for a belief to be appropriately based such as to meet that condition. See Boghossian (2014) and the reply by Wright (2014) for two approaches to the issue, and discussions of the difficulties.

neutral on this issue, since my argument is aimed at inferentialism generally, and it is consistent with inferentialism to endorse the less demanding view. Now, *which* circumstances are included in this qualification? Paul explicitly refrains from offering an exhaustive list (p.14), but she notes that plausible candidates include “knowledge of ability, conduciveness of circumstances, and one’s history as an agent [who reliably does what she intends]” (p.15), as well as knowledge that “many action descriptions function in such a way as to apply truly in large part because that description is what the agent intends” (p.16). Yet, inferentialist views might differ with respect to which circumstances they take as relevant, so I also want to remain neutral in this regard.

Inferentialism economically explains agential knowledge, relying on the independently plausible claims that we have privileged access to our intentions, and that knowledge acquired on the basis of such access can serve as a basis for inferences. Yet, one might worry that inferentialism fails to capture the characteristic directness of agential knowledge since its acquisition is on this account inferentially mediated. Paul takes this objection head on, arguing that the directness is accounted for by the fact that inferences “can take place rapidly and automatically at a non-conscious level, without the mindful entertaining of premises or feeling of drawing a conclusion” (p.10). One could object that agential knowledge should not be mediated even by unconscious processes, but absent further reasons to believe this, the objection is unpersuasive.¹²

A different line of objection has been pursued by Setiya (2008, 2007). It begins with the contention that there is a necessary connection between intentional actions and certain cognitive entities, such as beliefs: if one is Φ ing intentionally, one must believe that one is Φ ing.¹³ If any

¹² This reply can also serve as the basis to answer O’Brien’s (2007) objection that “it does not *seem* to me that I know I am raising my arm as a result of inferring that I am from my knowing that I intend to raise my arm” (p.177). The unconscious shape that inferences can take might make knowledge of our actions seem immediate without their actually being so.

¹³ Setiya argues for a weaker connection in terms of confidence, but since the specifics are not important for our purposes, for simplicity I focus on this one. Similar principles are defended by Grice (1971), Harman (1986a, 1997, 1986b), Velleman (2014, 2001), and Marušić and Schwenkler (2018).

such necessary connection holds, Setiya argues, inferentialism is false. The reason is that according to inferentialism the connection between practical and cognitive entities is mediated by an inference, but since we can always fail to draw an inference, the connection between the action and belief turns out to be contingent rather than necessary on an inferentialist view.

The problem with this argument is that it is questionable whether there are such necessary connections. Indeed, Paul (2009b) has provided persuasive arguments against this position, taking it rather as an *advantage* of inferentialism that it does not require such necessary connections (Paul 2009a). Hence, Setiya's argument is at best dialectically weak.¹⁴

Nevertheless, I shall argue that inferentialism should be rejected, and that practical knowledge is non-inferential. The objection I pursue is in the form of a dilemma: in short, inferentialists must either hold that we can know our actions only from an alienated standpoint, or they must admit that their account is subject to a vicious regress.

In presenting the dilemma, I begin with knowledge of future action, where the problems are most salient. This would be enough to call inferentialism into question, which purports to offer an account of agential knowledge in general, both knowledge of action in progress and knowledge of future action. As I shall argue, however, once we see that inferentialism has trouble accounting for knowledge of future action, it will be plain that it also has trouble accounting for knowledge of action in progress.

Consider, then, an ordinary case of knowledge of future action, such as my knowledge that I will take a walk after lunch. According to inferentialism, I know this on the basis of an inference with the following structure:¹⁵

(4) I intend to take a walk after lunch.

¹⁴ It is strong only against a view like Grice's (1971), who upheld both inferentialism and cognitivist principles.

¹⁵ For simplicity, I omit the qualification about circumstances. Let's assume they are met throughout.

(5) If I intend to take a walk after lunch, I will take a walk after lunch.

(6) So, I will take a walk after lunch.

As noted above, one of the seeming advantages of inferentialism is its ability to explain agential knowledge by appeal to independently plausible claims. For, according to inferentialism, the knowledge that (4) and (5) give me of (6) is ordinary inferential knowledge, on a par with my knowledge of (9) on the basis of (7) and (8).

(7) The fire is burning in the chimney.

(8) If the fire is burning in the chimney, the room will warm up.

(9) So, the room will warm up.

Yet, there is a crucial disanalogy between the two inferences, pertaining to their conditional premises: whereas the connection in (8) between the fire's burning and the room's warming up obtains through an independent process that can run its course without any contribution on the part of the agent, the connection in (5) between the intention and the action is one that necessarily requires the person to execute it.¹⁶ Another way to put this is that the truth of the conclusions are oppositely related to the agent's involvement: (9) will obtain unless the agent intervenes, whereas (6) won't obtain unless she does. This crucial difference gives rise to a dilemma.

First Horn

On the one hand, the agent might hold the belief in (5) on the basis of a mere empirical generalization, as one might hold (8).¹⁷ She realizes that in situations such as the one she finds

¹⁶ The process must be the particular process of *execution* since otherwise the resulting action won't be intentional, even when caused by an intention with the appropriate content (Davidson 2001; Frankfurt 1978).

¹⁷ This seems to be the best way to ensure that the belief in the conditional claim is not dependent on belief in execution. The two horns of the dilemma at a more general level are as follows: either (a) belief in the conditional claim is not dependent on belief in the execution or (b) it is so dependent. (a) leads to alienated knowledge, and (b) to a regress (assuming inferentialism).

herself in, where, for instance, she knows she has the ability and opportunity to execute her intention, she generally executes it (just as, generally, when the fire is on, the room warms up). However, in an ordinary case it is up to the agent whether she executes her intention or not. Unlike the case of the fire, where nature simply needs to run its course, the action will not execute itself, so whether the conditional premise holds true depends on what she does.¹⁸ True, not going for a walk (say, out of laziness) might constitute a form of practical irrationality, but the point is that it is up to the agent—the very one considering what she is going to do—to act rationally or irrationally, as she should recognize. To recognize this, to take the execution of the intention as being up to oneself, is constitutive of the agential standpoint; to think otherwise, to take the execution as something that is simply settled given one’s circumstances (even circumstances that include one’s intentions), is a paradigm example of bad faith, taking an alienated standpoint on a phenomenon that is within one’s control.¹⁹ Such an alienated standpoint can of course yield knowledge: the addict can, after all, come to know that she will be taking a hit before the end of the day by reflecting on her addicted urge. However, such knowledge would not be agential knowledge.

It might be objected that the addict’s situation is crucially different from the situation of the person inferring that she will take a walk. After all, the addict’s knowledge is not grounded in an intention. She has not yet formed an intention, and if she had, it would be unclear whether her knowledge is alienated in a way incompatible with agential knowledge. This difference is crucial,

¹⁸ This point bears on an objection raised to me by Sarah Paul: Can’t the agent simply have it as a background assumption that she won’t change her mind (as one might have it as a background assumption that there will be sufficient oxygen for the fire to keep burning)? I answer that she cannot: the agent would at least have to assume the stronger claim that her mind is set on doing what she intends, i.e. that she will execute her action (given the considerations in n14). More importantly, the same worry would arise for such a background assumption: on what basis would the agent assume (or be warranted in assuming) that she won’t change her mind? Given that she is capable of changing it, she cannot from the agential standpoint assume that her intention settles that she will not.

¹⁹ My understanding of bad faith is indebted to Moran’s (2001) now classic work on the topic.

since intentions are plausibly taken as the loci of agency. Hence, one might argue that if an agent knows about her action based on knowledge of her intention (and knowledge that the relevant circumstances obtain) that suffices to guarantee that her knowledge of the action is spontaneous.

To see why those conditions are insufficient, consider the following case.

Compelling Intentions: Moria often forgets her plans, so she carries around an agenda where she writes down what she has decided to do as soon as she forms an intention. On its pages, you can read phrases like: *At 5pm I shall go to the bank*; or *After lunch, I shall go to the gym*. These notes have for Moria the force of absolute and irrevocable commands: provided that on the given day she has the ability and opportunity to execute these actions she will do so. Thus, if Moria read on a given day, *Today I will go to the gym*, she would not deviate from that plan even if you offered her a million dollars to do so. Once written down, the matter is settled for her. This is why these notes are such a reliable source of information. If you asked Moria what she was doing on a given day, she would check her notes and reason as follows: ‘I wrote that I am going to Φ ; I generally do what I wrote down (when I’m able to, etc.); hence I am going to Φ ’. On that basis, for instance, she might conclude and come to know that she is going to the gym later on.

It is clear that Moria can generally acquire knowledge of her actions on the basis of what she reads in her diary. But it is equally clear that this knowledge is not agential knowledge, since it is alienated. The knowledge is alienated despite being grounded in an intention that the agent forms while correctly assuming that the appropriate circumstances will obtain and that she will not change her mind about what to do. This shows that even knowledge of an action that is grounded on an intention can be alienated.

Moria's case shows that knowledge gained on the basis of an intention does not guarantee spontaneity. Granted, the case is far from standard. However, it is hard to see how the differences between it and standard cases could be exploited by an inferentialist theory to explain why Moria lacks spontaneous knowledge of her actions. In particular, notice that we can assume that Moria correctly writes what her intentions are because she has the special access to them that agents normally have. In addition, we could modify the case so that only a few milliseconds go by between the time at which Moria writes down her intention and the time when she reads her agenda. Moria's knowledge would be equally alienated in such a case. The central difference between Moria's case and an ordinary one is therefore in the manner in which the agents keep track of their intentions: agents standardly do so through memory, while Moria does it through her agenda. Clearly, though, this feature does not explain why Moria's knowledge is alienated. After all, I can have agential knowledge of my actions even if I have to check my calendar to remind myself what I have planned!²⁰

The intuitive reason why Moria's case differs from the standard one (including the standard one where we have to aid our memory by the use of an agenda) is this: when Moria reflects about what she is going to do, she does so from a standpoint that fails to recognize that what she is going to do depends on what she thinks about the matter *from that very standpoint*. Moria thus regards her actions from the standpoint of a mere spectator, simply waiting to see how a certain process unfolds, albeit one that begins with an intention formed by her past self. But that is not the standard way in which we regard our intentions. Our intentions are standardly maintained through our

²⁰ This also addresses another worry. Might it be that Moria's case is alienated because she needs to consciously rely on external factors to figure out what she intends, rather than (as I claim) because she takes herself to have no further say about what to do given her previously formed intentions? The problem is that the proposed explanation overreaches: an agent who relies on her calendar to remember what she plans to do (as most of us do), need not have alienated knowledge, despite her conscious awareness of relying on the calendar. Her knowledge won't be alienated if she takes the decision to still be up to her. Thanks to Kenny Easwaran and Sergio Tenenbaum for pushing me to address this problem.

ongoing commitment to act in a certain way, and the agential standpoint is the standpoint from which our actions are regarded as so dependent on that commitment. It is thus constitutive of the agential standpoint from which practical knowledge is available to regard the execution of an action as in our power.

I have so far assumed that spontaneity is an essential feature of agential knowledge. However, this assumption can be disputed.²¹ For it might be suggested that the description ‘the knowledge an agent characteristically has of her intentional actions’ sufficiently captures what is essential to this knowledge, and that we can give content to this description without having to say that the knowledge is spontaneous. Agential knowledge does *appear* as spontaneous, which explains why we are inclined to think it so (e.g. in the cases of Moria and the addict). However, this appearance should be explained in the same way that Paul aims to explain the appearance of directness: it results from the fact that the beliefs that constitute such knowledge arise from very fast and unconscious inferences based on knowledge of intention. The speed and unconscious character of these inferences makes it appear as though there wasn’t any inferential transition. And, since intentions are (by all accounts) spontaneous states, the resulting beliefs also appear as spontaneous. Yet, these are *mere* appearances, and neither directness nor spontaneity are genuine features of the resulting knowledge.

However, this suggestion suffers from a number of difficulties that in my view make it untenable. First, although it is easy to see how the speed and unconscious character of a transition can explain why we think the resulting belief is unmediated, it is hard to see how such features could explain the fact that a property of the intention, spontaneity, is automatically ascribed incorrectly to the belief. Indeed, it is clear that generally these transitions *don’t* have such an effect.

²¹ I thank an anonymous referee for raising this objection and suggesting the appeal to unconscious processes to strengthen it.

Going for a run with a friend, I may form the intention to turn left on the next street, and rapidly and unconsciously form the belief that she will also turn left (because I tacitly believe that she will follow me); but I do not thereby take it that where she turns is up to me, in the way that turning left is up to me. Second, we can certainly make sense of a mere appearance of directness; but we shouldn't assume the same is true for spontaneity. After all, there are certain notions, like pain, for which the distinction between appearance and reality seems to dissipate (if it appears to you that you are in pain, then you are in pain). Now, recall that spontaneity consists in the way the agent regards her relation to the world, whether she takes what will happen to be up to her. Can we make sense of a distinction between the appearance of regarding something as up to us and so regarding it? That is far from obvious; and even if we can sometimes draw such a distinction, the relation between appearance and reality here seems intimate enough as to rule out the possibility of the kind of systematic error that would be needed for the proposal to work (it needs to be systematic, because the mere appearance is posited as a general feature of agential knowledge). Finally, and setting aside these difficulties, the proposal fails to explain all the phenomena in question. Although the appeal to appearances can explain why it would seem *to the agent* that her knowledge is spontaneous where it isn't, it cannot explain why someone else would take spontaneity as a criterion for *attributing* agential knowledge. Yet, where the argument relies on spontaneity, what matters are the external attributions. For instance, it is on the basis of *our* assessment of Moria's or the addict's cases as ones of alienated knowledge that I have claimed agential knowledge is spontaneous in a way that cannot be captured by inferentialism. Thus, I do not see a viable way for the inferentialist to escape the first horn through the proposed strategy.

Although I have presented the argument by focusing on knowledge of future actions, the problems stemming from alienation equally apply to knowledge of actions in progress. This is

because these actions too are sustained through our ongoing commitment to them. An agent can regard what she is doing from the standpoint of a mere spectator, as a process that merely unfolds before her eyes, or she can regard it from the standpoint of an agent, as a process that is itself sustained by and dependent on her commitment to execute it. And her knowledge will be spontaneous only if she regards her action from this second standpoint.

The problem with taking (5) as based on a mere empirical generalization—whether to gain knowledge of future or ongoing action—is now clear. It requires the agent to regard her action from an alienated standpoint, since it requires her to think of the execution of her action as something inevitable given her previous intention. Since belief from this alienated standpoint would, on the inferentialist view, necessarily be part of the justification that one has for believing that one will act in a certain way, this means that on the inferentialist view one could know about one's actions only from an alienated standpoint. But the phenomenon of agential knowledge is supposed to be special precisely in being knowledge from the agential standpoint. Hence, agential knowledge cannot be gained on the basis of an inference relying on (5), if (5) is understood to hold just on the basis of an empirical generalization.

Second Horn

Suppose the agent's belief in (5) is held from a non-alienated standpoint, compatible with the recognition that the connection between her intention and her action is dependent on her execution.

In that case, my knowledge of (5) would seem to depend on my further knowledge that:

(10) I will execute my intention to take a walk after lunch.

Now, either one's execution of an intention to Φ is identical with one's Φ ing, or it is not. Suppose it is identical (as I think we should hold for reasons considered below). In that case my executing

my intention to take a walk after lunch is identical to my taking a walk after lunch. But this would mean that (10)=(6). However, since on this view (10) is supposed to be the basis for holding (5), this means that the inference (4)-(6) is viciously circular, since the conclusion, (6), is based on a set of premises that includes itself. This shows that if execution and action are identical, and if one can have knowledge of conditionals linking one's intentions to one's actions from the agential standpoint, that knowledge cannot be prior to the knowledge of one's actions. This, of course, is incompatible with inferentialism.

Suppose, on the other hand, that the execution is not identical to the action. Then, in the case at hand we have a second action, viz. my execution of the intention to take a walk after lunch. However, how do I know that I will carry out *this* action; that is, on what basis do I hold (10)? There are two options: either the action is intentional or it is not. It would be hard to argue that it is not intentional: on the view under consideration the execution stands in the instrumental relation characteristic of intentional actions (presumably, I execute my intention to take a walk in order to take a walk after lunch), and the action is under my agential control: if I stop intending to execute, I will stop executing my action. More importantly, holding that it is not an intentional action would not help the inferentialist, since taking this view would push us back into problems related to the first horn of the dilemma: the agent would have to regard her action as settled by what's not under her control in the way characteristic of intentional action, namely, her execution, thus understood.

This only leaves one option open for the inferentialist: the execution of an intention to Φ is a different intentional action from the action of Φ ing. Since it is another intentional action, the inferentialist would have to hold that I know (10) on the basis of two further beliefs, namely:

- (11) I intend to execute my intention to take a walk after lunch; and
- (12) If (11), then (10).

Yet, the problems that arose with regards to (5) clearly arise also with regards to (12): how do I know this conditional claim? We can rule out that I know it as a simple generalization, since this would have me knowing (12) from an alienated standpoint. Hence, I must know it on the basis of the claim that I will execute *this* further action, i.e. on the basis of the further belief that:

(13) I will execute my execution of my intention to take a walk after lunch.

Clearly, though, this gives rise to the same problems as before, and we are caught in a vicious regress, since we are unable to give epistemic grounds that would provide warrant for (5).

One could reply that this is a problem for everyone. For the problem arises once we realize that to culminate in an action, the intention needs to be executed. If we take execution as an intentional action (as we seemingly should), then the agent must intend to do it as well; but then she must execute this further intention; and so on. However, this is only a problem if we think that the execution of the intention is not the same as the action. Avoiding such a regress seems to me a good reason to hold the identity claim instead. Yet, as we saw in the discussion of the second horn, the inferentialist cannot hold this without making the inference circular; but this is only because the inferentialist holds that knowledge of a conditional claim like (5) is more basic than knowledge that one will act. I suggest we reject that assumption.²²

Of course, the inferentialist could bite one of the bullets. For instance, she could hold that agential knowledge is available only from an alienated standpoint. However, we saw above that there are good reasons to believe agential knowledge must be spontaneous, and no apparent way for the inferentialist to account for the phenomena that motivate this view. Another option would be to admit that there is a regress but deny its viciousness. This would, at a minimum, require an

²² Thus, for instance, we could endorse a view like Setiya's (2011), on which knowledge of our intentions is posterior to knowledge of our actions. However, the argument is consistent with a view that gives no priority to either knowledge, and takes them as equally basic.

endorsement of epistemic infinitism, the position that a belief can be warranted on the basis of an infinite inferential chain.²³ Infinitism is an unpopular position, and commitment to it seems to me too steep a price to pay.²⁴ There is thus no good way to escape the dilemma while holding onto inferentialism. I therefore conclude that agential knowledge is non-inferential.²⁵

References

- Anscombe, G. E. M. ([1957] 1963). *Intention* (2nd ed.). Oxford: Basil Blackwell.
- Boghossian, P. (2014). What is inference? *Philosophical Studies*, 169(1), 1-18.
- Clark, P. (2020). Intentions, Intending, and Belief: Noninferential Weak Cognitivism. *Pacific Philosophical Quarterly*, doi:10.1111/papq.12308.
- Davidson, D. (2001). Intending. In *Essays on actions and events* (pp. 83-102). Oxford: Clarendon Press.
- Frankfurt, H. G. (1978). The Problem of Action. *American Philosophical Quarterly*, 15(2), 157-162.
- Grice, H. P. (1971). Intention and Uncertainty. *Proceedings of the British Academy*, 57, 263-279.
- Harman, G. (1986a). *Change in View* (Vol. 1): MIT Press.
- Harman, G. (1986b). Willing and Intending. In R. E. G. a. R. Warner (Ed.), *Philosophical Grounds of Rationality* (pp. 363-380). Oxford: Oxford University Press.
- Harman, G. (1997). Practical reasoning. *Review of Metaphysics*, 29(3), 431--463.
- Hornsby, J. (2016). Intending, knowing how, infinitives. *Canadian Journal of Philosophy*, 46(1), 1-17.
- Klein, P. D. (2011). Infinitism. In S. Bernecker, & D. Pritchard (Eds.), *Routledge Companion to Epistemology* (pp. 245-256): Routledge.
- Marušić, B. (2015). *Evidence and agency: Norms of belief for promising and resolving*: Oxford University Press, USA.
- Marušić, B., & Schwenkler, J. (2018). Intending Is Believing: A Defense of Strong Cognitivism. *Analytic Philosophy*, 59(3), 309-340.
- Moran, R. (2001). *Authority and estrangement: An essay on self-knowledge*: Princeton University Press.

²³ I say “at a minimum” because even infinitists recognize forms of epistemic regress that are vicious. Hence, the inferentialist would also have to show that the present regress is among the non-vicious ones.

²⁴ Peter Klein has been the foremost defender of the view among contemporary epistemologists. See Klein (2011) for a helpful survey of the literature.

²⁵ This paper is based on work from my dissertation, and was presented at the 2020 Eastern Session of the APA. I thank Randolph Clarke, Kenny Easwaran, Blakely Phillips, and Paul Silva for their questions, and especially Eugene Chislenko for his excellent comments and insightful conversation on the occasion. I also thank two anonymous referees for their comments, as well as David Charles, Jennifer Daigle, Stephen Darwall, Daniel Greco, Michael Della Rocca, Reier Helle, Daniel Moerner, Allison Piñeros Glasscock, and Jason Stanley for discussion. Finally, I owe special thanks to Sarah Buss, Sarah Paul, Sergio Tenenbaum, and Gideon Yaffe for discussion and written comments on earlier drafts. Research on this project was supported by the University of Toronto-Mississauga Postdoctoral Fellowship Award Program, sponsored by their Philosophy Department and the UTM Office of the Vice-Principal, Research.

- Moran, R. (2004). Anscombe on 'Practical Knowledge'. In J. Hyman, & H. Steward (Eds.), *Royal Institute of Philosophy Supplement* (Vol. 55, pp. 43-68): Cambridge University Press.
- O'Brien, L. (2007). *Self-knowing agents*: Oxford University Press.
- O'Shaughnessy, B. (1980). *The will* (Vol. 2): Cambridge University Press Cambridge.
- O'Shaughnessy, B. (2003). The Epistemology of Physical Action. In J. Roessler, & N. Eilan (Eds.), *Agency and Self-Awareness: Issues in Philosophy and Psychology*: Clarendon Press.
- Paul, S. K. (2009a). How We Know What We're Doing. *Philosophers' Imprint*, 9(11), 1-24.
- Paul, S. K. (2009b). Intention, belief, and wishful thinking: Setiya on "practical knowledge". *Ethics*, 119(3), 546-557.
- Paul, S. K. (2012). How we know what we intend. *Philosophical Studies*, 161(2), 327-346.
- Piñeros Glasscock, J. S. (ms.). Authoritative Knowledge.
- Schwenkler, J. (2012). Non-Observational Knowledge of Action. *Philosophy Compass*, 7(10), 731-740.
- Setiya, K. (2007). *Reasons without rationalism*: Princeton University Press.
- Setiya, K. (2008). Practical knowledge. *Ethics*, 118(3), 388-409.
- Setiya, K. (2009). Practical Knowledge Revisited. *Ethics*, 120(1), 128-137.
- Setiya, K. (2011). Knowledge of intention. In A. Ford, J. Hornsby, & F. Stoutland (Eds.), *Essays on Anscombe's Intention* (pp. 170--197): Harvard University Press.
- Thompson, M. (2008). *Life and action*: Harvard University Press.
- Velleman, J. D. (2001). *Practical Reflection* Stanford, CA: David Hume Series: CSLI Publications.
- Velleman, J. D. (2014). What good is a will. *Rational and Social Agency: The Philosophy of Michael Bratman*, 83.
- Wright, C. (2014). Comment on Paul Boghossian, "What is inference". *Philosophical Studies*, 169(1), 27-37.