

Teleological Momentum and Non-Basic Actions

J. ADAM CARTER *

Abstract: When an agent is performing a non-basic action (e.g., walking to the summit, writing a novel, baking a cake) what determines at any given time whether they are *still* performing that action as it unfolds? The puzzle of the imperfective aspect reveals that intention and physical movement are insufficient: two agents can share both the intention to ϕ and the same initial behavior, yet, only one is in the process of ϕ -ing. I develop a novel account on which an agent is performing a non-basic action A if and only if they possess sufficient *teleological momentum* toward A. Teleological momentum is, like physical momentum, a composite vector quantity comprising three elements: intentional mass, practical velocity, and guidance alignment. The TM framework explains the differential truth-conditions of progressive action ascriptions while avoiding what are shown to be pitfalls of purely modal, dispositional, intentionalist, and causalist accounts. It also has wider implications for how responsibility, akrasia, and collective action should be understood in cases of temporally extended agency, by clarifying how a non-basic action can persist over time through partial progress and counterfactual vulnerability in virtue of sufficient teleological momentum toward its end.

“Motion is the mode in which the future belongs to the present”

— Aristotle (via St. Thomas Aquinas)

1 | THE PUZZLE OF IMPERFECTIVE ACTION-ATTRIBUTION

One and the same stretch of behaviour can be described under many action-descriptions: *taking a step, walking to the summit, training for a marathon, keeping a promise*. Such descriptions are typically *non-basic* in the sense that if the agent is doing them at all, she is doing them *by* doing other things.

Many philosophically central non-basic actions are also what the aspectual literature calls *accomplishments* (Vendler 1957, 148). An accomplishment is *dura-*

*Draft version. Comments welcome. This work is supported by the European Research Council grant number 101199255 (KNOW-HOW).

tive; it unfolds over time and has intermediate stages; it's also *telic*: it is individuated by a characteristic endpoint or culmination. *Walking to the summit*, *finishing the book*, and *building the house* are paradigms. Unlike mere *walking* or *typing*, these descriptions build a goal into the kind of action they ascribe.

This paper asks a question that's easy to state and hard to answer: *when does someone count as performing a goal-directed (telic) non-basic action, rather than merely doing something else that could – under favourable conditions – be extended into it?*

The answer matters here. If I'm *building a house* (not merely stacking bricks), an interruption can be an obstruction; if I'm *writing a book* (and not merely *typing*), the activity is assessable by standards of progress and completion. Which action is underway affects which reasons rationalise what is happening and what forms of criticism or responsibility are appropriate.

Trying to draw the needed distinction runs straight into a special case of the *imperfective paradox* (Dowty 1977, 45). For accomplishment predicates, the progressive (or imperfective) form – English “is/was ϕ -ing” – can be true even when no ϕ -event is ever completed. After the walk sign flashes, a pedestrian may truly be said to be *crossing the street* even if the crossing is never completed, perhaps the pedestrian is picked up halfway, or struck and injured. In general, the progressive does not entail completion: “S is ϕ -ing” does not entail “S has ϕ -ed,” and does not even entail “S will ϕ ” (see, e.g., Dowty 1977; 1979, 134–35; Landman 1992, 2–5; Szabó 2004, 32–33).

The philosophical difficulty, however, is not merely this failure of entailment. A natural thought is that the progressive reports a proper part of a ϕ -event. But in the relevant cases there may be no completed ϕ -event at all for anything to be a part of. Still, progressive attributions are not truth-conditionally indiscriminate. Saying that the pedestrian *was crossing the street* isn't equivalent to saying that they (merely) *were walking in the street*. The puzzle, then, is how – *prior* to any culmination – the progressive fixes one telic description rather than another. What makes it the case, before the endpoint is reached, that the unfolding process counts as a street-crossing, rather than as some other ongoing activity which would have culminated in a crossing had it continued?

In the case of intentional action, the puzzle becomes sharper because the relevant “in progress” claims are themselves action-ascriptions. Consider two hikers, *Aria* and *Bron*. Both stand at the base of Ben Nevis. Both have formed the intention to walk to the summit. Both take their first step up the trail, where these are physically indistinguishable movements. Now stipulate that *Bron* is a systematic “false starter”: although he continues to intend to reach the summit, he reliably turns back after that first step because a recurrent panic takes over. We can also

stipulate that nothing external blocks his way: suppose that if Bron *did* continue beyond the first step where the panic characteristically kicks in, the normal continuation of that very route would take him to the summit. Even so, it is natural to say that (knowing this background), at the time of the first step, only Aria is already *walking to the summit*; Bron is doing something else: merely *taking a step while intending to reach the summit*. What grounds this difference?

The Aria/Bron case illustrates what I'll call the *Differentiation Puzzle*: two agents can share (i) a standing intention to ϕ and (ii) matching behaviour at time t and yet, (iii) differ in whether they are ϕ -ing at t . Standard factors – mere possession of the intention, bodily movement, even causal history up to t – can all be held fixed while progressive action-attribution diverges.

The Differentiation Puzzle is a special case of a more general *Guidance Problem* for non-basic action:

Guidance problem: What is it, in the actual world, for an end to be guiding an agent's present activity so that the activity counts as a stage of that very action?

To answer this in a satisfactory way, we need an account that is not merely future-looking (what would happen in “normal” continuations), but that identifies an actual-world feature of the agent's present condition that makes the progressive apt.

Because the Guidance Problem is general, it helps to keep a controlled diagnostic case in view. The Differentiation Puzzle plays that role. It isolates, at a single time t , the question whether the end is already guiding the agent's activity, while holding fixed the standing intention and the outward movement. In what follows, I therefore return repeatedly to Aria and Bron, given that failure on the Aria/Bron contrast is failure to answer the Guidance Problem; but, success on that contrast would still need to be shown to generalise.

The existing literature suggests three familiar strategy types. *Modal* accounts appeal to “inertia” or “normal continuation” conditions so as to tie “S is ϕ -ing” to what happens *if the present process continues without disruptive interference* (Dowty 1979, 149–50; Vlach 1981, 285f; Portner 1998, 773–77; Varasdi 2017, 313). In the philosophy of action, a closely related idea has been defended for process-claims of the form “x is ϕ -ing” (see Wolfson 2012). *Intentionalist* accounts are relevant even in the Aria/Bron set-up, but not because they can simply cite an intention to ϕ , given that the puzzle holds fixed that both agents have such an intention. Their distinctive ambition is to say what it is for an intention to be in execution, that is, operative in organizing and guiding present behaviour (for

instance, as an intention-in-action, or as a state that structures ongoing control). *Dispositional* accounts appeal to a propensity to continue in ways that would constitute ϕ -ing. Each of these approaches has at least some structural limitation; none satisfactorily specifies what, in the actual case, makes one set of continuations relevant rather than another, or what makes the present behaviour belong to ϕ -ing rather than merely coincide with it. That limitation matters because it is exactly what the Aria/Bron case exploits.

I'll argue that the rival accounts fail because they miss a crucial factor, put roughly for now, the *teleological organisation* of the agent's activity. Just as an intention can be present without being operative, likewise, behaviour (at any stage in a given action's trajectory towards a goal) can be goal-conducive without being goal-guided. What's needed is a condition that registers whether the agent is not merely oriented *in thought* toward an end, but is *already in a state of directed practical engagement toward it*.

The positive proposal to be defended holds that an agent is performing a non-basic action A at time t if and only if she has sufficient *teleological momentum* directed toward A 's end. Teleological momentum is a composite, direction-sensitive property with three components that are broadly analogous to the constituents of physical momentum: first, there's what I call *intentional mass* (how robust, stable, and priority-laden the relevant intention is), second, *practical velocity* (the *rate and direction* of the agent's current activity with respect to the end), and, thirdly, *guidance alignment* – viz., the degree to which the activity is controlled by, and responsive to, the end *as end*). When these components jointly exceed a context-sensitive threshold, progressive action-attribution is true: the agent is A -ing and not just doing something that could be extended into A -ing. The account thereby aims to supply an actual-world basis for the “normal continuations” appealed to by modal and dispositional theories.

The plan of the paper is as follows. Section 2 critiques modal, intentionalist, dispositionalist, and causalist approaches to our guiding problematic. Section 3 develops the teleological momentum framework. Section 4 applies it to the Aria/Bron case and related tests. Section 5 addresses objections. Section 6 concludes and outlines implications for moral responsibility, akrasia, and collective action.

2 | CRITIQUE OF RIVAL ACCOUNTS

Before developing the positive view, we should see why the main rival approaches fail to solve the Guidance Problem for non-basic actions. In this section, I consider the leading candidates in turn: modal, intentionalist, dispositional, and causalist accounts; it's shown that none adequately explains cases like Aria and Bron. Since

the Differentiation Puzzle is the Guidance Problem in miniature, that case will serve as a recurring stress-test.

2.1 | Modal Accounts

Modal accounts hold that progressive truth involves a counterfactual component: roughly, what is happening now must be of a kind that *would*, under appropriate continuation, issue in completion (see, e.g., Dowty 1977, 1979; Landman 1992, 25; Portner 1998, 773–77). Although modal accounts differ in their technical implementations, they converge on the same structural idea. Let's abstract away from the details and state that schematic idea explicitly.

Modal Schema. “S is ϕ -ing at t ” is true only if (and on standard versions: *iff*) ϕ is completed in the relevant *normal continuations* of what is happening at t .

In standard implementations, “normal continuations” are modelled by quantification over possible worlds (or world–event continuations) that match the actual world up to t and then develop as the ongoing process “would be expected” to develop, absent disruptive interference. Dowty's (1977, 1979) term is *inertia worlds* – viz., worlds whose future unfolds in a way maximally compatible with the past course of events (see Dowty 1979, 148).¹ The point is to filter out derailments (e.g., an earthquake may stop a house from being built without thereby making it false that the house *was being built* before the quake.)

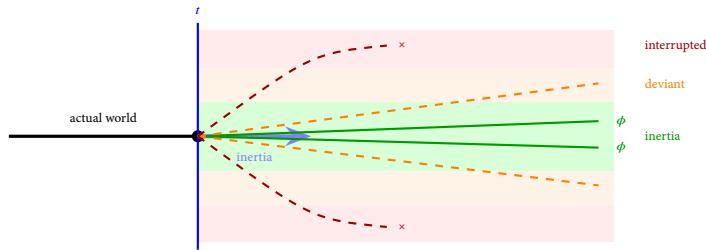


Figure 1: Modal semantics evaluates the progressive over *inertia worlds*: continuations maximally compatible with the ongoing process, filtering out interruptions.

Later refinements do not alter this basic explanatory ambition. Landman's

1. For some criticisms of Dowty's inertia-world model, see, e.g., Parsons (1990), Vlach (1981), and Engelberg (1999).

(1992, 25) influential revision distinguishes between interruptions that are *external* to an event's unfolding (for example, earthquakes, power outages, etc.) and factors that are *internal* to how the event itself develops. Portner's (1998, sec. 4.2) Kratzer-style treatment² recasts the same idea in more general modal terms: a *modal base* fixes the salient background circumstances, while an *ordering source* ranks continuations by how well they satisfy “no-interruption” conditions. The progressive is then evaluated not over all possible continuations, but over the best ones—those that preserve the relevant circumstances while deviating as little as possible from the ongoing process.

These refinements improve the formal articulation of the modal view, but they leave its core structure intact. In each case, progressive truth is tied to what happens in a privileged class of interruption-free continuations of the present process. The problem for present purposes is not that this picture is unmotivated in general. It is that, as a solution to the Guidance Problem for non-basic actions, it leaves a crucial parameter idle. Consider that the schema appeals to “normal continuations,” but it doesn’t actually tell us which continuations are supposed to count as normal ones in the agentive cases where we often most need illumination, namely, cases in which the relevant “interruptions” are *partly constituted by the agent’s own capacities, dispositions, and patterns of control*. That is exactly what the Aria/Bron case exploits.

Let’s return to Bron under the stipulations built into the set-up. At time t , Bron (i) has the intention to reach the summit of Ben Nevis and (ii) takes the same first step just the same as as Aria does. Yet (iii) Bron reliably turns back after that first step because a recurrent panic takes over (this happens every time). Also stipulate that there is no external impediment: if Bron continued beyond the first step, the ordinary continuation of that very route would take him to the summit. Now ask what the modal account says about the sentence:

(B) *Bron is walking to the summit at t .*

The modal framework can’t answer whether (B) is true or not without settling, one way or another, whether Bron’s panic is meant to count as an “interruption” to be screened off. But that’s precisely the issue on which the Differentiation Puzzle

2. Kratzer-style modal semantics treats modal expressions as sensitive to two parameters: a *modal base*, which specifies the relevant background facts (e.g. facts about the agent, environment, and ongoing process), and an *ordering source*, which ranks accessible worlds by how well they conform to certain norms or ideals (here, conditions of non-interruption or normal continuation). In applications to the progressive, this machinery is used to restrict evaluation to the *best* continuations of the present process, rather than to all possible ones. See Kratzer (1981) for the general framework, and Portner (1998) for its application to the progressive.

turns.

Here's the problem, then: If panic gets treated as an *interruption*, then the relevant inertia continuations are going to be worlds where the panic doesn't occur, Bron continues along the trail, and (since the route is clear) reaches the summit. On this construal, the modal schema makes (B) come out true. But if panic gets counted as part of Bron's normal course, then worlds where he turns back count among the best or most compatible continuations, and in those worlds he of course doesn't reach the summit. So on this construal, the modal schema makes (B) come out false.

Of course getting the second result would be better. But the point isn't that a modal theorist can't *choose* the second construal in order to recover the verdict that (B) is false. It's rather that as a purported explanation of why (B) is false, the choice would be left ungrounded. Thinking of things this way, the modal schema really just shifts the explanatory burden onto a classification – *viz.*, specifically, that of panic as interruption or as normal unfolding – which is exactly what we needed a theory of progressive action to *justify*. To put it in other words, *if modal accounts are offered as a solution to the Guidance Problem*, then the Differentiation Puzzle shows that they owe an account of what, in the actual world at t , fixes the relevant interruption profile, *viz.*, – *what is it* that makes it the case that the summit is *already governing Aria's activity but not governing Bron's*.

A defender of the modal approach might try to close the explanatory gap in familiar ways but none ultimately succeeds. One strategy here would be to just, in slogan form, “build Bron's psychology into the inertia profile.” One might say that the inertia worlds *hold fixed the agent's internal constitution and dispositions and not merely the external set-up*. Since Bron is disposed to panic, the best continuations, on this line, are going to include panic; therefore (B) is false.

While this looks like a principled thing for the modalist to say, such a manoeuvre, on closer inspection, really just restates the problem in different vocabulary. Consider that the question wasn't whether Bron's psychology matters but rather it was *how* it matters, more specifically, *which* internal goings-on count as interruptions rather than as part of the normal unfolding. Sudden cramps, fainting, heart attacks, etc. are internal events too, but they are naturally treated as *interruptions* of an action in progress (Portner 1998, 773–78; see Zhang 2023 for a review), not as constitutive of its “normal course.” So simply gesturing at “the agent's internal constitution” doesn't yet tell us why panic, in this case, belongs on one side of the line rather than the other. Any non-ad hoc criterion here is going to need to appeal to whether the relevant end is operative in guiding the present activity. But that's the very notion the modal approach was meant to analyse away.

The proponent of the modal approach might come back at this point and say:

“Let conversational context fix what counts as interruption, including which bits of a person’s psychology count and which bits don’t”. Does that work? Well it looks initially promising, given that the modalist can help themselves to Portner-style semantics which makes room for context-sensitivity via the choice of modal base and ordering source (see Portner 1998, 474–83). So perhaps, as this line of thought would go, depending on context, we sometimes ignore panic-like derailments and sometimes don’t. Here’s the problem, though. Even if that’s right as a piece of linguistic semantics, it still doesn’t answer the action-theoretic question this paper is pressing. The Differentiation Puzzle concerns what grounds an actual-world difference between Aria and Bron at t under a fixed action-description (*walking to the summit*) and, ex hypothesi, *in the same context*. If the modal story leaves it to context whether Bron’s panic is screened off, then it hasn’t explained *why* the action-ascription is apt for Aria and not for Bron in the case as described (where we’re holding fixed the evaluation context) but has rather, as it were, kicked the can further down the road.

2.2 | Intentionalist Accounts

The family of views we can call *intentionalist accounts* locates the source of an ongoing action’s directedness in the agent’s own practical attitude – typically, an intention, plan, or in Anscombean practical knowledge that *structures* what is happening now as something being done *for the sake of* a certain outcome.

Return to Aria and Bron. By stipulation, both (i) have the intention to walk to the summit and both (ii) take the same first step at time t . Yet we are inclined to say that only Aria is already *walking to the summit* at t , whereas Bron is, at t , merely *taking a step while intending to reach the summit*. Obviously if simply *having* the right intention were enough, Bron would get the same ‘gold star’ as Aria. So the intentionalist’s key move is to distinguish between (i) a standing intention that is merely present in the background and (ii) an intention that is (in some way) *in execution*, which is to say, an intention that’s not just an antecedent mental state, but is presently operative in organising and guiding behaviour.³

On the intentionalist picture, the Aria/Bron contrast is explained as follows. Aria’s summit-intention is already *in execution* at t : her first step is taken *as part of* the activity of going to the summit, guided by her plan, through such things as route-following, responsiveness to reasons bearing on the summit, readiness

3. This is a familiar thought in the philosophy of action, and it is naturally articulated within an intentionalist framework. Michael Thompson (2008) emphasises that an intentional activity genuinely in progress bears a *non-accidental* connection to its completion, such that when the activity is under way, it wouldn’t be a matter of luck if it reached its characteristic end (see Thompson 2008, Ch. 8). I return to this point in §4.

to correct course, and the like. Bron, on the other hand, possesses the summit-intention only as a standing commitment that doesn't yet guide his present conduct. Even if the bodily movement at t is indistinguishable, the intentionalist holds that Bron's step is not integrated (at least not well enough) into a summit-directed activity in the relevant way, either because some other practical orientation (e.g., sensitivity to potential panic inducers, panic responsiveness, etc.) is currently what's operative, or because the summit-directed plan isn't functioning strongly enough as the organizing principle of what he is doing, not compared at least to what else (after t) takes priority for him (turning away from the mountain). The difference at t on this view, is therefore not *whether* there is an intention to reach the summit (that's held fixed *ex hypothesi*), but whether that intention is playing the distinctive *guiding role* characteristic of intention in execution.

Intentionalist accounts in taking this line put their finger on something that modal accounts leave as a black box, which is the idea of an intention being *operative* rather than merely present. But as a solution to the Guidance Problem, they face two connected difficulties, one concerning the role of the world, the other which concerns the content of the appeal to “in execution.”

Here is the first problem. An intention can't, in and of itself, make a progressive non-basic action-description true if the world doesn't cooperate. In ordinary usage we retreat from “S was ϕ -ing” to “S was trying to ϕ ” when it turns out that ϕ was impossible given the actual conditions. Thus, someone who says “I am unlocking the door” but is holding the wrong key is not unlocking the door, even if her intention fully structures what she is doing (e.g., she inserts the key, turns it, and monitors for resistance, etc.). Had she been holding the right key, she *would* have been unlocking the door; but given that she's not, she is at most *trying* to do what she intends to do. Likewise, someone who says “I am writing in French” but does not know French is not writing in French, however sincerely the intention is held or however systematically it guides the activity. Falvey (2000, 24) himself insists that it is “surely a necessary condition” for “I am ϕ -ing” to be true that it be *possible* for the agent to ϕ successfully.⁴ You're not, no matter what you

4. Falvey however notes that there might be some complications here. In a note to his comment, he suggests the complication (that he sets aside) that a person might knowledgeably say that they're writing a book even if they *know* that they're going to die before they finish it (see Falvey 2000, 42, fn. 6). He suggests such cases suggest something potentially weaker, which is that to be in the process of doing something you have to not know it can't be completed, and so, it must at least be an epistemic possibility for you that someone else will complete it for you. In the dying author example, Falvey takes the weaker condition to be met so long as the dying author thinks someone else might finish it before he dies. I'm not sure what to think about the original case; for one thing, can one really know they'll die before they finish the book if they don't know that they *won't* later decide to declare (performatively) that “it's finished”, making it so through that speech act?

might think, counting to 100 if you start at 5 and start counting downwards. If it becomes clear that ϕ can't be achieved (e.g., because the agent lacks a crucial ability, a tool, or enabling conditions) then the claim that ϕ is in progress must be withdrawn (even if the intention persists).⁵ A lunatic can presumably *try* to do anything, as long as they have enough false beliefs about what it would take to bring it about. So intentionalist accounts must incorporate an objective enabling-condition constraint, and with it a tacit *ceteris paribus* clause, namely, that the intention's "guarantee" of success is only defeasible.

Second, and more directly relevant to Aria and Bron, once enabling conditions are admitted, the intentionalist still owes a non-circular account of what it is for an intention to be *in execution*. In the Aria/Bron case, the enabling conditions are already, by stipulation, in place: the route is clear, and if Bron did continue he'd reach the summit by ordinary continuation. So an account that appeals merely to *possibility* won't by itself exclude Bron. The intentionalist thus needs some further ingredient – not merely that success is possible, but that the summit-intention is *currently operative* in guiding behaviour. Yet if the story that's given at this juncture is just that Aria's intention is operative while Bron's isn't, then we haven't yet been told what actually grounds that difference at t in a way that doesn't presuppose the very action-ascription at issue. And saying that Bron's intention isn't "really" in execution because he will in fact turn back threatens to slide back into the modal pattern (a future-oriented diagnostic), while saying that it isn't in execution because panic "interrupts" threatens to reintroduce the interruption-classification problem in psychological dress. The explanatory burden is, after all, to say what, *now*, makes an intention count as executing rather than merely (idly) standing.

Granted, none of this shows that intentionalist accounts are hopeless but rather that, to solve the Guidance Problem in cases like Aria/Bron, they must do more than *merely* insist (correctly) that intentional action is teleologically organised from the agent's point of view. They must explain (without merely relabelling) what it is for an intention to be *in execution* in a way that integrates the agent's perspective with the actual-world conditions under which that perspective can (for practical knowledge views) amount to knowledge, and under which the intention can genuinely guide rather than merely accompany a goal-consistent movement. Anything short of that, and the intentionalist risks

5. There is some philosophical division about whether one can intend the impossible. Results from Buckwalter et al. (2021) suggest that folks are inclined to attribute intentions to do impossible things. That said, intentionalists under the Anscombean practical knowledge banner will be more restricted here: since practical knowledge is factive, one can't have practical knowledge that they are doing "X" in order to "Y" when Y can't be brought about by X; they'd have, at most, knowledge that they are doing X to *try* to bring about Y.

collapsing into the uncontroversial (but unilluminating!) truth that intentional actions are directed toward their ends; that statement is true, but it's not yet an answer to the Guidance Problem; the Differentiation Puzzle is simply the place where that inadequacy on the part of the intentionalist is hardest to ignore.

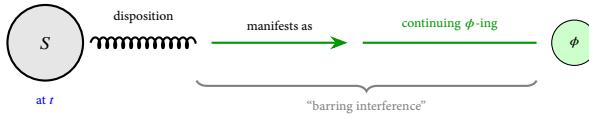
2.3 | Dispositional Accounts

Let's look at a different strategy, which is to ground progressive action-ascriptions in the agent's *dispositions* rather than in counterfactual "normal continuations" or in an intention's being "in execution." The guiding thought is familiar from both ordinary talk and philosophical theory, which is that to say that someone is, e.g., *walking to the summit* is, at least in part, to say that she is *set* to continue in the relevant way. She is not merely performing something like "a summit-compatible motion at an instant" but is in a state that would carry the process forward, were it allowed to unfold. Versions of this idea occur in the vicinity of Ryle's (1949, Ch. 2) dispositional analyses of intelligent action⁶ as well as in Kenny (2003, Ch. 8). We can put the dispositionalist proposal in a form parallel to the modal schema:

Dispositional Schema. "*S* is ϕ -ing at t " is true only if *S* has (at t) a suitable disposition such that, if it were to manifest under appropriate conditions, the manifestation would constitute *S*'s continuing ϕ -ing in a way that (barring interference) culminates in ϕ .

Understood in this way, dispositionalism aims to take the openness of the progressive seriously while at the same time keeping its truth grounded in actual features of the agent, viz., in what the agent is disposed to do. So far so good.

6. The overarching idea that what explains whether one's action is an intelligent action (of which intentional action is a species) as opposed to dumb movements appears throughout Ryle's (1946, see esp. 13-14) lecture on know-how and know-that, as well as in the *Concept of Mind* (1949, Ch. 2) where he notes (for instance): "But in looking beyond the performance itself, we are not trying to pry into some hidden counterpart performance enacted on the supposed secret stage of the agent's inner life. *We are considering his abilities and propensities of which this performance was an actualisation.* Our inquiry is not into causes (and a fortiori not into occult causes), but into capacities, skills, habits, liabilities and bents" (Ryle 1949, 32, my italics).



Dispositional Schema: “*S* is ϕ -ing at *t*” is true iff *S* has a disposition that would manifest as continuing ϕ -ing toward completion.

Figure 2: The Dispositional Schema: progressive truth grounded in a present disposition to continue.

The trouble though is that, once stated explicitly, the schema reveals the same kind of structural gap we found in modal accounts. Consider that dispositions are always dispositions *to do something under some conditions*.⁷ But the Aria/Bron case is designed to force the question: *which* conditions, and *which* disposition?

Return to the stipulated set-up. Aria and Bron are behaviourally matched at *t* (same first step up the same trail), and they share the standing intention to reach the summit. Yet at *t* only Aria is already *walking to the summit*, while (ex hypothesi habitual false starter) Bron is merely *taking a step while intending to reach the summit*. A dispositionalist needs to therefore identify some disposition that is present in Aria and absent in Bron and in a way that doesn't just smuggle in the conclusion (“Aria is summit-walking, Bron is not”). So what's that disposition, then?

The most obvious candidate is just a bare persistence disposition:

D1: a disposition to *keep going* up the trail if nothing interrupts.

But D1 is too weak. Bron is, in a perfectly ordinary sense, disposed to *keep going for a while* as long as nothing external prevents him. He's already taken a step up the trail after all. Even when we point out that Bron reliably turns back, it remains easy for D1 to come out true of him at *t*; his turning back might itself require some further change (fatigue, fear, discouragement, etc.) *none of which* need be present at the first step. In other words, a persistence disposition under a generic “no interruption” condition fails to track what we want it to track given that it threatens to classify any initial summit-directed motion as summit-walking, so long as it would continue absent a suitably described disturbance. And that just collapses exactly the distinction between mere goal-consistent motion and goal-guided engagement that the Differentiation Puzzle forces on us (and so leaves the

7. One very standard way of thinking about disposition is to ‘link’ them to conditionals, an idea initially developed by Carnap (1967) and which continues to hold sway in various areas of philosophy, including and especially in virtue epistemology (see, e.g., Sosa 2010, 466). Though whether any conditional analysis of dispositions can hold up is controversial – on this see Vetter (2011).

Guidance Problem unsolved.)

The dispositionalist can of course respond here by just *strengthening* the disposition so that it's not just a persistence disposition. So, as a next pass: a disposition to not merely “keep going”, but to “keep going *in the way characteristic of summit-walking*”, for instance:

D2: a disposition to continue in a manner that would *sustain* summit-directed progress – e.g., adjusting route, correcting deviations, responding to obstacles – so as to culminate in reaching the summit, so long as nothing interferes.

This is better, but only because it makes explicit what was tacit in D1: the relevant disposition is not just to persist, but to *maintain a goal-directed pattern under variation*. However, that is where the dispositionalist now faces a dilemma.

The first horn of the dispositionalist's dilemma is this: If D2 is specified only in outcome-based terms, as a disposition whose manifestation would *in fact* culminate in summit arrival when conditions are “right,” then we are right back to the modal problem in a kind of “dispositional dress”. The notion of “right” or “non-interfering” conditions would then be doing all the work. But in agentive cases (where what's at issue is the performing *by agents* of non-basic actions) the crucial candidates for “interference” aren't just external obstacles like avalanches and hurricanes (threats from hostage takers, police barricades, etc.) but also such things as panic, distractions, loss of resolve, shifts of practical priority, etc. The case was set up precisely so that some such internal development (Bron's characteristic panic) will occur. Now, whether *that* development counts as an “interruption” (something that stops an action already underway) or as part of the ordinary unfolding of what is happening (something that shows the action was never underway) is exactly what we're trying to explain, and so not something we can help ourselves to.

Here's the other horn: If D2 is specified not in outcome-based terms but instead in *guidance terms*, as a disposition to regulate one's behaviour by the end (by such things as monitoring and revising in the light of the goal), then it starts to look like a label for the very phenomenon at issue. For that disposition isn't a neutral, independently specifiable property as opposed to just a disposition to be guided by the end *as end*.⁸ So the dispositionalist, by specifying things this way, hasn't so much answered the Guidance Problem as renamed it.

The point can be sharpened by looking directly at the Aria/Bron contrast. Sup-

8. That is close to what the paper will later call *guidance alignment*.

pose we try to distinguish Aria and Bron by citing Bron's characteristic panic: Bron is disposed to turn back; Aria is not. Even if that's entirely true, it still doesn't explain *why* the progressive is false of Bron at t . There's familiar difference between (i) a disruption of an action already underway (cramps, a stroke, sudden external restraint) and (ii) a systematic lack of engagement such that the action was never underway in the first place (e.g., a deluded person intending to fly and moving their arms, or for that matter, a person dreaming they are doing so, barely twitching their muscles).

A dispositionalist who simply says "Bron is disposed to turn back, so he is not walking to the summit" treats a robust propensity to abandon as a standing defater of the progressive at t . But that then conflates two different roles a propensity can play. Some propensities *merely mark an action as fragile without showing it was never underway*. After all, a chef at high risk of a serious aneurysm is still making a soufflé; granted, collapsing from, say, an abdominal aortic aneurysm would thoroughly interrupt the soufflé making, but it wouldn't reveal the absence of ever having been soufflé-making before the aneurysm. The thought then is that if mere propensities to abandon were sufficient to falsify "S is ϕ -ing" at t , progressive truth would be retrospectively hostage to later derailments such that the very same earlier stretch of behaviour would count as ϕ -ing in cases where the propensity doesn't manifest and as not ϕ -ing in cases where it does. But that collapses the distinction we need between interrupting an ongoing ϕ -ing and never having begun ϕ -ing at all.

The dispositionalist might try to avoid this kind of "retrospective-veto" problem by *narrowing* the relevant disposition. What matters, the dispositionalist might say, isn't so much a tendency to turn back *later*, but a disposition whose *present* manifestation is already in train (perhaps: being in a "summit-walking mode", as opposed to merely producing summit-compatible motion). But once the view gets refined in that way, it then becomes unclear what remains distinctively *dispositional* about the explanation. For the contrast between a "mode" that is presently sustaining a summit-directed pattern and a condition that is not, just is a contrast in *current guidance*; in the former case, the end is already doing organising work in the agent's on-going control of what she's doing; in the latter it is not. That's not a further *dispositional* fact that can be cited *in place of* guidance but is basically the Guidance Problem stated in different words.

In sum, then, dispositional considerations are clearly part of the story, but the dispositionalist faces a dilemma, which we've seen looks like the following in short: Characterise dispositions thinly (as a tendency to continue compatibly with the goal) and the view is then too weak in the sense that it can't exclude Bron; but characterise them thickly (e.g., as dispositions to be guided by the end) and the

view builds the explanandum into the explanans.

2.4 | Causalist Accounts

Let's consider one more strategy type which we find in the causalist tradition in the philosophy of action. On a Davidsonian (1970) picture, an action is intentional under a description if it is caused, in the right way, by the agent's reasons (beliefs, desires, intentions) that rationalise it under that description.⁹ That framework is *prima facie* attractive here because it promises to ground action-ascriptions in something firmly actual-world, namely, a causal history.

Applied to the present topic, the causalist will be tempted to say the following: that Aria's first step is a step *in walking to the summit* because it is appropriately caused by her summit-directed intention (or by the belief–desire complex that constitutes her reason for walking), whereas Bron's first step isn't appropriately so caused.

Of course, the immediate difficulty is that the Aria/Bron case holds fixed precisely what the naïve causalist wants to vary. By stipulation, Bron has the summit-intention at t . His first step is a deliberate step up the trail to Ben Nevis and not some sort of accidental flailing about. So it's not obvious, on the face of it, why the summit-intention can't be among the causes of that step in Bron's case as well. (We can even suppose akrasia would kick in if his initial intention wasn't impressive enough by his lights, as summing it is.) So even when it's pointed out that Bron's panic will later undermine execution, that doesn't by itself show that the summit-intention is causally idle *at the very first step* where we are evaluating him.¹⁰

So the causalist is going to need more than the claim (in distinguishing Aria and Bron at t) that the relevant intention is among the causes. The standard response at this juncture will be a familiar one: intentional action requires the *right kind of causal connection*, not just any. Otherwise we get deviant causal chains where the intention causes the behaviour, but not in a way that constitutes the agent's acting for that reason.

That's a familiar move (as far as a thesis about intentional action goes), but in the present dialectic, it yields an explanation only when the theory supplies conditions for *right way* that are informative in the relevant range of cases. Now consider what happens if the causalist does not supply such conditions. What we're

9. For a recent more ecumenical way of thinking about causation in connection in non-basic action, see Kelley (2022), who allows non-basic actions to be performed through productive and not only constitutive causal means.

10. Indeed, many intentions do their work without constant conscious focus, and they can influence immediate behaviour even when they will not be sustained.

offered is basically: “Aria’s step is caused by her summit-intention in the right way. Bron’s step is not.” Given the set-up, those claims don’t tell us what distinguishes the cases at t . The original question concerns what grounds the difference between an unfolding summit-walk and mere motion that could be extended into one. A bare invocation of “right way” doesn’t yet provide an actual ground for that difference.

Notice that the pressure becomes sharper once we attend to the *temporality* of non-basic action. “Walking to the summit” is an extended process whose stages must be integrated under the action-description. So a causalist account adequate to such cases needs to therefore distinguish at least two kinds of causal roles an intention can play: There is (i) *initiatory causation*, where the intention helps bring about the first step; and (ii) *sustaining causation* where the intention continues to shape what happens as the process unfolds, where this will involve (e.g.) selecting means, correcting deviations, responding to obstacles, and then stopping when the end is reached or becomes impossible.

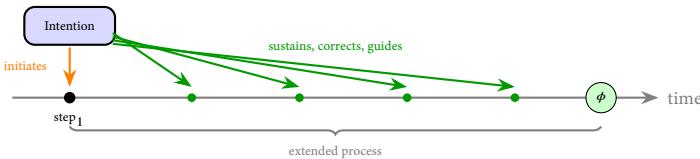


Figure 3: Initiatory vs. sustaining causation. The intention *initiates* the first step (orange) but must also *sustain* the process throughout (green)—selecting means, correcting deviations, guiding to completion.

This distinction type will be (by analogy) familiar from epistemological discussions of causation in connection with the epistemic basing relation, where the difference between a good reason’s merely causing the initial uptake of a belief and causing/sustaining that belief is thought to matter for whether the belief continues to be based on the good reason that initially caused it, at any given time that follows the initial formation (see, e.g., Bondy (2016) for an overview).

The difference (initiatory causation and sustaining causation) matters for our purposes because it is consistent with the case description that Bron’s summit-intention plays some initiatory causal role at t while failing to play the sustaining role that would make the activity a genuine summit-directed process. If so, then the causalist can’t identify summit-walking with mere causal initiation by the intention. The causalist must instead appeal to some structure of *ongoing control*.¹¹

11. That is broadly analogous to how proponents of causal-sustaining theories of the basing relation in

But (and here's the catch) once the causalist appeals to such ongoing control, the same question returns in a more determinate form: what, *at time t*, grounds the claim that the intention is exerting the relevant kind of control? It *can't* be enough to say that the intention *will continue to exert it* because that just pushes the explanation into the future and makes progressive truth depend on what happens later. But nor is it going to be enough to just say that the intention *would* exert it in "normal" circumstances, as that move then reintroduces the normalcy parameter that modal and dispositional accounts left unexplained. The causalist needs an actual-world marker of control at *t* that distinguishes a merely triggered movement from a movement under guidance.

The causalist's 'hammer' for most problems across different areas in philosophy is the hammer of its non-deviance proviso.¹² So, one might think the present juncture is exactly where causalist resources about "deviance" would potentially help: perhaps Bron's case involves a kind of deviance (what he does is certainly weird). But to call what's going on with Bron's initial movements deviant (specifically, to call the causal connection between his intention and his taking the step up the hill "deviant") is misleading. The puzzling feature isn't that Bron's intention causes his movement by some freak route. It's that, even with perfectly ordinary causal relations between intention and movement, *the movement can still fail to belong to the extended action*.

At that point the best causalist move looks like it's just to concede that what's missing is an account of the intention's *guiding* role – viz., its role as a control state that structures the unfolding of behaviour in response to reasons and information about the end. Put negatively: the causalist needs some principled way to distinguish (i) mere causation by an intention from (ii) causation that constitutes rationally intelligible, goal-directed execution.

That is precisely the gap the subsequent positive theory to be developed in the next section is meant to fill. The *Teleological Momentum Theory* which I'll turn now to setting out and defending, can be read, if one wants, as offering the causalist the "missing middle", viz., a way of articulating what it is for an intention's causal influence to take the form of ongoing guidance rather than mere initiation. The key idea (put roughly for now) is that what matters for progressive action-ascription isn't just that the intention causally precedes the behaviour, but that the behaviour is embedded in a present structure of directed practical engagement

epistemology think your belief is based on your good reason you have for it only if that good reason continues to influence your maintenance of the belief over time, which includes, e.g., responsiveness to evidence, reasons to inquire further, etc.

12. For discussion on this point, see Sosa (2015, Ch. 1).

whose central component is (what I'll call) *guidance alignment*.

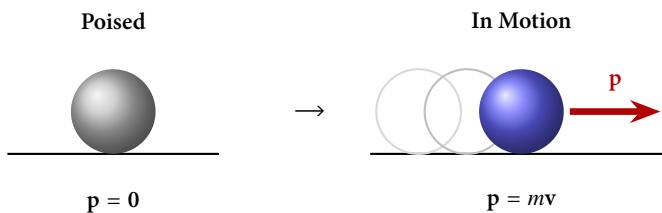
So the moral of the causalist detour we've just navigated isn't that causation is irrelevant (it is relevant!) but that causation without a theory of guidance, rather than just a gesture to guidance, is simply too coarse-grained for the present problem.

3 | TELEOLOGICAL MOMENTUM

Having identified the need for an actual-world, structural feature that distinguishes genuine ongoing action from mere intention-plus-motion, I propose a new framework whose central idea is that performing a non-basic action at a given time requires having sufficient *teleological momentum* directed toward the action's end. In this section, I introduce and elaborate this idea, and in doing so will use an extended analogy with physical momentum to serve as an explanatory scaffold.

3.1 | The Concept of Teleological Momentum

Consider how, in physics, we explain the difference between an object that is merely poised to move and one that is *actually in motion*. A rolling ball possesses *momentum*, a product of its *mass* and *velocity*, which gives the ball a tendency to keep going in its current direction.



A heavy ball rolling slowly and a light ball rolling quickly might have equal momentum. Momentum is also *directional*. A ball rolling east is not thereby disposed to roll north. To put it in the simplest terms: momentum has a *magnitude* (how much) and a *direction* (which way), and it's the *combination* that matters for what counts as "already underway".

The proposal I'll begin to lay out now takes as a starting point that something structurally analogous is at work in the case of *non-basic goal-directed action*. What we need, to solve the Differentiation Puzzle, is a present, actual-world feature that distinguishes mere goal-*compatible* movement from genuine goal-*guided* activity, something in virtue of which it would not be a matter of luck if the agent reached the relevant end, so long as nothing interferes. That feature is what I'll

call *teleological momentum*.

Rather than building up to the view, I'll state it up front and then proceed to unpack it, show how it explains non-basic actions in progress, and demonstrate how it resolves the problems at hand.

Formally, we can define teleological momentum toward an end E as follows:

Teleological Momentum: Teleological momentum $TM(S, E, t)$ of an agent S at time t is a composite vector property of S , representing S 's *ongoing directedness* toward bringing about E . It has three components:

- i). **Intentional Mass (M_I)**: the robustness, stability, and priority of S 's intention to bring about E .
- ii). **Practical Velocity (V_P)**: the rate of S 's current activity directed toward E , together with its orientation (toward or away from E).
- iii) **Guidance Alignment (G_A)**: the degree to which S 's present behavior is *controlled by* and *responsive to* E as the goal.

I'll explain each component in detail in the next section. For now, though, two clarifications will be useful. First, calling teleological momentum a “composite vector property” is not to posit a literal vector in physical space. The idea is just to register a structural point about *how the relevant property behaves*. In the present setting, the “direction” of a momentum attribution is fixed by the *parameter* E : we are always asking about the agent's practical state *as evaluated relative to that end*. But, as with physical vectors, the crucial feature is that *more activity isn't automatically more progress*. What matters is *aligned contribution*. A convenient way to think of this is by analogy with *projection*. If we imagine an agent's ongoing activity as having a certain practical orientation, then what contributes to momentum *toward E* is the component of that orientation that “points” toward E . If the agent's activity is largely orthogonal to E (busy, but irrelevant), the relevant component can be close to zero. If it's oriented against E (the agent is undoing what would be needed for E), the relevant component can be negative *with respect to E*. The vector talk is meant to capture exactly this kind of sensitivity to *alignment* – viz., the same amount of “effort” can count very differently depending on how it's organised in relation to the end.

Second, the three components (i.e., intentional mass, practical velocity, and guidance alignment) aren't intended as three separable “tests” whose independent satisfaction settles whether S is already doing A . Rather, they are introduced as

dimensions along which an agent can be more or less *under way* toward an end, and they are meant to *combine into a single assessment of directed engagement*. In particular, *guidance alignment* isn't merely one contributing factor among others but rather it's what makes the difference between (i) intention plus motion that would reach the end only by good fortune, and (ii) intention plus motion in which the end is already doing *organising work* in the agent's present control of what the agent is doing. That's why an agent could potentially have a firm intention and substantial activity motivated by that intention, and yet have little teleological momentum toward E , if the activity isn't being steered by the goal in the relevant way. Conversely, an agent could be moving slowly or doing comparatively little at a given moment, and yet have significant teleological momentum, if what they're doing is tightly integrated under the guidance of the end. Teleological momentum is meant to mark, in an actual-world way, the difference between mere goal-consistent motion and goal-guided engagement.

With that notion in place, we can state a general criterion for progressive action-ascriptions. Let E_A be the characteristic end of an action-type A (for example, for *walking to the summit*, E_A is *reaching the summit*). Then, the idea is this:

Teleological Momentum Thesis (TMT): An agent S is performing non-basic action-type A at time t if and only if S has teleological momentum toward E_A at t that meets or exceeds a threshold T .

Here the threshold reflects a familiar feature of ordinary action-ascriptions: there are going to be borderline cases in which it is indeterminate whether someone is *really* doing A or is merely preparing or acting in ways that *would* count as A only given further uptake (e.g., is Sergio García's initial waggle of the golf club, as he stands up to the ball, *part of the swing*, or is something that precedes it and initiates it?) Different action-types plausibly set different standards for this. Being such that one is (actually) *writing a doctoral dissertation* typically requires more sustained and better-integrated engagement than being *walking to the corner store*. Conversational purposes can also raise or lower the relevant bar in predictable ways (for example, when it's high stakes whether someone has genuinely begun – say, for funding to be released – rather than merely set the stage or prepared to begin). The teleological momentum view doesn't try to eliminate such vagueness; it just claims that, where a progressive action-ascription is true, its truth is going to be grounded in the agent's having *enough* teleological momentum, at that time, with respect to the end characteristic of the action. That's what it is to be performing a non-basic action.

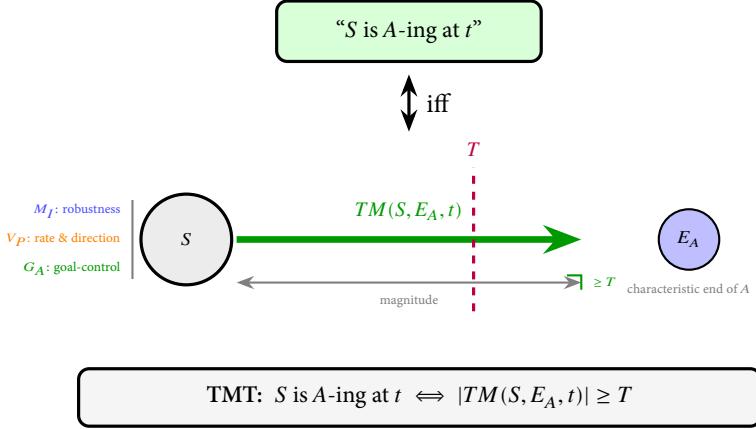


Figure 4: The Teleological Momentum Thesis. S is performing non-basic action A at t if and only if S has teleological momentum toward the characteristic end E_A that meets or exceeds threshold T . The three components combine into a single composite vector property representing S 's ongoing directedness.

3.2 | The Three Components Elaborated

This has all been somewhat high-altitude so far; let's now look squarely at each component of teleological momentum in turn and clarify its distinctive role.

3.2.1 Intentional Mass (M_I)

Intentional mass captures the *weight* of the agent's intention within their practical psychology. Not all intentions are created equal; some intentions are heftier, so to speak, and thus contribute more to momentum, others less so. For example, a well known result in empirical psychology is that despite the significant health benefits associated with eating healthily, the vast majority of people who form intentions to change their diet fail to do so (van Osch et al. 2009).¹³ Not all of our intentions are flimsy like this. An intention with high M_I will have at least three important features: it is (a) *robust*, (b) *stable*, and (c) of *high priority* for the agent. Regarding robustness: one factor that feeds into the 'mass' of an agent's intention concerns whether they would maintain the intention across a range of contrary pressures or changing circumstances. A robust intention doesn't evaporate at the first sign of difficulty or temptation; for example, like those would-be dieters in the van Osch

13. See Sheeran et al. (2005) for discussion of how implementation intentions (intentions to implement means to a given goal) are sensitive to the strength of underlying goal intentions.

et al. (2009) study: for some who intended to change their diet, the presence of a tempting treat offered to them was enough. Robustness can be distinguished from *stability*, which also matters for what I'm calling intentional mass. When an intention is *stable*, it persists over time and across distractions or delays. I might, after reading Frege's "Über Sinn und Bedeutung" in graduate school, briefly form the intention to read everything he ever wrote, only to drift on to something else (after taking a new class) despite neither encountering or considering any reason to abandon the plan. This contrasts sharply with the stability of Michael Dummett's intention, early in his career, to read everything Frege wrote, which was actually sustained and carried through to completion. Thirdly, high-mass intentions will have for the agent *high priority* – viz., a high place in the agent's hierarchy of ends. The agent is willing to allocate substantial time, energy, and resources to it (what the high priority intention is an intention to do), possibly foregoing other competing projects.

Let's introduce two new characters, Alice and Beth, each who also has an intention to summit Ben Nevis. Alice forms the intention casually, perhaps after seeing a photo in a Scottish tourism guide. But this intention doesn't lead to very many downstream *implementation intentions* – viz., intentions to bring about determinate means toward one's (goal) intention (see Gollwitzer 1999; Gollwitzer and Sheeran 2006; Sheeran et al. 2005). Beth, on the other hand, has made this climb a central goal in that she's rearranged her schedule, packed supplies, informed her friends that she'll be unreachable for a couple of days, and so on. While Alice's intention doesn't play much if any role in how she does other things in her life, Beth's intention is the organising principle of her current plans. If both are at the base of Ben Nevis intending to climb, Beth already has more *momentum* toward success by virtue of her firm commitment, and she has this on the present view even before any physical steps are taken.¹⁴

3.2.2 Practical Velocity (V_P)

Practical velocity represents the *rate and directionality* of the agent's current activity with respect to the end E . This the speed of *practical progress* (rather than just

14. This notion of intentional mass connects to Michael Bratman's (1987) insight that intentions play a characteristic functional role in *planning* agency. An intention to ϕ is a commitment to ϕ that resists reconsideration and coordinates one's activities (Bratman 1987). The thought, for the present purposes, is that the more an intention is playing that role – resisting reconsideration, structuring the agent's decisions and attention – the more "mass" the intention has. On TMT, M_I quantifies at least roughly how much "oomph" or sustaining force the intention provides toward the end. So a higher intentional mass means the agent's intentional stance toward E is weighty and resilient in an such a way that is going to contribute strongly to the overall momentum.

simply the speed of one's actions while one thinks they are bringing one closer to E). It has two sub-aspects: *magnitude* and *direction*.

Let's look first at magnitude (or speed). If someone has an intention to bring about some end E , we might ask how *fast* is the agent moving or progressing toward E ? This could be measured in physical terms (miles per hour toward the summit) or in terms of task completion (tasks done per hour, pages written per day, if one is writing a novel for National Novel Writing Month (NaNoWriMo) etc.), depending on the action. Importantly, even if the agent is temporarily stationary or moving slowly, they might still have a nonzero practical speed if what they're doing is meaningfully advancing the project (e.g. carefully planning the route or taking a brief rest might still contribute to eventual completion). Arguably, Kant had non-zero practical speed (towards writing the *Critique of Pure Reason*) when he entered in 1770 what scholars often call his “silent decade” (roughly 1770–1781); while he had during those 11 years nothing much to show for himself in terms of word count, he was in deep reflection about the foundations of his views on reason and metaphysics (and, in this way, not unlike Beth who has already some momentum towards her hike, but not Alice, prior to any steps being taken.)

On the other hand, rapid physical motion might not translate to practical progress if it's aimless or misdirected. When it comes to *direction* (as a component of practical velocity) we can ask: is the activity oriented *toward* the goal, or is it taking the agent away from it, or perhaps just orthogonal/sideways to it? Direction here means the following: given the content of the goal, are the agent's actions such that they reduce the remaining distance (literal or metaphorical) to the goal? For instance, walking in circles at base camp has high physical speed but zero practical velocity toward the summit. Running downhill (away from the summit) would be negative practical velocity with respect to the summit goal.

V_P thus captures how dynamically and effectively the agent is moving *now* toward completing the action. It contributes to teleological momentum specifically by *increasing the momentum's magnitude when the agent is actively doing things that make a dent in the task*. So if practical velocity drops to zero (i.e., if the agent is idle with respect to the project, neither planning nor acting) or becomes negative (the agent's activity undoes or counteracts progress, as in backtracking on the trail), then momentum toward E diminishes accordingly.

3.2.3 Guidance Alignment (G_A)

Guidance alignment is perhaps the most crucial and distinctive component of teleological momentum. It measures the extent to which the agent's *present, ongoing activity* is *controlled* by their representation of the end E . Put another way, G_A

assesses the *guidance relationship* between the goal and the current action. Even if one has a firm intention (high M_I) and is exerting effort (some V_P), *if that effort is not guided by the intention, the agent lacks teleological orientation toward the goal*.

The notion of *guidance alignment* borrows its governing idea from control theory (Phillips and Parr 2011; Bechhoefer 2005; Glad and Ljung 2018), which studies dynamical systems that evolve over time and the mechanisms by which they are steered toward a target state. In the central (and most useful) cases, guidance is a matter of *closed-loop* control where the system doesn't just move in a target-conducive direction, but continually uses feedback about its current state to regulate what happens next. A target (or reference value) is represented; the system tracks the *error* between its present trajectory and that reference; and then, control inputs are adjusted so as to reduce that error over time.

What matters, therefore, is not just that the current motion is compatible with reaching the target, but that the ongoing behaviour is *counterfactually sensitive* to deviations under the agent's ordinary operating conditions in the setting (e.g., liable to correct course when disturbances arise, when the environment shifts, or when the agent's own performance begins to drift, etc.) Guidance alignment, in this sense, is the degree to which an end functions as the *control variable* of the agent's present activity, given the agent's standing control profile in the setting, in such a way that the end is intended, served by what happens next, and (crucially) plays an actual organising role in the feedback structure that governs the unfolding dynamical process.

When the *guidance relationship* between the goal and the current action is well aligned, the agent's current behaviour would adjust if the end were to change, or in response to information about the end (we'll return to this in (§4)) under conditions that fall within the agent's ordinary range in the setting. So, for instance, if E were suddenly achieved or for that matter rendered impossible, the agent's actions would change accordingly. High G_A means, in short, *what the agent is doing right now depends on the goal*. (If Beth discovered that the trail she's on will not lead to the summit, she would alter her course immediately; Alice, by contrast, might not react to such news (if her current walking is governed by a different aim), indicating for her *low sensitivity*.)

A second core feature – beyond the above kind of sensitivity – that matters for guidance alignment concerns *corrective feedback and error monitoring*. In the good case (i.e., high G_A), the agent is organised to *detect deviations from the path to E and to correct them*. This is a hallmark of guidance in the sense that when something goes wrong relative to the goal, a guided agent compensates within the ordinary range for the agent in the setting. (For instance, Beth stumbling or realising she's slightly off trail prompts an immediate correction *because* she is

trying to get to the summit.)

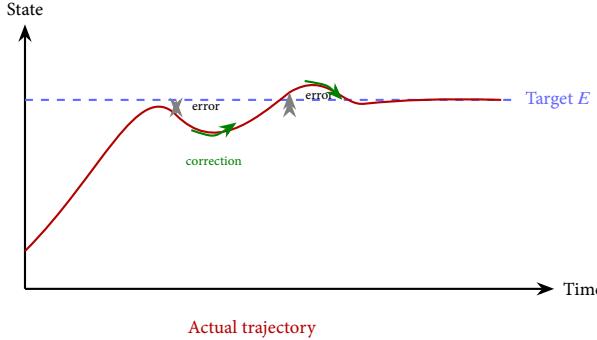


Figure 5: Guided action tracks the error between current trajectory and the target, adjusting control inputs to reduce deviation over time.

In sum, guidance alignment distinguishes merely *having* a goal while pursuing it from *actively pursuing* it in the sense that it ensures – in the terms of our framework here – that the other components (intentional mass and practical velocity) are *properly directed*. Without G_A , intentional mass and practical velocity toward E can be present yet the agent isn't *really* doing the action. High G_A thus means the agent's condition right now is appropriately described as “working toward E ”. As we'll shortly see, outward sameness at a time can mask a deeper difference in whether the agent's current condition actually supports guidance alignment for the task at hand in the relevant agent-and-setting-relative sense.

4 | SOLVING THE DIFFERENTIATION PUZZLE

Having now fixed the three components of teleological momentum, we can now apply the TMT thesis in earnest – with the core operative idea being that an agent S is performing non-basic action A at time t if and only if S 's teleological momentum toward A 's characteristic end meets a context-sensitive threshold. This section demonstrates how the framework solves the Differentiation Puzzle by identifying the actual-world basis for the divergence between Aria and Bron.

4.1 | Guidance Alignment and Constitutive Resilience

At this point in the dialectic, the TM framework has a very specific job to do; at the outset, it's worth registering that, within the framework, *intentional mass* and *practical velocity* don't have what it takes, just by themselves, to resolve the Aria/Bron contrast, without appeal to guidance alignment. The case is constructed so that (i)

$M_I(t)$ can be held fixed across the agents (grant Bron a robust, stable, high-priority intention *ex ante*), and (ii) $V_P(t)$ is matched at the same snapshot (they take the same first step, with the same immediate rate and direction of progress). So, if TMT is to satisfactorily explain why only Aria is already *walking to the summit* at t , an explanatory burden falls on the remaining component, *guidance alignment*. The present section isn't going to *reintroduce* G_A from scratch, but I am going to develop it further, and in some more technical detail, in order to make explicit some features of G_A that help do important discriminating work in early-stage cases like Aria/Bron.

Recall now the earlier control-theoretic gloss on guidance alignment: $G_A(t)$ is high to the extent that the end E_A functions as a *control variable* of the agent's unfolding activity, so that the agent's present policy is *closed-loop* with respect to the goal (that is, ongoing behaviour is regulated by feedback about its deviation from E_A , rather than merely initiated and left to run open-loop). So, concretely, the agent (when $G_A(t)$ is high) represents (perhaps implicitly¹⁵) a target state associated with E_A ; monitors her evolving state relative to that target; and then modulates what she does in ways that are systematically responsive to deviations, where this will involve, e.g., adjusting route, pace, means, attention, and subplans, etc., so as to reduce goal-relative error as the process unfolds. (Think again here of the heat seeking missile that adjusts as the target changes its spatial orientation).

This control-theoretic picture already contains a modal dimension: closed-loop control is partly characterised by what the system would do under disturbances, and so whether it *would* correct, stabilise, and continue. The present point is that, to do the work the TM framework needs in the Aria/Bron case, *high* $G_A(t)$ must involve more than a momentary match in, as it were, "looking goal-directed"; it needs to involve what we can call *constitutive resilience*. The thought here is that an agent's present control organisation needs to *be such that it can sustain* the goal-guided process across the ordinary perturbations characteristic of the action-type for the agent in the setting.

Let's make the idea of constitutive resilience more precise. Let \mathcal{P}_A be the ordinary perturbation profile for A in context. This will include (e.g.,) minor route deviations, small obstacles, transient distractions, ordinary fatigue shifts, and the

15. With this caveat I take it that the view here accommodates cases like Paul's distracted driver, driving home on autopilot (see Paul (2009, 5)). Paul's point with the case is that intentional action, in so far as it requires non-observational knowledge, doesn't require *occurrent belief*, but mere tacit or implicit belief would suffice; in a similar vein, I take it that examples like Paul's driver show (on TMT) that a distracted driver could also be truly *driving home* at a time prior to arriving, and even while not hosting *occurrent beliefs* about what they are doing. I take up this point in some more detail in the Objections and Replies.

like. Treat \mathcal{P}_A as a control-theoretic parameter: the bounded class of disturbances relative to which stability and corrective responsiveness are assessed. The profile can therefore include predictable idiosyncratic disturbances that recur for a given agent in the relevant setting. Bron's panic qualifies on the stipulated case description, since it's systematic and immediate (whereas a rare panic episode would not; it would count as an interruption). Then the relevant thought can be stated as follows:

Constitutive Resilience (CR): An agent's activity at t is constitutively resilient with respect to E_A only if, given the agent's *present* control organisation, the E_A -tracking process remains stable across the perturbations in \mathcal{P}_A : deviations would, in the ordinary range, trigger compensatory adjustments that preserve (and typically restore) a trajectory of progress toward E_A .

To be clear, CR shouldn't be read as a fourth component of teleological momentum. We should think of it rather as a constraint on what counts as *high* guidance alignment. Put in control-theoretic terms, just consider that a goal can enter an agent's activity in different ways. In some cases it merely figures in the agent's practical orientation at a moment (a fighter pilot, say, who momentarily focuses on flying faster than the Ice Man, rather than staying on mission). In others it's embedded in the organisation that governs how the activity unfolds over time. In the latter case, information about the agent's current relation to the goal feeds into the regulation of what she does next, so that ordinary deviations elicit compensatory adjustments rather than loss of direction. Guidance alignment is high only (for an agent and a goal-directed task) when this goal-sensitive control organisation has sufficient stability to sustain directed activity across the ordinary disturbances characteristic of the action-type for the agent in the setting. This is what allows $G_A(t)$ to be a genuine actual-world basis for distinguishing (i) being underway in an extended action from (ii) merely producing an initial movement that is goal-compatible.

Making CR explicit allows us to be clearer about what's issue at the very outset of an extended action. The question at that stage is how an agent can already count as performing a non-basic action when only the first step has been taken and very little progress has yet been made. On the TM framework, this turns on the *overall magnitude* of the agent's teleological momentum at a time, which (per TMT) gets represented as the result of some aggregation of our three components:

$$|TM(S, E_A, t)| = F(M_I(t), |V_P(t)|, G_A(t)),$$

where the only assumptions about F that matter here are reasonably modest ones: first, that momentum increases as any of the three components increases; and second, that if one component is close to zero, it strongly limits the contribution of the others (the intuitive “weak-link” feature of composite directedness). A simple toy model that makes this structure vivid is multiplicative:

$$|TM(S, E_A, t)| \approx M_I(t) \cdot |V_P(t)| \cdot G_A(t).$$

Nothing substantive hangs on this exact form; its role is just to make explicit how the components jointly constrain overall momentum.

Now, consider what this implies at the *beginning* of an extended action whose trajectory unfolds nontrivially over time. This will include the initial stages of many of our non-basic actions where many distinct means need to be implemented en route to the goal. At such early moments – indeed, paradigmatically, at the very first step – practical velocity $|V_P(t)|$ is necessarily *small*: little progress has yet accrued, and so velocity can’t by itself account for a substantial amount of teleological momentum. Here “small” is to be understood relative to the action-type: $|V_P(t)|$ is evaluated on a task-relative scale, so first-step values are low but not arbitrarily close to zero. Even where intentional mass is significant, whether the action already counts as underway will therefore depend on how the remaining component (guidance alignment) contributes at that time. So if, as the TM view holds, progressive truth requires teleological momentum to meet or exceed some threshold T appropriate to the action-type, then at an early time t_0 guidance alignment must be strong enough to compensate for the limited contribution made by initial progress:

$$|TM(S, E_A, t_0)| \geq T \quad \Rightarrow \quad G_A(t_0) \gtrsim \frac{T}{M_I(t_0) \cdot |V_P(t_0)|}$$

The precise inequality is not important. What matters is the qualitative result it makes explicit, which is just that when practical velocity is small, meeting the threshold requires guidance alignment to be comparatively strong. This yields a more precise version of what was earlier an informal claim, which is this: at the earliest stages of an extended action, teleological momentum depends disproportionately on guidance alignment, not because guidance alignment always matters most, but because practical velocity has not yet had time to accumulate. In short, at the outset of an extended action, what makes it true that the action is already underway is principally that the goal-tracking control mode is already *online* in a robust way.

Now let’s return to the Aria/Bron case under the stipulations that fix $M_I(t)$ and match $V_P(t)$ at the first step. If the TMT framework is to discriminate the

cases at t , it needs to locate a difference in $G_A(t)$, and the preceding points show what that difference amounts to. Aria's first step is embedded in a robust closed-loop organisation aimed at E_A (the summit) - in that ordinary deviations would register as error relative to the goal and would trigger compensatory corrections, sustaining summit-directed progress across the normal perturbations in \mathcal{P}_A . That is precisely constitutive resilience, and it is what high guidance alignment amounts to in the present setting.

Bron lacks this. Although his initial movement matches Aria's, his practical organisation is such that a predictable internal disturbance (panic) reliably seizes control immediately after the first step. In control-theoretic terms, Bron's summit-tracking organisation lacks adequate robustness margin; it's poised to switch to an avoidance policy under conditions that fall within the normal range for him. The fact that such panic might well be statistically atypical among summit-walkers generally is irrelevant to this robustness assessment. The summit does not function as a stably operative control variable of the unfolding activity. That is not best modelled as an "external interruption" of an already-guided summit-walk; it is a manifestation of the absence of the constitutive resilience required for high $G_A(t)$ in the first place.

This explains why, *even at the first step*, where $|V_P(t)|$ is too small to do much discriminating work, Aria's teleological momentum can exceed the threshold for *walking to the summit* while Bron's can't. Aria has high $G_A(t)$ (and thus constitutive resilience), and therefore sufficient teleological momentum at t . Bron's $G_A(t)$ is low, since the summit doesn't robustly govern his activity, and given the early-stage smallness of $|V_P(t)|$, his momentum falls below threshold. This provides the TM-based ground for the verdict that Aria is already walking to the summit at t , whereas Bron is not.

4.2 | Resolving the Open Questions

We've got the main answer to our guiding question on the table now; but there remain some residual questions that some of the rival accounts left open along the way (§2), and so now let's see how the TMT framework can offer some principled answers to them. One such question concerns the status of panic as interruption or ordinary unfolding. Recall modal and dispositional accounts needed to decide whether Bron's panic counts as an "interruption" to be screened off. The TM view supplies a criterion, namely, that *internal developments count as interruptions when they overcome a goal-directed control structure that's already in place*; they count as ordinary unfoldings, and thus revelations that the action was never underway, when the agent's constitution is such that goal-directed activity could not be sustained. And our view was that Bron's panic falls into the second category

whereby it doesn't interrupt a summit-walk but instead manifests the absence of one.

TMT also gives us an answer to the question of what it is for an intention to be "in execution." Intentionalist accounts appealed to intentions being "in execution" but didn't specify what this comes to (at least, not in a way that let us make any traction on the Differentiation Puzzle) without circularity. TMT provides a substantive answer here, which is that intention is in execution at t when the agent's practical organisation has the guidance alignment to sustain goal-directed activity such that the end is doing organising work in a control structure with sufficient resilience to carry the process forward.

What about initiatory versus sustaining causation? Causalist accounts needed to distinguish mere initiation of behaviour by an intention from the ongoing control that constitutes genuine action. Bron's summit-intention may play an initiatory causal role at t in the sense that it causes his first step. But initiatory causation isn't not sufficient for action-in-progress as matters for non-basic action underway. What's required is that the intention be embedded in a practical organisation with the resilience to sustain goal-directed activity (and for TMT the result is that Bron's organisation lacks this resilience and thus his intention initiates but does not sustain.)

5 | OBJECTIONS AND REPLIES

I consider now several challenges and offer replies on behalf of TMT:

Objection 1: What about absent-minded actions? We often do things like walking to the store while our mind is elsewhere, i.e., lost in thought, not consciously thinking of the store at every step, etc. (Paul 2009, 5). According to TMT's emphasis on guidance, would such a person lack guidance alignment and thus not be walking to the store? That seems wrong; surely they *are* walking to the store even if daydreaming.

Reply: Guidance alignment doesn't equate to constant conscious attention to the goal. The absent-minded agent can still have high guidance alignment in the functional sense that their unconscious or background guidance systems are operating properly. For example, even while daydreaming like the Sarah Paul's case of the distracted driver (Paul 2009, 5), if the agent reaches the street where they need to turn to get to the store, they reliably make the turn (perhaps automatically). If they took a wrong turn, at some level they would notice and correct (maybe with a start: "Oh, I need to go the other way!"). This shows the end is still controlling behaviour, just at a sub-personal or less explicit level (see, e.g., Sosa 2021, 24–31,

52–58, 64, 110, 118).¹⁶ Guidance alignment thus can operate through habits and automatic processes.¹⁷

Objection 2: “The threshold T is arbitrary.” The Teleological Momentum Thesis introduces a threshold of momentum that must be met to count as doing an action. But the theory doesn’t give anything even approximating a specific value for T . Isn’t this a weakness? How can we know how much momentum is “enough” for a given action, and isn’t this just smuggling in a vague intuitive judgment?

Reply: It’s true that the threshold T is context-sensitive and not sharply defined. But, as gestured to initially in passing, this more plausibly just reflects the reality of our concepts rather than a flaw in the theory. As noted, *writing a dissertation* (or doing a PhD) demands a higher level of sustained engagement to truly be “writing it” (as opposed to procrastinating or merely intending to write it) than *walking to the store* or for that matter *heading to the kitchen* does. The theory acknowledges this by leaving T flexible, which is appropriate. Our ordinary language is replete with such context-sensitive thresholds (consider: how hungry is “hungry enough to eat a horse”? When is someone “tall”?).

Importantly, then, not specifying an exact threshold is not fatal. And beyond that, the theory still provides a *framework* to discuss borderline cases; one can analyse which component is lacking when we hesitate to say someone is ϕ -ing. For instance, if someone hits a few notes on a piano sporadically (with big ideas but modest talent), do they have enough teleological momentum to say they are “writing an etude”? Perhaps not. Intentional mass or practical velocity might be too low. Zooming out: the indeterminacy in hard cases is acceptable and even expected. It mirrors our intuitive indecision in those cases and doesn’t undermine the clear ones. In clear cases (like Aria vs. Bron), the theory is decisive. In borderline cases, it at least explains why they are borderline (some components present but not robust, etc.).

Objection 3: “Does this account introduce an infinite regress of actions?” The worry here is: to be ϕ -ing, we require guidance alignment, which involves the agent attending, correcting, etc., in light of the goal. But aren’t those attending and correcting activities *themselves* actions? If so, do we then require a further teleological momentum toward those sub-actions, and so on ad infinitum? For

16. In fact, much of our intentional action is plausibly guided by intentions without needing continuous conscious focus (see, e.g., Mele 1992).

17. Pacherie (2008) provides a framework of multiple levels of control that range from conscious deliberation down to subpersonal motor routines.

example, if Aria is adjusting her walk by consulting the map (a sub-action) for the sake of reaching the summit, do we need to apply the whole criterion to *that* map-consulting action as well? This seems to lead to a vicious regress of ever-smaller actions requiring momentum – one that the reader will notice bears a structural similarity to Ryle's (1949, 19–20) regress from the *Concept of Mind*.

Reply: This regress is avoided by noting that guidance alignment, while it can involve sub-actions, is fundamentally a *relation* rather than a separate action that itself must be justified. Guidance processes often operate automatically or at a basic level. For instance, keeping one's balance or correcting one's footing are usually not actions in their own right that we deliberate about; they are sub-agential processes or basic actions that occur within the execution of the larger action.¹⁸ If an agent performs a genuine basic action (assuming such exist) or a highly automatic adjustment, we don't need an account of *how* that adjustment is done intentionally beyond the fact that it's done in service of the larger intention. The intention to ϕ can directly organise many low-level behaviours without any new intention for each micro-step. Put another way, *guidance alignment is a structural/constitutive feature of the action, not an endless series of new actions*. This is no more regressive than saying a thermostat regulates temperature (we don't ask “who regulates the regulator?” if the regulator is designed to do so by itself in a basic way).

6 | CONCLUDING REMARKS

The view developed here starts from a simple question: what makes it the case, at a given time, that an agent is performing a non-basic action with a characteristic end, rather than merely doing something that could, under favourable conditions, be extended into it? The answer, on the TM view, is: teleological momentum toward the action's characteristic end. Teleological momentum is a composite property with three components (intentional mass, practical velocity, and guidance alignment) such that, when their contribution meets a context-sensitive threshold, the action is underway.

Importantly the TM framework gives an actual-world basis for the privileged continuations invoked by modal and dispositional accounts and it also makes room for the openness of the progressive without placing the whole explanatory burden on what happens later. A progressive ascription can be true even when completion never occurs, because its truth depends on the agent's present teleological organisation, not on any guarantee of culmination. At the same time, the

18. To borrow terminology here from Sosa (2015), we can think of such actions as analogous with what he calls “functional beliefs” (Ch. 3), as opposed to judgmental beliefs, which we form in response to deliberation about ‘whether p’ question.

framework explains why intention and matched initial movement can fail to suffice for being ϕ -ing at t . Guidance alignment (understood in control-theoretic terms developed here) can differ while intention and first-step behaviour match, and crucially *vis-à-vis* the Differentiation Puzzle, that difference can show up at the very start of an extended action, since practical velocity is then necessarily small in an action-type-relative sense. The Aria/Bron case brings that point into focus.

The proposal was motivated by a metaphysical puzzle adjacent to the imperfection paradox, but it has wider implications in several areas of philosophy, and I'll wrap up by registering just a few of them. Consider, for one thing, that many ordinary assessments of responsibility for temporally extended agency presuppose a distinction between (i) being engaged in an ongoing action and (ii) merely being in a preparatory or merely intention-involving condition. TMT supplies resources for drawing that distinction at a time, and it does so in a way that also allows degrees. The relevant facts needn't reduce to *whether* an intention is present (consider, e.g., *mens rea* in legal philosophy along with broader discussions of whether an agent acted with (or without) a given morally relevant intention). They concern as well whether the intention is weighty enough, whether the agent's activity is progressing in the relevant direction, and whether the end functions as a control variable of what happens next. This bears on various things, including responsibility for abandonment, for the maintenance of our commitments (e.g., including derivative norms related to promise keeping), and more generally for the management of risk in extended projects. A responsibility attribution can track the point at which teleological momentum rises above threshold (the action begins), the period over which it remains above threshold (the action persists), and the point at which it falls below threshold (the action ends). Those distinctions can matter to familiar questions about attempts, omissions, and excuses (e.g., Williamson 2016) in cases where the agent's control organisation fails in a systematic way.

A second area where TM has direct ramifications is debates about akrasia (see, e.g., Davidson 2001; Holton 1999). Akrasia has a temporal profile. It often involves stability in evaluative commitment alongside failure in execution at a particular moment. TMT offers a natural diagnosis. An akratic agent may retain high intentional mass toward the judged-better end while practical velocity and guidance alignment shift toward an incompatible end. At the relevant time, the teleological momentum that meets threshold can be directed toward the tempting end rather than the endorsed one. That picture, an implication of TMT, treats akrasia as a *disturbance in goal-governance*, and it also explains why the akratic agent can report a persisting intention and yet at the same time act against it. The intention stays present, but it doesn't structure the control loop that governs the unfolding activity at the crucial time.

Teleological momentum also suggests a useful way to think about temporally extended collective agency (e.g., [Palermos and Tollesen 2018](#); [Gilbert 2013](#); [Bird 2010](#); [Hutchins 1995](#)), which matters given that the very question of whether a group is currently (in the process of) doing one thing rather than another often has far reaching implications. What TMT tells us is that a group can have teleological momentum toward a collective end when there is (i) an appropriate form of shared commitment or joint intention¹⁹ (collective intentional mass), (ii) coordinated progress toward the end (collective practical velocity), and (iii) a stabilising organisation that keeps the end operative as a control variable for the group's unfolding activity (i.e., "collective guidance alignment"). The third element connects to familiar phenomena in collective action such as, e.g., division of labour, mutual monitoring, role structure, communication, and error-correction. It also offers an elegant way to distinguish cases in which many individuals happen to do end-conducive things from cases in which the group is engaged in a single joint activity. The same framework then bears on questions of collective responsibility (that is, it supports a focus on the organisation that sustains goal-governed coordination over time, as opposed to just a focus on individual intentions taken one by one.)

These implications are, of course, only sketched here and aren't meant to be exhaustive (or even particularly representative). Still, they help illustrate why questions about action in progress, and for the purposes here TMT's way of addressing them, matter beyond familiar climbing cases or the imperfective paradox. The basic idea of teleological momentum offers what promises to be a unifying way of linking the semantics of the progressive with the control structure of temporally extended agency more generally, and also with the normative and social facts that turn on whether an action is genuinely underway at a given time.

REFERENCES

Bechhoefer, John. 2005. "Feedback for Physicists: A Tutorial Essay on Control." *Reviews of Modern Physics* 77 (3): 783–836. <https://doi.org/10.1103/RevModPhys.77.783>.

Bird, Alexander. 2010. "Social Knowing: The Social Sense of Scientific Knowledge." *Philosophical Perspectives* 24: 23–56.

Bondy, Patrick. 2016. "Counterfactuals and Epistemic Basing Relations." *Pacific Philosophical Quarterly* 97 (4): 542–69. <https://doi.org/10.1111/papq.12136>.

Bratman, Michael. 1987. "Intention, Plans, and Practical Reason." <https://philpapers.org/rec/braipa>.

Buckwalter, Wesley, David Rose, and John Turri. 2021. "Impossible Intentions." *American Philosophical Quarterly* 58 (4): 319–32. <https://doi.org/10.2307/48>

19. See here, e.g., [Gilbert \(2013\)](#).

619317.

Carnap, Rudolf. 1967. *The Logical Structure of the World*. University of California Press.

Davidson, Donald. 1970. *Essays on Actions and Events: Philosophical Essays Volume 1*. Clarendon Press.

Davidson, Donald. 2001. "How Is Weakness of the Will Possible?" In *Essays on Actions and Events*, edited by Donald Davidson. Oxford University Press. <https://doi.org/10.1093/0199246270.003.0002>.

Dowty, David R. 1977. "Toward a Semantic Analysis of Verb Aspect and the English 'Imperfective' Progressive." *Linguistics and Philosophy* 1 (1): 45–77. <https://www.jstor.org/stable/25000028>.

Dowty, David R. 1979. *Word Meaning and Montague Grammar: The Semantics of Verbs and Times in Generative Semantics and in Montague's PTQ*. Vol. 7. Springer Science & Business Media.

Engelberg, Stefan. 1999. "The Imperfective-Paradox'Paradox and Other Problems with the Semantics of the Progressive Aspect." *Deseret Language and Linguistic Society Symposium* 25 (1): 5.

Falvey, Kevin. 2000. "Knowledge in Intention." *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 99 (1): 21–44. <https://www.jstor.org/stable/4321043>.

Gilbert, Margaret. 2013. *Joint Commitment: How We Make the Social World*. Oxford University Press.

Glad, Torkel, and Lennart Ljung. 2018. *Control Theory*. CRC press.

Gollwitzer, Peter M. 1999. "Implementation Intentions: Strong Effects of Simple Plans." *American Psychologist* 54 (7): 493. <https://psycnet.apa.org/journals/amp/54/7/493>.

Gollwitzer, Peter M., and Paschal Sheeran. 2006. "Implementation Intentions and Goal Achievement: A Meta-Analysis of Effects and Processes." *Advances in Experimental Social Psychology* 38: 69–119.

Holton, Richard. 1999. "Intention and Weakness of Will." *The Journal of Philosophy* 96 (5): 241–62. <https://doi.org/10.2307/2564667>.

Hutchins, Edwin. 1995. *Cognition in the Wild*. MIT press.

Kelley, Mikayla. 2022. "How to Perform a Nonbasic Action." *Noûs*, ahead of print. <https://doi.org/10.1111/nous.12440>.

Kenny, Anthony. 2003. *Action, Emotion and Will*. Routledge.

Kratzer, Angelika. 1981. "The Notional Category of Modality." *Words, Worlds, and Contexts* 38: 74.

Landman, Fred. 1992. "The Progressive." *Natural Language Semantics* 1 (1): 1–32. <https://doi.org/10.1007/BF02342615>.

Mele, Alfred R. 1992. *Springs of Action: Understanding Intentional Behavior*. Oxford University Press, USA.

Osch, Liesbeth van, Mariëlle Beenackers, Astrid Reubaet, Lilian Lechner, Math Candel, and Hein de Vries. 2009. "Action Planning as Predictor of Health Protective and Health Risk Behavior: An Investigation of Fruit and Snack Consumption." *International Journal of Behavioral Nutrition and Physical Activity* 6 (1): 69. <https://doi.org/10.1186/1479-5868-6-69>.

Pacherie, Elisabeth. 2008. "The Phenomenology of Action: A Conceptual Framework." *Cognition* 107 (1): 179–217. <https://doi.org/10.1016/j.cognition.2007.09.003>.

Palermos, S. Orestis, and Deborah P. Tollefson. 2018. "Group Know-How." *So-*

cially Extended Epistemology, 112–31.

Parsons, Terence. 1990. *Events in the Semantics of English: A Study in Subatomic Semantics*. MIT Press.

Paul, Sarah K. 2009. “How We Know What We’re Doing.” *Philosophers’ Imprint* 9: 1–24.

Phillips, Charles L., and John M. Parr. 2011. *Feedback Control Systems*. Pearson.

Portner, Paul. 1998. “The Progressive in Modal Semantics.” *Language* 74 (4): 760–87. <https://doi.org/10.2307/417002>.

Ryle, Gilbert. 1946. “Knowing How and Knowing That: The Presidential Address.” *Proceedings of the Aristotelian Society* 46: 1–16. <https://doi.org/10.1093/aris/46.1.1>.

Ryle, Gilbert. 1949. *The Concept of Mind*. Hutchinson & Co.

Sheeran, Paschal, Thomas L. Webb, and Peter M. Gollwitzer. 2005. “The Interplay Between Goal Intentions and Implementation Intentions.” *Personality and Social Psychology Bulletin* 31 (1): 87–98. <https://doi.org/10.1177/0146167204271308>.

Sosa, Ernest. 2010. “How Competence Matters in Epistemology.” *Philosophical Perspectives* 24 (1): 465–75.

Sosa, Ernest. 2015. *Judgment & Agency*. Oxford University Press UK.

Sosa, Ernest. 2021. *Epistemic Explanations: A Theory of Telic Normativity, and What It Explains*. Oxford University Press.

Szabó, Zoltán Gendler. 2004. “On the Progressive and the Perfective.” *Noûs* 38 (1): 29–59. <https://www.jstor.org/stable/3506154>.

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

Varasdi, Károly. 2017. “Worlds, Events, and Inertia.” *Journal of Logic, Language and Information* 26 (3): 303–32. <https://doi.org/10.1007/s10849-017-9253-3>.

Vendler, Zeno. 1957. “Verbs and Times.” *The Philosophical Review* 66 (2): 143–60.

Vetter, Barbara. 2011. “On Linking Dispositions and Which Conditionals?” *Mind* 120 (480): 1173–89. <https://doi.org/10.1093/mind/fzr077>.

Vlach, Frank. 1981. “The Semantics of the Progressive.” In *Tense and Aspect*. Brill. <https://brill.com/downloadpdf/edcollbook/title/23221.pdf#page=291>.

Williamson, Timothy. 2016. “Justifications, Excuses, and Sceptical Scenarios.” *The New Evil Demon*. Oxford University Press, Oxford.

Wolfson, Ben. 2012. “Agential Knowledge, Action and Process.” *Theoria* 78 (4): 326–57. <https://doi.org/10.1111/j.1755-2567.2012.01146.x>.

Zhang, Mengjie. 2023. “A Review of Semantic Studies on the English and Chinese Progressive Operators.” *International Journal of Language and Linguistics* 11 (5): 148–55. <https://doi.org/10.11648/j.ijll.20231105.11>.