

Internalism without cartesianism privilege

Darren Bradley

To cite this article: Darren Bradley (08 Feb 2026): Internalism without cartesianism privilege, *Inquiry*, DOI: [10.1080/0020174X.2025.2454470](https://doi.org/10.1080/0020174X.2025.2454470)

To link to this article: <https://doi.org/10.1080/0020174X.2025.2454470>



© 2026 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



Published online: 08 Feb 2026.



Submit your article to this journal [↗](#)



Article views: 219



View related articles [↗](#)



View Crossmark data [↗](#)

Internalism without cartesianism privilege

Darren Bradley

PRHS, University of Leeds, Leeds, UK

ABSTRACT


Contemporary philosophy has seen a widespread rejection of Cartesianism, roughly, the thesis that we always know our own minds [Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press]. A number of philosophers, notably [Wedgwood. 2002. "Internalism Explained." *Philosophy and Phenomenological Research* 65 (2): 349–369], [Srinivasan, A. 2015. "Normativity without Cartesian privilege." *Nous-Supplement: Philosophical Issues* 25 (1): 273–299], [Weatherson, B. 2019. *Normative externalism*. Oxford University Press] and [Lasonen-Aarnio, M. 2019. "Guidance, epistemic filters, and non-accidental ought-doing." *Philosophical Issues* 29: 172–183], have argued that anti-Cartesianism undermines the main argument for normative internalism, the thesis that what we ought to do depends on our internal mental states. I will defend a novel version of normative internalism which is well-motivated despite the rejection of Cartesianism. I will argue that this version of internalism has important advantages over externalism, showing how normative internalism can survive the absence of Cartesian privilege.

ARTICLE HISTORY Received 12 May 2024; Accepted 13 January 2025

KEYWORDS Internalism; justification; norms; guidance; Uber-norm

1. Introduction

Contemporary philosophy has seen a widespread rejection of Cartesianism, roughly, the thesis that we always know our own minds (Williamson 2000). A number of philosophers, notably Srinivasan (2015), Weatherson (2019) and Lasonen-Aarnio (2019), have argued that anti-Cartesianism undermines the main motivation for normative internalism, the thesis that what we ought to do depends on our internal mental states.¹ This motivation is that

CONTACT Darren Bradley  d.j.bradley@leeds.ac.uk  University of Leeds, Woodhouse Lane, Leeds, LS2 9JT, UK

¹For more on internalism see Conee and Feldman (1985), Pappas (2017) and Finlay and Schroeder (2017).

When the consequent of the norm is an action or intention it is a practical norm and when the consequent is a belief it is an epistemic norm. My arguments are intended to apply to both types of norm.

© 2026 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

internalist norms can always guide us in a way externalist norms cannot. I will defend a version of normative internalism which is well-motivated despite the rejection of Cartesianism. I argue that even if there is no one internal norm that can always guide us, nevertheless, in all situations, some or other internal norm can guide us. I then defend an Uber-norm which can provide guidance in any given situation.²

Section 2 explains the background. Section 3 reconstructs the normative externalist argument and section 4 shows how the argument can be avoided. Section 5 defends an Uber-norm, section 6 considers objections and section 7 concludes.

2. Setting the stage

In this section I explain and motivate normative internalism, and explain the objection from Anti-Cartesianism.

Following Srinivasan (2015, 277), we can define a norm as follows:

A norm is a universal generalisation about how an agent *S* is obligated or permitted to act, of the schema < if *C* then *S* ought to ϕ >, where ϕ ranges over the beliefs and basic actions available to *S*, and *C* is the norm's triggering condition.

This is a narrow-scope norm with a descriptive antecedent and a normative consequent. I will follow Srinivasan in focusing on epistemic normativity rather than moral, prudential etc. normativity, though I think my arguments could be extended to other domains.

Let's call a norm with *C* as the triggering condition (henceforth just 'condition') a *C*-norm. It will be important whether *C* describes an *internal* or *external* state. Let an *internal* state be a state necessarily shared by duplicates; external states need not be. Call a *C*-norm an *internal norm* if and only if *C* is an internal state.

We should distinguish two types of normative internalism (henceforth just 'internalism'):³

Mentalism: All justifiers are internal mental states

Accessibilism: All justifiers are internal mental states we have access to

We can take the access involved here to be: is in a position to know by reflection alone whether one is in the mental state. And for our dialectical purposes

²The closest precedent I can find is Greco (2017), who defends a contextualist version of internalism which is compatible with the view I develop, but Greco doesn't investigate what norms would look like given contextualism.

³See Feldman and Conee (2001), Bergman (2006, ch. 3). One qualification the internalist needs is to understand 'justifier' as 'contingent justifier'; logical principles are plausibly (necessary) justifiers, but are not mental states (Goldman 1999, 216–217). This fits with supervenience formulations of internalism e.g. Conee and Feldman (1985), Pryor (2001, 104), Gibbons (2006, 20).

we can work with agents who move from being in a position to know to actually knowing, allowing us to move between ‘S knows P’ and ‘S has access to P’. The Mentalist/Accessibilist distinction can be traced to Conee and Feldman (1985) who accepted only Mentalism. Accessibilism is considerably stronger and so more difficult to defend than Mentalism. I am not convinced that Accessibilism is true, but as my overall goal is to defend internalism, it makes dialectical sense to defend a logically strong version of internalism, so I will argue that Accessibilism is well-motivated despite the rejection of Cartesianism.

The relevant motivation for Accessibilism is based on the intuitive connection between norms and guidance – norms must be able to guide us (Pollock 1986, 10; Wedgwood 2002, 2017, ch. 7). (There is also an intuitive connection between norms and blame (Steup 1988, 1999). I will focus on guidance to keep things simple, but I think a similar argument could be developed for the connection between norms and blame.) Let’s put the point as follows:

Guidance Internalism:

For all agents, S, and possible actions, phi:

there is some non-trivial internal norm which can guide S’s deliberation about whether to phi.

(‘Non-trivial’ is Williamson’s term, meaning ‘contingent’, and is needed to allow that when it comes to states that all agents are *necessarily* in, we might always be in a position to know by reflection alone whether we are in them.⁴ A non-trivial norm is a norm with a non-trivial triggering condition. This will be important in section 5.) It is not plausible that *external* norms have such close connections to guidance. Assuming that such connections are requirements of genuine norms, *Guidance Internalism* provides support for internalism.

However, this story doesn’t seem to work given the widespread contemporary rejection of Cartesianism:

Cartesianism:

There are some non-trivial mental state conditions such that we are always in a position to know by reflection alone whether we are in them.

Most contemporary philosophers reject Cartesianism, including myself. Williamson (2000) gives roughly the following influential argument against *Cartesianism*. Terminology: Agent A’s belief that P is safe iff in all nearby worlds where A believes that P (and forms the belief via the same method as in the actual world), P is true. Assume the safety condition for knowledge: knowledge that P requires safe belief that P. Williamson argues that for all states,

⁴Call a condition non-trivial if and only if it obtains in some cases but not in all.’ Williamson (2002, 89).

there are borderline cases where in nearby worlds we are not in the state but still believe that we are; thus in such cases the belief that we are in the state is not safe, and therefore not knowledge. So, for example, an agent who is just about cold and believes they are cold might not know they are cold because there is a nearby world in which they are not cold but still believe (now falsely) that they are cold.⁵

I will grant that this argument succeeds. We arrive at:

Anti-Cartesianism:

There are no non-trivial mental state conditions such that we are always in a position to know by reflection alone whether we are in them.

The next section explains the argument that Anti-Cartesianism will lead to the rejection of Accessibilism, and the section after shows how the argument can be avoided.

3. The externalist argument

There are surprisingly few explicit arguments from Anti-Cartesianism to the rejection of Accessibilism. One explicit argument is given by Srinivasan (2015) so this section draws on her work and fills in the gaps where needed.⁶

Consider the following norm:

COLD

If S is cold then they should dance

Given Anti-Cartesianism, COLD can fail to provide guidance because you can be cold without being in a position to know by reflection alone that you are cold. The externalist argument extends the point to *all* norms as follows:

Anti-Cartesianism:

There are no non-trivial mental state conditions such that we are always in a position to know by reflection alone whether we are in them.

Guidance-Linking premise:

For all C-norms, if an agent is not in a position to know by reflection alone that they are in condition C then the C-norm fails to provide guidance

Therefore

Unusability:

⁵As I am granting the conclusion I am attempting to give the gist of Williamson's argument in a readable way. Williamson's version of the argument does not assume continuity in methods for safety, and also uses confidence levels rather than outright belief.

⁶Weatherson's (2019) rejection of his 1–4 argument on p.37 seems to be based on Srinivasan's argument. I connect Lasonen-Aarnio's argument at the end of section 5 and Wedgwood's argument in section 6.

Any non-trivial norm can fail to provide guidance

This threatens to break the tight connection internalists want between norms and guidance as expressed in *Guidance Internalism* (repeated for convenience):

Guidance Internalism:

For all agents, *S*, and possible actions, ϕ :

there is some non-trivial internal norm which can guide *S*'s deliberation about whether to ϕ .

I do not dispute *Unusability*; nevertheless, I will argue in the next section that Accessibilism maintains the close connection to guidance that internalists want, as there is always *some or other* non-trivial internal norm which provides guidance. Then in section 5 I make the stronger claim that one trivial norm always provides guidance.

4. The quantifier shift

Even though any non-trivial norm can fail to provide guidance (*Unusability*), I maintain that there is always *some or other* norm that can provide guidance, and that this is enough to motivate Accessibilism.

To appreciate the distinction we can appeal to the quantifier shift fallacy (Smith (2003, 32–33)). It is well known that changing the order of quantifiers can result in a fallacy, as demonstrated by the following example:

Every mountain has been climbed by someone. (AE)

Therefore, someone has climbed every mountain. (EA)

The fact that every mountain has been climbed by someone does not entail the stronger claim that one heroic climber has climbed every mountain. With this in mind, recall:

Guidance Internalism:

For all agents, *S*, and possible actions, ϕ :

there is some non-trivial internal norm which can guide *S*'s deliberation about whether to ϕ .

This could be given either of two readings. Start with the logically stronger, EA reading:

EA Guidance Internalism:

For all agents, *S*, and possible actions, ϕ :

there is *one specific* non-trivial internal norm which can guide *S*'s deliberation about whether to ϕ .

This says that one heroic norm can always guide us. I grant that this is refuted by the Externalist Argument that any non-trivial norm can fail to provide guidance (*Unusability*). However, the AE reading remains:


AE Guidance Internalism:

For all agents, *S*, and possible actions, ϕ :

there is *some or other* non-trivial internal norm which can guide *S*'s deliberation about whether to ϕ .

I claim that that *AE Guidance Internalism* survives the Anti-Cartesian objections and gives internalists the connection they want between guidance and norms. Accessibilism does not require that there is *one* condition such that agents always have access to the fact that they are in it (if they are); Accessibilism only requires that in all situations agents have access to *some or other* condition that they are in. And *that* provides the triggering condition of *some or other* norm that provides guidance.

Let's go through an example. Consider an agent who is cold, though in a borderline state of being cold, and not in a position to know by reflection alone that they are cold (#).



Let's assume:

COLD

If *S* is cold then they ought to dance⁷

We can grant that this agent cannot be guided by COLD, but this is no counter-example to:

AE Guidance Internalism:

For all agents, *S*, and possible actions, ϕ :

⁷This is a toy example of a norm. One feature which real norms would need is a totality clause, which I set aside until section 5. There also needs to be reference to a time, which I will sometimes leave out for the sake of readability.

there is *some or other* non-trivial internal norm which can guide S's deliberation about whether to phi.

Although they cannot be guided by COLD, there will be some *other* non-trivial internal norm which they can be guided by. It will plausibly be a norm with a logically weaker antecedent and correspondingly weaker consequent e.g. if S is not hot then they are permitted to dance.

And in other situations COLD will be a norm that they can be guided by. For example, there are cases where you are *very* cold, and where, in all nearby worlds, you are cold and also believe that you are cold. So you are in a position to know by reflection alone that you are cold. Call these cases of being Safely Cold. The cases where you are Safely Cold are a proper subset of cases where you are cold. So from:

COLD

If S is cold then they ought to dance

a fortiori, it follows that:

SAFELY COLD

If S is Safely Cold then they ought to dance

Consider the level of coldness indicated by the asterisk (*). The agent is Cold and Safely Cold (though not safely Safely Cold). They cannot be guided by SAFELY COLD as they are not in a position to know by reflection alone that they are Safely Cold (because they are not safely Safely cold). But they can be guided by COLD as they are in a position to know by reflection alone that they are cold.

This is analogous to a familiar internalist move in which norms are changed from involving the external world to involving just beliefs or perceptions e.g. the consequentialist might not be able to choose the best outcome, but they can choose what they expect will be the best outcome; the brain in a vat cannot respond to the external world, but they can respond to the appearance of what seems to be an external world.⁸ The difference is that here we are only dealing with mental states, and we are moving from mental states the agent does not have access to, to mental states the agent does have access to.

So far I have argued that there is always a non-trivial norm that can provide guidance, and I think this undermines the externalist argument. Those who

⁸See Bales (1971), Williams (1981) and Jackson (1991) for related classic discussions of internal reasons and the subjective ought. Guidance norms will express internal reasons/subjective oughts. There are plausibly cases where the guidance norm tells you to do something you know to be objectively wrong e.g. giving a low risk, low reward drug to a patient whose disease is unknown, rather than either of two drugs which will either kill or cure the patient, depending on which disease they have (Jackson 1991).

are just looking for a way to block the externalist argument can stop reading here. In the next section I will go further to defend the bolder claim that *one* norm can always provide guidance – the Uber-norm.

5. The Uber-norm

5.1. *The positive view*

I have conceded that there is no specific *non-trivial* norm that can always provide guidance (EA). However, in this section I will argue that there *is* a specific *trivial* norm that can always provide guidance. So there is a true EA guidance principle.

We first need to distinguish total and non-total states. Let the *total* internal state of an agent consist of every internal state the agent has at some moment; let a *non-total* internal state consist of some proper sub-set of the total internal state. I will argue that there is a *trivial total* mental state condition that we are always in a position to know by reflection alone that we are in.⁹

Next, consider the disjunction of every total mental state that is epistemically possible.¹⁰ By definition, it is trivial (i.e. necessary) that we are in one of those states, so we are in a position to know by reflection alone that we are in one of them. For example, one possible total mental state is the combination of {an intention to make dinner, a borderline experience of being cold, a non-borderline experience of seeing red}.

Next, consider the disjunction of every epistemically possible set of total *evidence* e.g. all the possible experiences. Here we're just setting aside all the non-evidential mental states and zooming in on the evidential mental states. By definition, it is trivial that we are in one of those total evidential states, so we are in a position to know by reflection alone that we are in one of them. Using our example, as an intention is not a mental state, the total evidential state is {a borderline experience of being cold, a non-borderline experience of seeing red}

Let's zoom in once more, to the disjunction of every epistemically possible total evidential *accessible* state. This is the disjunction of every total evidential state that an agent could know they are in by reflection alone. The disjunction is something like: S's total evidential accessible state is C or S's total evidential accessible state is D or ... By definition, it is trivial that we are in one of those total evidential accessible states, so we are in a position to know by reflection alone that we are in one of them. Using our example, as a borderline experience of being cold is inaccessible, the total evidential accessible state is just

⁹Total states are needed for the same reason epistemologists appeal to total evidence – a sub-set of the total evidence will not determine what the agent ought to believe or do.

¹⁰I have in mind the epistemic possibilities developed by David Chalmers in many places e.g. (2002).

{a non-borderline experience of seeing red}. (I assume that the experience of seeing red is not inaccessible for the some other reason.)

The next step is to connect each total evidential accessible state with a normative consequent e.g.

< if S's total evidential accessible state at t is C then S ought to phi at t >

< if S's total evidential accessible state at t is D then S ought to psi at t >

< if S's total evidential state accessible at t is E then S ought to tau at t >

...

We can put these conditionals together into a single conjunctive norm, the *Uber-norm*:

Uber-norm

< For any person S, time t, and proposition p,

if S's total evidential accessible state at t is C then S ought to phi at t;

and if S's total evidential accessible state is D at t then S ought to psi at t ... >¹¹

I suggest that this is a *trivial* norm which can always provide guidance. It will be useful to put the Uber-norm in terms of knowledge (for readability I use 'knows' rather than 'is in a position to know', as explained in section 2):

Uber-norm

< For any person S, time t, and proposition p,

if the total evidential state S knows they are in at t is C then S ought to phi at t;

and if the total evidential state S knows they are in at t is D then S ought to psi at t ... >

That's the positive view. The rest of the paper elaborates and discusses objections. In the rest of this section I will discuss what I will call the Inapplicability Objection and the Ignorance Objection.

5.2. The inapplicability objection

The Inapplicability Objection says that for some agents no conditional will apply to them. Specifically, the conditionals might not apply to an agent because the agent might not have access to what their total evidential state is. For example, the total evidential state of an agent might be that

¹¹See Christensen (2010) and Bradley (2019) for discussions of the Uber-norm.

they are cold without the agent having access to the fact that this is their total evidential state.

This would be a problem if the antecedents had the following structure:

< ... if S has access to the fact that their total evidential state is C ... >

or

< ... if S knows that their total evidential state is C ... >

But the Inapplicability Objection is based on a misunderstanding, for the antecedents in fact have the following structure:

< ... if S's total evidential accessible state is C ... >

or.

< ... if the total evidence S knows is C ... >

Put in terms of access, the antecedent refers to the total evidential state S has access to, and there will always be some total evidential state S has access to. Put in terms of knowledge, the antecedent refers to the total evidence that S knows. The truth of the antecedent does not require that S knows *that* C is their total evidence; it requires that C is the total evidence they know. They can know C, while C is the total evidence they know, without knowing that C is their total evidence. C is the strongest evidential proposition they know, but they might not know that C is the strongest evidential proposition they know.

For example, an agent's total evidential state might be that they are very cold, but they only know the logically weaker fact that they are cold. So although their total evidential state is that they are very cold, they do not know that their total evidential state is that they are very cold. Nevertheless, they can know that they are cold, and, assuming that that is all they know, the total evidential state they know they are in is coldness. So there is no problem concerning conditionals not applying to agents due to agents not knowing what their total evidence is; the conditionals apply to agents as long as there is some total evidential state they know they are in. And for all agents there is some total evidential state they know they are in.¹²

5.3. The ignorance objection

The Ignorance Objection says that agents might not know which antecedent is true (or have access to which is true). Even if the total evidential state an agent is in is coldness, it does not follow that they know that the total

¹²Technically there might be agents with literally no knowledge, but we can just take their knowledge to be the empty set, and the relevant conditional will include '... if the total evidential state S knows they are in is the empty set ...'

evidential state they are in is coldness. They might be ignorant of whether they are in a richer, logically stronger evidential state, such as being very cold, in which case they would not know the antecedent ‘the total evidence S knows is that they are cold’. And if they don’t know which antecedent is true, they don’t know which conditional applies to them. So someone might object that the *Uber-norm* cannot be used for guidance, as guidance would require knowing which conditional applies to oneself.

In response, Accessibilists need not demand that knowledge of the antecedent is needed for guidance; the relevant knowledge is *in* the antecedent. The Accessibilist need only claim that the conditionals relevant for guidance involve knowledge of internal mental states. That is, an agent can be guided by ‘if the total state S knows at t is that they are cold then S ought to dance at t’ even if they don’t know that the antecedent is true. The agent must know they are cold, but they do not have to know that they know they are cold. Higher order knowledge is not needed.

It is arguable that such an agent does have some kind of *implicit* higher-order knowledge that they know they are cold. Compare the way most speakers cannot be explicitly guided by the rules of grammar, but are still guided by them when speaking. It is plausible that they have implicit knowledge of the rules of grammar. Or consider how Huckleberry Finn is guided by the true moral rules in not turning in an escaped slave, even if he wrongly believes that he should turn in the slave (Arpaly 2002). It is plausible that he has implicit knowledge of the moral rules. In these cases the agent can be guided by their implicit knowledge by being *responsive* to this knowledge. Similarly, the agent who lacks explicit knowledge that they know they are cold might still be have implicit knowledge that they know they are cold, such that they can be guided by their implicit knowledge.¹³

Someone might object that an agent who conforms to the *Uber-norm* without explicit knowledge of which antecedent applies to them is following the *Uber-norm* by accident, so they are not really being guided by the *Uber-norm*. In response, I deny that non-accidentality requires explicit knowledge. Non-accidentally might require a modally robust responsiveness to the situation, but agents without explicit knowledge of the relevant norms can have such responsiveness. For example, Huckleberry Finn might reliably do the right thing across a wide range of situations despite a lack of explicit knowledge of what the right thing to do is. We need not take a stand on exactly what is necessary for guidance, but some kind of responsiveness to the relevant facts seems sufficient.¹⁴

¹³Could an agent bootstrap their way to explicit higher order knowledge by examining their own behaviour? This might be possible in some cases (e.g. working out grammatical rules), but in Williamson-style borderline cases, the agent will presumably not know that they have followed the conditional that applies to them.

¹⁴See Howard (2019) for an argument that higher-level attitudes are not important to normative evaluations. For dissent, see Quine (1970) and Johnson King (2018).

The objector might stress that we are interested in norms that guide *deliberation*, and insist that for internalist norms to have an advantage over externalist norms in guiding deliberation, the agent must have explicit knowledge of the norm.

But it's not obvious that explicit norms are needed to guide deliberation. For we can deliberate about which sentence of English to use and be guided by rules of grammar which are only implicit. Huckleberry Finn can deliberate about what to do, and be guided by the moral rules, despite his lack of explicit knowledge of them. The objector might be right that agents cannot always have the relevant norm explicitly in mind when they are deliberating, but this is a point the internalist can grant.

This is a good place to respond to Lasonen-Aarnio's (2019) argument. At the end of her paper (section 3.3) she considers a version of the strategy defended here, in which the triggering conditions which might have referred to the external world are replaced by triggering conditions which refer to how things seem to the agent.¹⁵ Lasonen-Aarnio objects that:

one still confronts cases with the structure that gives rise to the anti-luminosity argument:

If (1) it can seem to me that *p*, but I cannot discriminate my case from a very similar case in which it doesn't seem to me that *p*, and

if (2) the [triggering conditions state] whether or not it seems to me that *p*,

then (3) I am in no position to track the [triggering conditions] in a way that is not accidental.¹⁶ (180)

My response is that the Accessibilist can deny that (1) and (2) are compatible. The Accessibilist will hold that the triggering conditions of the norm that provides guidance cannot depend on which indiscernible fact is true, so if the agent cannot discriminate between it seeming that *p* and seeming that not *p*, then a norm starting 'if the total evidence *S* knows is that it seems that *p* ...' will not provide guidance to the agent. Instead, guidance will be provided by a norm with a triggering condition the agent *can* discriminate from its negation, which is likely to be some less fine-grained statement of how things seem e.g. 'if the total evidence *S* knows is that it seems that *p* or *q* ...'.¹⁷

It might be useful to compare Mentalists, who allow that internal mental states the agent does not have access to can be justifiers. Mentalists will

¹⁵Lasonen-Aarnio's 'normative reasons' are roughly my 'triggering conditions'. Lasonen-Aarnio also discusses Gettier cases to argue that whether an agent knows cannot always determine what they ought to do. I think this shows that *access* internalism should not be understood in terms of knowledge but in terms of access or evidence.

¹⁶I've added the numbers.

¹⁷Of course, there will be agents in the borderline states unable to discriminate *p*-or-*q* from its negation. Such agents need a conditional with an even weaker antecedent

allow that agents can be guided by conditionals that have internal mental states in the antecedent, and they will not demand that agents *have access* to these internal mental states. Internal mental states are privileged in a way that non-internal or non-mental states are not, and this allows internal mental states to be triggering conditions. If someone objects that the agent might not have access to which internal mental states they are in, the Mentalist will deny that that matters. Similarly, the Accessibilist holds that *accessible* internal mental states are privileged, and will deny that it matters whether the agent has access to which accessible states they are in. Put in terms of knowledge, the Accessibilist holds that it doesn't matter whether the agent knows which is the total evidential state they know they are in.¹⁸

What makes a state privileged in a way that allows the state to be a triggering condition? One plausible answer is that the state must be sufficiently integrated into the rest of the agent's systems. In particular, perhaps one important feature is that the state is sufficiently integrated with motor systems that control behaviour. For example, suppose an agent who is looking at a pointer which is pointing at 53 can only say that the pointer is somewhere between 52 and 54, but when they reach out to grab it, they reach directly for the location of 53 (Greco 2019). Then despite a lack of explicit knowledge that they know the pointer is at 53, they can still be guided by implicit knowledge that they know the pointer is at 53, and this is enough for the implicit knowledge to be a triggering condition. States about which agents have implicit but not explicit knowledge emerge as tricky cases, as they are more integrated than fully external states, but less integrated than states of which agents have explicit knowledge. This explains why they have led to controversy; but I have argued that these states are integrated enough to be triggering conditions.

6. Infinity and the uber-norm

In this section I elaborate upon and defend the Uber-norm from objections connected to infinity. The triggering conditions relevant to humans will usually be long conjunctions e.g. the visual experience of seeing green and the gustatory experience of a coffee after-taste and ... etc. Someone might object that these triggering conditions will need to be infinitely long.

In response, there is no need for finitude. We can allow infinite sets of beliefs e.g. 'I believe P, I believe that I believe P, etc.' (Treanor 2013, 580; Fales 2014) Indeed it is somewhat artificial to talk about the number of

¹⁸Someone might worry that these considerations lead us to a position which better supports Mentalism than Accessibilism. I agree. But, as I said above, I am assuming Accessibilism to make life difficult for myself, and showing that externalist objections are not fatal to Accessibilism.

beliefs agents have, so I don't think this is a concern. (Compare: How many propositions does a map represent? Plausibly, an infinite number.)

Furthermore, if there is a problem with access to infinitely long descriptions, the Uber-norm automatically avoids the problem, as it refers to the strongest state the agent has access to. If the agent's access is capped at, say 1000 conjuncts, then the thousand-conjunct proposition will be the strongest state the agent has access to.

A related worry is that Accessibilism cannot avoid a vicious regress. Such regress problems are well-known objections to Accessibilism. Wedgwood (2002) writes:¹⁹

the following worlds w_1 and w_2 both seem possible. In both w_1 and w_2 , you believe p on the basis of certain reasons, but in w_1 you [have access to the fact] that you believe p on the basis of those reasons, while in w_2 you [do not]; otherwise, you are in just the same mental states in both w_1 and w_2 . So, according to [Accessibilism], the fact that you believe p on the basis of these reasons may be part of what makes the belief rational in w_1 , but it cannot be part of what makes the belief rational in w_2 . Hence, [Accessibilism] must say that it could be the case that this belief is rational in w_1 but not rational in w_2 . But then the fact that [you have access to the fact that these reasons are] the basis for your belief in w_1 is itself one of the facts on which the rationality of the belief supervenes. The fact that you are in a position to know these facts is not merely a precondition of these facts determining whether the belief is rational. It is itself one of the facts that determine whether the belief is rational. So, according to [Accessibilism], you must also [have access to] that fact by reflection alone. Thus, if there is any set of facts that determines whether a belief is rational, then for every fact F that belongs to that set of facts, that set must include, not just F , but also the fact that one [has access to] $F \dots$, the fact that one [has access to the fact that one has access to F], and so on, ad infinitum. (352)

I take it the idea is that when we identify a purported justifier, X , we have to add to the justifier the fact that the agent has access to X . So the new purported justifier is that the agent has access to X . And if that is the new purported justifier, we have to add that the agent has access to *that*, so the new purported justifier is that the agent has access to the fact that the agent has access to X , and off we go to infinity.

Notice Wedgwood moves from an existential claim (offering one case in which we must add the agent's access to the justification) to a universal claim (stating that in all cases we must add the agent's access to the justification). We cannot generally move from an existential to a universal claim, so more needs to be done to show that the existential claim generalises.²⁰

¹⁹See also Fumerton 1995, 55 and 81 and Wedgwood 2017, 166–167.

²⁰This is a quantifier-related fallacy, but it is different from the quantifier shift fallacy above. Here, each sentence has only one quantifier rather than two, and the illicit move is analogous to inferring from an *instance* of a black raven to the *universal* claim that all ravens are black.

Moreover, we saw in the previous section that Accessibilism need not be committed to the agent having access to every fact on which the rationality of their belief supervenes.²¹ The rationality of the belief might supervene on a fact about access; the agent need not have access to that fact about access. In the example above, the rationality of *S*'s dancing supervenes on *S*'s access to their coldness, but *S* does not need to have access to the fact that they have access to their coldness.

Rather than being *ad hoc*, this fits with the motivations of the Accessibilist. The Accessibilist is motivated by a desire to restrict justifiers to a privileged set of states, specifically to rule out external justifiers. Wedgwood's objection is that there end up being too many justifiers. But the Accessibilist is in the business of restricting justifiers, so they will be happy to restrict justifiers even further in response to Wedgwood's objection. A full discussion of Wedgwood's objection would require a separate paper, but I hope I've done enough to show how the Accessibilist can reply.

7. Conclusion

I have argued that internalism is well-motivated despite the rejection of Cartesianism. Internalism need not be committed to the thesis that there is any one non-trivial internal norm which always provides guidance. Internalists can hold that there is always *some or other* non-trivial internal norm which always provides guidance. Furthermore, I argued that there is a trivial norm which always provides guidance – the Uber-norm. Although a trivial norm, it is still informative when combined with the agent's access to their internal states, however revealing or non-revealing that access might be.

For the purpose of open access, the author has applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising from this submission.

Acknowledgements

Thanks to audiences at the University of Leeds and the Leeds-Dianoia meta-ethics workshop in Rome for helpful feedback, and to Adam Elga and Dan Greco for very helpful discussion of earlier drafts.

Disclosure statement

No potential conflict of interest was reported by the author(s).

²¹See Hatcher (2018) for a version of this strategy as a response to Wedgwood.

Funding

This work was supported by John Templeton Foundation: [Grant Number 62603].

References

- Arpaly, Nomy. 2002. *Unprincipled Virtue: An Inquiry into Moral Agency*. Oxford: Oxford University Press.
- Bales, R. E. 1971. "Act-Utilitarianism: Account of Right-Making Characteristics or Decision-Making Procedure?" *American Philosophical Quarterly* 8 (3): 257–265.
- Bergman, M. 2006. *Justification without Awareness*. Oxford: OUP.
- Bradley, D. 2019. "Are There Indefeasible Epistemic Rules?" *Philosopher's Imprint* 19 (3): 1–19.
- Christensen, D. 2010. "Higher-Order Evidence." *Philosophy and Phenomenological Research* 81 (1): 185–215. <https://doi.org/10.1111/j.1933-1592.2010.00366.x>
- Conee, Earl, and Richard Feldman. 1985. "Evidentialism." *Philosophical Studies* 48 (1): 15–35. <https://doi.org/10.1007/BF00372404>
- Fales, E. 2014. "Turtle Epistemology." *Philosophical Studies* 169 (2): 339–354.
- Feldman, Richard, and Earl Conee. 2001. "Internalism Defended." *American Philosophical Quarterly* 38 (1): 1–18.
- Finlay, Stephen, and Mark Schroeder. 2017. "Reasons for Action: Internal vs. External." In *The Stanford Encyclopedia of Philosophy (Fall 2017 Edition)*, edited by Edward N. Zalta. <https://plato.stanford.edu/archives/fall2017/entries/reasons-internal-external/>.
- Fumerton, Richard A. 1995. *Metaepistemology and Skepticism*. Lanham, MD: Rowman & Littlefield.
- Gibbons, J. 2006. "Access Externalism." *Mind* 115 (457): 19–39. <https://doi.org/10.1093/mind/fzl019>
- Goldman, A. 1999. "Internalism Exposed." *The Journal of Philosophy* 96:271–293. <https://doi.org/10.2307/2564679>
- Greco, D. 2017. "Cognitive Mobile Homes." *Mind; A Quarterly Review of Psychology and Philosophy* 126 (501): 93–121.
- Greco, D. 2019. "Fragmentation and Higher-Order Evidence." In *Higher-Order Evidence: New Essays*, edited by M. Skipper and A. Steglich-Petersen. New York: Oxford University Press.
- Hatcher, M. 2018. "ACCESSIBILISM DEFINED." *Episteme; Rivista Critica Di Storia Delle Scienze Mediche E Biologiche* 15 (1): 1–23. <https://doi.org/10.1017/epi.2016.36>.
- Howard, Nathan Robert. 2019. "One Desire Too Many." *Philosophy and Phenomenological Research* 102 (2): 302–317. <https://doi.org/10.1111/phpr.12642>
- Jackson, F. April 1991. "Decision-theoretic Consequentialism and the Nearest and Dearest Objection." *Ethics* 101:461–482. <https://doi.org/10.1086/293312>
- Johnson King, Zoe. 2018. "Accidentally Doing the Right Thing." *Philosophy and Phenomenological Research* 100 (1): 186–206.
- Lasonen-Aarnio, M. 2019. "Guidance, Epistemic Filters, and non-accidental Ought-Doing." *Philosophical Issues* 29 (1): 172–183. <https://doi.org/10.1111/phis.12166>
- Pappas, George. 2017. "Internalist vs. Externalist Conceptions of Epistemic Justification." In *The Stanford Encyclopedia of Philosophy (Fall 2017 Edition)*, edited by Edward N. Zalta. <https://plato.stanford.edu/archives/fall2017/entries/justep-intext/>.
- Pollock, J. 1986. *Contemporary Theories of Knowledge*. Totowa, NJ: Rowman and Littlefield).

- Pryor, James. 2001. "Highlights of Recent Epistemology." *British Journal for the Philosophy of Science* 52 (1): 95–124. <https://doi.org/10.1093/bjps/52.1.95>
- Quine, W. V. O. 1970. "Methodological Reflections on Current Linguistic Theory." *Synthese* 21 (3-4): 386–398. <https://doi.org/10.1007/BF00484806>
- Smith, P. 2003. *An Introduction to Formal Logic*. Cambridge: Cambridge University Press).
- Srinivasan, A. 2015. "Normativity without Cartesian Privilege." *Nous-Supplement: Philosophical Issues* 25 (1): 273–299. <https://doi.org/10.1111/phils.12059>
- Steup, M. 1988. "The Deontic Conception of Epistemic Justification." *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 53 (1): 65–84.
- Steup, M. 1999. "A Defense of Internalism." In *Theory of Knowledge: Classical and Contemporary Readings*, edited by L. in Pojman, 373–384. Belmont, CA: Wadsworth.
- Treanor, N. 2013. "The Measure of Knowledge." *Noûs* 47 (3): 577–601.
- Weatherson, B. 2019. *Normative Externalism*. Oxford: Oxford University Press.
- Wedgwood. 2002. "Internalism Explained." *Philosophy and Phenomenological Research* 65 (2): 349–369. <https://doi.org/10.1111/j.1933-1592.2002.tb00206.x>
- Wedgwood, R. 2017. *The Value of Rationality*. Oxford: Oxford University Press.
- Williams, Bernard. 1981. *Internal and External Reasons*. Cambridge: *Moral Luck*, Cambridge University Press.
- Williamson, T. 2002. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.