Know-How in Motion

J. ADAM CARTER *

Abstract: We do all sorts of things. But how much of our behaviour manifests genuine know-how? This practical *scope question* mirrors a familiar epistemological question—how much knowledge there is—and it invites a parallel kind of scepticism: perhaps far less of what we do really manifests know-how than we ordinarily assume. A prominent constraint in the literature appears to deliver such a sceptical answer. According to the *success-only constraint* (SOC), one manifests know-how to ϕ only if one successfully ϕs. This paper argues that SOC is false for *non-basic telic tasks*—extended, multi-step undertakings with a distinctive finish line. In such cases, know-how can be on display in the agent's intelligently controlled route *toward* the goal, even when the goal is never reached. I develop a trajectory-based alternative: a *Good Trajectories* view, on which manifestation is tied not to completion but to the agent's intentional control of a suitably substantial stretch of robustly progress-making activity. The resulting picture preserves the thought that manifestations of know-how must be non-accidental and competence-guided, while allowing interrupted, abandoned, and deliberately "thrown" performances to manifest know-how in the relevant sense. It also supports a less sceptical—and, I argue, more accurate—answer to the scope question: our behaviour manifests significantly more know-how than SOC would allow; when we act (indeed whether we succeed or fail) the manifestation of know-how is normal.

1 | INTRODUCTION

How much of our behaviour actually manifests genuine know-how? The question is the practical analogue of a familiar epistemological question about *scope*: *how much knowledge is there?* We can ask the epistemological question globally (about a community, a culture, a species, etc.) or locally (about knowledge of the external world, moral knowledge, scientific knowledge, and so forth). It also admits a sceptical twist: perhaps we possess far less knowledge than we ordinarily suppose; perhaps the conditions for knowledge are more demanding than everyday attributions acknowledge.

The parallel sceptical challenge arises for *practical intelligence*. Even granting that people often know how to cook, drive, teach, negotiate, play chess, and so on,

*Draft 18/01/26. Comments welcome. This work is supported the European Research Council grant number 101199255 (KNOW-HOW).

we can ask a different question: *how much of what people do manifests whatever know-how they possess?* How often is know-how actually *on display* in action? How wide is the footprint of *manifested* know-how in ordinary life?

This question matters because it bears directly on explanation—and, as I will suggest in §8, on how plausible it is that know-how can function as a norm in practical life. Much of what we do is purposive and structured: we pursue ends by means, improvise when circumstances shift, correct course when something begins to go wrong. A natural picture is that know-how is not merely a standing state that sits behind the scenes, but something that regularly shows itself in the fine grain of our conduct. But a deflationary picture is always available: perhaps know-how is comparatively rare in manifestation. Perhaps much of what looks like intelligent action is routine habit, or a lucky run, or something that falls short of genuine know-how.

A prominent route to the deflationary picture begins with a simple constraint on manifestation: one manifests know-how to ϕ only if one successfully ϕs. Call this the *success-only constraint* (SOC). If SOC were correct, the footprint of manifested know-how would shrink to the footprint of successful completions. Since so much of our goal-directed activity is interrupted, abandoned, strategically forfeited, or otherwise fails to culminate in success, SOC would deliver a strongly sceptical answer to the scope question.

I will argue that SOC is mistaken, at least for a central class of cases: extended, multi-step, goal-directed activities—what I will call *non-basic telic tasks*. In such cases, know-how can be on display in the agent's intelligently controlled route toward the goal even when the goal is never reached. I then develop a trajectory-based alternative—the *Good Trajectories* account—on which manifestation consists, roughly, in the agent's intentional control of a sequence of robustly progress-making moves. Rejecting SOC supports a less sceptical (and more realistic) verdict about practical intelligence: we often manifest know-how in action even when we do not complete what we set out to do.

The plan is as follows. Section 2 formulates the scope question more carefully and motivates SOC. Section 3 explains why the debate is most pressing for telic, non-basic tasks, where it is natural to assess performance "along the way" rather than only at the finish line. Section 4 presents a set of cases that put pressure on SOC. Sections 5–6 develop and apply the Good Trajectories alternative. Section 7 answers objections and clarifies how the view avoids making manifestation either too cheap or too demanding. Section 8 returns to the scope question and draws the upshot about the extent of manifested know-how in ordinary life.

2 | THE SCOPE QUESTION AND THE SUCCESS ONLY CONSTRAINT

2.1 | Possession vs. manifestation

By "scope of know-how" I do not mean how much know-how agents simply *possess* irrespective of what they do. The guiding question is more specific:

> **Scope question:** *How much of what agents do is aptly described as manifesting whatever know-how they possess?*

Manifestation is a relation between a standing state (knowing how to do something) and an episode of action. For convenience, let $KH(S, \phi)$ abbreviate "S knows how to ϕ", and let $M(S, \phi)$ abbreviate "in episode E, S manifests know-how to ϕ".

Nothing here turns on a particular metaphysics of know-how—whether knowing how is propositional, ability-based, hybrid, and so on.[1] The issue is what constraints govern *manifestation*. One minimal constraint is that manifestation is, in a broad sense, action-involving:

> **Action involvement:** If $M(S, \phi, E)$, then E includes some relevant practical engagement by S (at least attempting, initiating, or exercising control in a ϕ-directed way).[2]

This constraint is deliberately weak. It allows that manifestation might occur in attempts, partial performances, or mid-stream control—precisely the kind of cases under debate.

2.2 | SOC: a hard constraint on manifestation

The thesis I will challenge is a one-way constraint:

> **SOC (Success-only constraint).** For any agent S, telic task ϕ, and episode E: If $M(S, \phi, E)$, then S succeeds in ϕ in E.

1. This methodological stance parallels the *particularist* response to the problem of the criterion in mainstream epistemology (see Chisholm 1977). Particularists prioritise "What do we know?" over "By what criteria do we decide?", whereas methodists reverse this order (see also Sosa (1980), 4). While the particularist relies on implicit criteria to identify clear instances of knowledge, those pre-theoretical platitudes are not intended to adjudicate between competing analyses of the nature of knowledge. Similarly, I rely here on clear cases of manifestation without presupposing a specific metaphysics of know-how.

2. I intend "action" broadly enough to include mental actions; for instance, if (as Sosa (2015) suggests) judgement is a kind of intentional action. For simplicity I focus mainly on non-purely mental action, but the account generalises.

SOC is attractive partly because of its simplicity. It makes the manifestation of know-how outcome-tight: know-how to ϕ manifests only when ϕ is successfully achieved. The footprint of manifested know-how is therefore bounded by the footprint of successful completions.

Something like SOC has recognisable precedent. Cunningham (2022) remarks that it sounds "very odd" to say a pianist manifested her knowing how to play a fugue without actually succeeding; the more natural description is that she *attempted* to display her know-how but failed. He links the point to an analogy with knowing-that: knowing how and knowing that are both, in some sense, "essentially successful" states; knowing-that p is essentially successful because it requires a belief that is "as it ought to be" (e.g. a true belief), and knowing how to F is essentially successful because it "can manifest itself only if its manifestation is a successful instance of F" (2022, 403). Earlier voices make similar claims. Gellner (1951) suggests that when we attribute know-how on the basis of observation, what we are interpreted as having observed was successful performance rather than failed attempt. More recently, Andrada (2024) has defended an ability-sensitive variant: for performance to manifest know-how it is not enough that it succeeds; its success must be appropriately due to the agent's well-regulated ability.

A wider cluster of views linking know-how to successful performance can seem to pull in the same direction. Ryle's (1949) talk of know-how being "exercised" in what one does can be read as privileging successful intelligent performances over failed attempts. Hawley's (2003) "successful action plus warrant" analysis builds success into the account of knowing-how itself, in such a way as to encourage the thought that manifestation is success-bound. Still more radical ability-based approaches imply that knowledge (including know-how) is an ability whose apt exercise yields success; on that picture, the manifestation of know-how would plausibly coincide with successful exercise.[3]

These considerations motivate a picture on which the *paradigm* manifestation of know-how is a successful performance of what the agent knows how to do, and failure is treated as at best an attempted display. That picture also makes SOC look like a natural answer to the scope question. For if SOC is true, then every incomplete attempt—that is, every performance cut short or abandoned before success—falls outside the realm of manifestation. Since such incomplete episodes are ubiquitous, SOC implies a sparse footprint of manifested know-how across

3. Sympathy for a success-only constraint is present in, for example, Cunningham (2022), Andrada (2024), Hawley (2003), Carr (1981), Damschen (2009), Roland (1958), Hartland-Swann (1956), and Carter and Pritchard (2015).

our lives.

That is exactly why the debate matters. If SOC were correct, manifested know-how would be the tip of the iceberg. Much of what we ordinarily regard as intelligent action—performances interrupted, abandoned, forfeited, or derailed—would count, in the strict sense, manifestation-wise, for nothing.

The remainder of the paper argues that this picture is mistaken, at least for the kinds of tasks where the manifestation question is philosophically most pressing.

3 | TELIC STRUCTURE AND WHERE SOC BITES

SOC is most tempting when the task is *telic*: it comes with a distinctive completion condition, a "finish line". Let the goal state for ϕ be G_ϕ. Then ϕ is completed just in case G_ϕ. obtains: checkmate is delivered; the risotto is plated; the summit is reached; the election is won.

By contrast, many activities are *atelic*: they are assessable as going well or badly, better or worse, but they do not have an intrinsic terminal condition of the relevant kind (being a good friend, maintaining a pleasant household) (see Setiya 2014, 12–13). Those cases raise their own questions. Since SOC is naturally formulated in the telic setting, I will focus on telic tasks.

Even within the telic domain, SOC does not face equal pressure across the board. The sharpest challenges arise for *non-basic* telic tasks: tasks that characteristically require a temporally extended sequence of choice-sensitive steps (see, e.g., Danto 1965, 1973; Bratman 1987; Lavin 2013; Altshuler 2025; Kelley 2022). For such tasks, it is natural to evaluate performance *mid-course*, not merely at the finish line.

We can roughly mark the contrast as follows (agent-relatively, as it should be). A telic task may be *basic for S* in circumstances C when S can achieve its goal by a single "doable-now" act at the grain of choice-sensitive control. A telic task is *non-basic for S* in C when achieving the goal characteristically requires a structured sequence of moves: intermediate stages can be navigated better or worse, and competent route-management is visible long before completion.

The agent-relative character of the distinction matters. What is basic for an expert may not be basic for a novice.[4] A concert pianist executes an arpeggiated passage as a single motor chunk; a novice must consciously manage each fingering and hand transition. The expert's control point spans a longer temporal arc.

4. This reflects the point about "chunking" in cognitive science: with expertise, agents come to represent and control what were previously many steps as single higher-level units, shifting what counts as "doable now". Classic statements include Miller (1956) and Chase & Simon (1973) on chess expertise.

Similarly, a grandmaster's "mate in one" may be "doable now" in a way it is not for a beginner whose control points are at a different grain.

Non-basic telic tasks also support a substantive notion of *progress* that is not merely "having done something that is necessary for success" in a trivial sense. One can progress in a way that is systematically guided; it can be responsive to feedback, regulated across intermediate stages, sensitive to contingencies and so forth. Those are paradigmatic marks of practical intelligence. They are also precisely what SOC threatens to ignore, because SOC forces the verdict "no completion, no manifestation" on every non-completing case alike.

None of this denies that successful completion is a *paradigm* manifestation of know-how, or at least, that it features in paradigm manifestations of know-how. The point is that it is not the only form. Even granting success as paradigmatic present, SOC can still be false.

4 | CHALLENGES FOR SOC

4.1 | Interruption without fault

Interruption cases put pressure on SOC as a constraint on manifestation (not merely on possession).

> **Interrupted Risotto.** Lina, an experienced chef, knows how to make risotto. She begins cooking in the normal way: heating the pan, toasting the rice, adding stock in the right increments, stirring at the proper cadence, tasting and adjusting. After several steps, an exogenous emergency (say, a gas leak in her building) forces an evacuation and she must stop. No risotto is plated.

It is natural to say that Lina's know-how to make risotto is *on display* in what she does up to the interruption. She's not just flailing about (as one would expect from a child); and nor is what Lina's doing here just educated guesswork. By hypothesis, everything she does prior to the evacuation is exemplary risotto-making: she navigates the relevant control points competently, makes appropriate adjustments, and guides the process in a way that is intelligibly directed toward producing a high-quality dish. The interruption prevents completion, but it does not retrospectively strip these earlier stages of their character as skilled, purposive action.

The SOC verdict is therefore: no completion, hence no manifestation. That verdict is implausible. The most obvious way for the defender of SOC to resist is to insist that manifestation can occur only at the endpoint, so that until completion

there is strictly no manifestation at all. But that is not merely counterintuitive; it undermines the explanatory role of know-how attributions. We routinely judge, while a performance is still underway, that the agent knows what she is doing. In a paradigmatically competent completion of a non-basic telic task, know-how does not suddenly appear at the final instant; it is displayed in the unfolding control of the route. Denying that is an excessive revision of ordinary and theoretical practice.

Two principles bring out what is at stake. First:

> **Prefix invariance (PI).** Whether S's know-how is manifested *up to time t* depends only on what happens up to time t—not on what happens later.

PI is a constraint against backward dependence: later accidents, later interruptions, later changes of mind should not retroactively determine the performance quality of earlier stages. If two episodes match perfectly up to time t in all task-relevant respects, they should receive the same verdict about know-how manifestation up to t. Second:

> **Non-instantaneous manifestation (NIM).** In paradigmatically competent completions of non-basic telic tasks, know-how is manifested already before the final instant.

NIM is modest. It does *not* say that success entails manifestation, or that every successful completion is a manifestation (fluke successes obviously need not be). It says only that in non-basic telic tasks performed at the gold standard to success, know-how is on display already in some earlier stretch of the performance.

Given PI and NIM, the pressure on SOC is straightforward. Take a paradigmatically competent completion of risotto, and consider an earlier time t before plating. By NIM, Lina's know-how is manifested up to t in that completion. Now consider an episode that is identical up to t but is interrupted at t by an exogenous emergency. By PI, Lina's know-how is still manifested up to t in the interrupted episode. But the interrupted episode ends at t and is not a completion. So we have know-how manifestation without success, contrary to SOC.[5]

5. Here is a precise and more general statement of the argument. Let E_{comp} be a paradigmatically competent completion of a non-basic telic task ϕ—that is, an episode such that both $\text{Succ}(S, \phi, E_{\text{comp}})$ and $M(S, \phi, E_{\text{comp}})$ hold. Let T be the completion time of E_{comp}. By (NIM), there exists some time $t^* < T$ such that $M(S, \phi, E_{\text{comp},\leq t^*})$. Now consider an interruption episode E_{int} that matches E_{comp} perfectly through time t^* in all ϕ-relevant respects (the same ϕ-directed moves under the same relevant conditions), and is cut short *immediately after* t^* by an exogenous interference (through no fault of S).

A further diagnostic point comes from comparing *Interrupted Risotto* with a superficially similar case in which non-completion reflects an internal breakdown rather than exogenous interruption:

> **Forgotten Risotto.** Lina begins in the same way, but later genuinely forgets what to do next. She stops in confusion, unable to proceed.

In *Forgotten Risotto* one may reasonably hesitate—an issue to which I return later[6]—to say that Lina's performance manifests know-how to make risotto, since her inability to continue plausibly reflects an outcome-determinative gap in the knowledge required for the task. In *Interrupted Risotto*, by contrast, the failure to complete is purely exogenous. SOC cannot distinguish the cases: both are non-completions, so both are "no manifestation". That is the wrong shape for an account of manifestation in non-basic telic domains.

4.2 | Deliberate non-completion: the thrown-success pattern

A defender of SOC might try to dismiss interruption cases by saying that external interference somehow "blocks" manifestation throughout. That response is unpromising, but we can avoid it altogether by considering cases in which the agent could easily complete yet deliberately chooses not to. Here the lack of completion is due neither to outside interference nor to an internal lapse of competence. It is a deliberate withholding of the endpoint by an agent who is plainly in command of the task.

> **Straight mate.** A chess grandmaster S is in a position from which she can force checkmate. She plays a forcing sequence that constructs a mating net and then delivers checkmate. This is paradigmatically elite level play.

> **Thrown mate.** Same initial position and same build-up. S plays exactly the same forcing sequence up to some point: she constructs the mating net. But then she deliberately changes course and plays

Since the episode ends before the telic terminus is reached, we have: $\neg\mathrm{Succ}(S, \phi, E_{\mathrm{int}})$. By (PI) and the perfect match through t^*, the manifestation verdict for the prefix through t^* is preserved. Hence: $M(S, \phi, E_{\mathrm{int},\leq t^*})$. But E_{int} ends at t^*, so: $E_{\mathrm{int}} = E_{\mathrm{int},\leq t^*}$. Therefore, $M(S, \phi, E_{\mathrm{int}})$. We have thus exhibited an episode E such that: $M(S, \phi, E)$ and $\neg\mathrm{Succ}(S, \phi, E)$. This contradicts the (SOC), which holds that manifestation entails success $[M(S, \phi, E) \rightarrow \mathrm{Succ}(S, \phi, E)]$. So SOC is false (at least for non-basic telic tasks).

6. This case admits of different variations which shift the extent and cause of the knowledge gap; these different variations are discussed in detail in §6.1.

> a move that releases the net and allows the opponent to escape. She does so intentionally, not through confusion or inattention. No checkmate occurs.

The natural verdict is that the grandmaster's mastery is on display in the net-construction itself. Indeed, in chess commentary one can recognise deep understanding in a sequence of moves even when the game ends before checkmate is delivered.[7] SOC forces the opposite verdict: since there is no checkmate, the episode cannot manifest know-how to checkmate. That is implausible. The grandmaster's practical intelligence in setting up the mate is not cancelled by her *later* decision to play differently.

Once again, the underlying structure is the same. In a paradigmatically competent completion of checkmating, the grandmaster's know-how is manifested before the final move (non-instantaneous manifestation), by NIM. The thrown episode matches the successful episode up to that point. By prefix invariance, what was manifested in the prefix remains manifested. But the episode does not end in success. So SOC is false.[8]

One might object that, since S intends not to checkmate, she is not "undertaking checkmating" at all. But this objection trades on an implausibly global view of intention. The relevant intention for undertaking a telic task can be local and time-indexed: S can intend to execute a checkmating sequence up to a point and also intend, for independent reasons, not to deliver the final move. Later changes of intention do not erase earlier control. That is the whole point of the case.

7. Chess commentators often describe a move or sequence as *deep*, connoting highly skilled planning and foresight. Deep sequences can feature in ultimately aborted mating nets. The natural judgement is that the depth was displayed in the sequence itself, not only in the eventual delivery of mate.

8. A more general version of the argument, put precisely, proceeds as follows: Let E_{straight} be a paradigmatically competent completion of the non-basic telic task ϕ, with completion time T, and assume that $\text{Succ}(S, \phi, E_{\text{straight}})$ and $M(S, \phi, E_{\text{straight}})$. By *Non-Instantaneous Manifestation (NIM)*, there exists some time $t^* < T$ such that $M(S, \phi, E_{\text{straight},\leq t^*})$. Now consider the corresponding *thrown* ϕ-episode E_{thrown} that matches E_{straight} perfectly through time t^* in all ϕ-relevant respects (the same ϕ-directed moves under the same relevant conditions), and then *immediately after* t^* the agent intentionally ceases to pursue ϕ (switching aims and playing to lose). Since the telic terminus is not reached in E_{thrown}, we have: $\neg\text{Succ}(S, \phi, E_{\text{thrown}})$. By *Prefix Invariance (PI)* and the perfect match through t^*, the manifestation verdict for the prefix through t^* is preserved. Hence: $M(S, \phi, E_{\text{thrown},\leq t^*})$. But E_{thrown} is the ϕ-episode, and it terminates when ϕ-directed activity ceases—i.e. at t^*. So: $E_{\text{thrown}} = E_{\text{thrown},\leq t^*}$. Therefore, $M(S, \phi, E_{\text{thrown}})$. We have thus exhibited an episode E such that: $M(S, \phi, E) \wedge \neg\text{Succ}(S, \phi, E)$, contradicting **SOC**, which says that manifestation entails success. Hence, SOC is false (at least for non-basic telic tasks).

4.3 | Abandonment and fatigue

Interruption cases show that completion can fail for external reasons. Abandonment cases bring out a different point: even when an agent's conduct up to some point is indistinguishable from the conduct of a paradigmaticly gold-standard performance, the agent may nevertheless just choose to stop.

> **Tired climber.** Let ϕ be the task of summiting this peak by the standard route under normal conditions. An expert mountaineer S undertakes ϕ intending to reach the summit. She proceeds in the characteristic way: chooses a sensible line, manages pacing and nutrition, adjusts to terrain and weather, corrects minor slips without loss of control, etc. After a long, competent ascent, she deliberately turns back because she is tired and prefers not to continue with a narrower margin for safety. She could very likely reach the summit if she continued. She abandons the summit attempt. No summit is reached.

Contrast this with an exhausted climber who turns back because she has mismanaged pacing or nutrition and can no longer continue without collapse. In that latter case it is much less clear that the earlier stretch manifests know-how to summit; the failure plausibly reflects serious gaps in the climber's skill set.

But in *Tired climber* the natural verdict is that S's know-how to summit this peak is on display in the summit-directed route-management prior to the turning point.[9] SOC again insists: no summit, therefore no manifestation. That verdict is implausible for the same reasons as before. Up to the turning point, the climber's conduct is exactly the sort of controlled summit-directed navigation that would ordinarily culminate in success. The later decision to stop does not erase what was already on display.

Here too SOC collapses distinct kinds of non-completion: exogenous interruption, internal breakdown, and deliberate abandonment. That is already enough to make SOC suspect as a general constraint on manifestation.

4.4 | Gradability: why endpoint-only views are too coarse

The previous cases supply direct counterexamples to SOC's conditional. There is also a more theoretical objection. SOC is too crude to play the explanatory

9. Nothing here requires the strong claim that S manifests know-how to summit throughout a longer episode that includes the descent. The claim is only that S manifests know-how up to the turning point, and that manifestation can cease once summit-directed control ceases. SOC cannot accommodate this time-sensitive verdict, since by its lights any failure to complete wipes out manifestation wholesale

role that "manifested know-how" ought to play, at least for non-basic telic performances.

A familiar datum is that know-how is *gradable*.[10] We quite ordinarily say that someone knows *in part* how to do something, knows *well enough* how, or only *sort of* knows how. As Ryle (1949) noted, it is perfectly ordinary to speak of someone "knowing in part how to do something," even though "partial knowledge" of a fact is only intelligible in the derivative sense of knowing *some* facts but not others (Ryle 1949, 59; see also Pavese 2017, 345–46). Bengson and Moffett (2011) likewise emphasize that we routinely say one "*really* knows how," "knows *well enough* how," or "*only sort of* knows how," and they explicitly connect this to a salient scale—*degree of mastery* (Bengson and Moffett 2011, 184, my italics).

Consider a simple case.

> **Circuitous but competent route.** There is an elegant, strongly progress-making route to G_ϕ. S instead follows an unnecessarily long and cautious route. Each step is only barely progress-making. S succeeds, but only barely within what would count as an acceptable manner or duration for the relevant task type.

Here we do not naturally say that *S* manifested *no* know-how (whatsoever). (A driver can still manifest know-how by getting across a city via a labyrinth of side streets and back roads, even when a faster motorway route was available; someone can still manifest know-how by assembling IKEA furniture, or say a LEGO model of the Roman Colosseum, through an over-careful, repeatedly re-read-the-manual approach, even if it takes twice as long as necessary). In such cases we say rather that S showed *some* know-how but perhaps not much: what S did was competent (enough) but *rather inefficient*, good control but a thin margin, and so on. These are all assessments of the *trajectory*—how well the agent manages the unfolding performance.

SOC's structure prevents it from recognising any gradations of manifest know-how in competent failures—near-misses in which partial mastery is on display despite the outcome. An account that cannot accommodate the trajectory-sensitive, gradable character of know-how manifestation is explanatorially inadequate for non-basic telic domains.

The best diagnosis, then, is that SOC is (i) extensionally wrong about non-basic telic tasks and (ii) structurally too crude.

10. For a now-canonical defence of the gradability of know-how—offered *in support of* intellectualism, despite gradability often being cited as an objection to it—see Pavese (2017).

5 | A GOOD TRAJECTORIES ACCOUNT OF MANIFESTING KNOW HOW

The counterexamples above suggest a different conception of where know-how manifests in extended action. In *non-basic telic* tasks (extended, multi-step tasks with a definitive goal state G_ϕ), know-how can be on display *mid-course* – in the intelligently controlled navigation of a route toward the goal – even if the goal is never reached. An adequate account of manifestation should therefore be sensitive to the *internal structure* of an attempt (how the agent performs at intermediate stages), not merely to whether the endpoint is attained.

To develop this idea, we need a principled way to evaluate *how* a performance unfolds. The guiding insight is that, in an extended task, practical intelligence is displayed not only in final success, and not merely in the fact of *trying*, but in the *choices and adjustments* by which the agent manages the task as it progresses. In other words, between mere trying to ϕ and achieving ϕ lies a middle level of granularity at which one's knowledge how to ϕ is evident.[11] We can capture this by analyzing an extended performance as a *trajectory* through a space of task states and available actions, and assessing the quality of the *moves* that constitute that trajectory. The account developed here will define key notions – *move*, *progress*, *robustness*, *good move*, *trajectory segment*, and *intentional control* – and use them to state a formally precise condition for manifesting know-how.

5.1 | Moves: The Right Grain of Evaluation

To make sense of "mid-course" assessment—and, in turn, of mid-course manifestation of know-how—we need to identify that level of description.

Let's introduce a task-relative unit – a *move*– between subpersonal micromotions (which we can't intentionally select) and whole-task completion:

> **Move (token).** A move token m for S in state s (relative to task T_ϕ) is a choice-sensitive action available to S at s that is individuated by its functional role in bringing about G_ϕ.

11. When completion is unsettled (or failure salient), we often shift from "S ϕ-ed" to "S tried to ϕ," and action theorists have emphasised how centrally 'trying' figures in our ordinary agency-descriptions (Hornsby 2010). But for *non-basic telic* activity, bare trying is still too undifferentiated to serve as the unit of assessment. A parallel moral shows up in epistemology: a belief-centred explanatory tradition is often said to account for an agent's *attempt* to get things right while treating success as something further (e.g. truth/knowledge) (for critical discussion see Williamson 2017; Miracchi and Carter 2022; Pavese 2020, 345). The notion of a *move* is meant to regiment this finer, *within-the-attempt* grain—precisely the grain at which it can be tempting (though, I argue, misleading) to describe competent non-completions as merely "attempted" displays of know-how (cf. Cunningham (2022), 403).

Note that different bodily executions can constitute the same move (e.g. one can stir a risotto clockwise or counterclockwise – physically different motions but functionally the same *stirring* move), and conversely the same physical motion (raising a finger) could count as a different move (e.g., asking for a cheque versus bidding on a painting) under a different task-description (because its contribution to G_ϕ shifts).

A move is also *choice-sensitive*: it represents a *control point* at which the agent faces a non-trivial set of alternative courses of action. In other words, given the agent's situation and intentions, there are multiple admissible moves available, and the agent must choose one.[12] At this level of granularity, it becomes appropriate to evaluate performance as better or worse: instead of merely saying "S is trying" or "S succeeded," we can say *"S is doing the right thing* here*, as opposed to other, worse options*."

By analysing an extended activity in terms of a *sequence of moves*, we obtain a structured trajectory through the state-space of the task. This trajectory-based structure allows us to talk about *mid-course progress* in a way that's sensitive to the *quality of the route* taken toward the goal.

5.2 | Good Moves

Even without reaching the final goal, some sequences of ϕ-directed moves clearly reflect greater knowledge how to ϕ than others. Intuitively, some moves put the agent in a *better position* to eventually succeed than others do. This suggests a way to regiment the notion of "moving toward the telos": we evaluate whether each move contributes to *making progress* toward G_ϕ, and how *secure* or stable that progress is.

> **Progress-Making Move:** A move M (taken from state s to a subsequent state s') is *progress-making* if and only if performing M *reduces the risk of eventual failure* (or equivalently, increases the security of eventual success) relative to state s, assuming a normal continuation of the attempt from s'.

12. The claim that moves are *choice-sensitive* does not require that the agent can explicitly survey, represent, or articulate the full set of admissible continuations at each point in an extended performance. In non-basic telic tasks, an intention to ϕ typically commits the agent only to a *partial plan*, whose remaining structure must be settled as circumstances unfold. As Bratman (1987) observes, "the plans for action contained in our intentions are typically partial and must be filled out in accordance with changing circumstances as the future comes" (Bratman 1987, 20). On the present proposal, *moves* correspond to the locally controllable sites of this filling-out: points at which the agent's intention leaves open multiple eligible ways of proceeding, and where practical intelligence, whether reflective or unreflective, selects how to continue.

In other words, holding fixed the relevant background conditions and assuming a normal continuation plan, M leaves the agent closer to success: it increases the margin for error, closes off some normal routes to failure, keeps viable paths to G_ϕ open, or secures an intermediate sub-goal that would typically be needed to reach G_ϕ. Conversely, a move is *progress-undermining* if it makes eventual success more *precarious*: for example, it opens new routes to failure that would otherwise be avoided, or forces the agent into such a narrow, fragile path forward that even slight perturbations could derail the attempt.

Not any kind of "progress" is sufficient for displaying know-how, however. An agent might stumble into making some progress by *accident* – for instance, a driver might get closer to a destination by blindly flipping a coin at each intersection and luckily choosing the right sequence of turns. Even if the outcome happens to be success, we would hesitate to credit the driver with genuine know-how in navigation. The success in such a case is *fragile* or *lucky* relative to what could easily have happened instead. To rule out crediting lucky flukes, a *robustness* condition on progress is needed.

> **Robustness of a Move:** A move M is *robustly progress-making* if and only if M's progress-conduciveness is *stable under minor variations* in circumstances.

Roughly, M should not be a one-off stroke of luck that would turn counterproductive under slightly different conditions. Formally, consider a set of nearby, task-relevant possible situations (small variations in timing, environment, opponent behaviour, etc. from the actual situation in which M is taken). M counts as robust only if in each such nearby situation, performing the analogous move M does *not* make success significantly less likely – ideally, M would continue to be progress-making, or at least would not undermine the agent's position. In short, a move is *safe* against slight perturbations: it would not easily become a bad move under conditions very close to the actual ones (compare the notion of a "safe" belief that remains true in nearby possible worlds in the sense of, e.g., Pritchard (2005)).[13]

13. I do not assume that the goodness of a move is *luminous* to an agent who knows how to ϕ, or even to an agent who is in the course of manifesting such know-how. To assume luminosity here would be implausible. Since what makes a move good can vary by imperceptibly small margins, combining luminosity with a standard *margin-for-error* principle (Williamson 2000, Ch. 5)—roughly, that if one knows a move is good in case α, then the relevant condition must obtain in all sufficiently nearby cases α^*—would quickly lead to a reductio. The present account therefore allows that agents can reliably make good moves, and thereby manifest know-how, without being in a position to know that the moves they are making are good.

We can now combine these ideas to define what it is for an individual move to be *good* in the relevant sense.

> **Good Move:** A move is a *good move* (relative to task ϕ) if and only if (i) it makes non-trivial progress toward G_ϕ in the actual situation (more than a negligible improvement), and (ii) it does so in a *robust* way: across the relevant neighbourhood of possible situations, performing that same type of move would not easily undermine progress.

In short, a *good move* is one that yields a genuine step toward the goal and is not merely a lucky accident suited only to the exact actual circumstances. Given the notion of a good move, we can extend it to sequences of moves. A *trajectory* in the task is any sequence of states and moves the agent goes through, when attempting a task, from the start toward the goal.[14] A *trajectory segment* is a finite portion of such a sequence – for example, the first n moves of an attempt (a chess *opening*), or any contiguous stretch of the agent's activity. We say what makes a segment *good* in terms of the goodness of its constituent moves:

> **Good Trajectory Segment:** A trajectory segment (within an episode of attempting ϕ) is *good* if and only if *every move* in that segment is a good move. Equivalently, it is a contiguous stretch of performance during which each step the agent takes is a robustly progress-making move.

Importantly, the goodness of a segment is evaluated *locally* in the sense that a segment can qualify as good even if the overall attempt later goes "off the rails" due to interruption, quitting, or subsequent poor moves. We only require that *up to that point* in the sequence, the agent's route-management consisted entirely of robustly progress-making choices. (Goodness is thus *prefix-sensitive* rather than all-or-nothing; a later failure doesn't erase the fact that an earlier part of the attempt was well-executed.)

14. As Anscombe (1957) observed, we use the imperfect aspect to describe individuals on such trajectories. So and so "is crossing the street". A person who would never have taken more than one step towards the street, is not on the same trajectory as one who is crossing the street, even if the initial first step is indistinguishable. I'll return to this point.

5.3 | Control and the Anti-Luck Element

So far we have characterised what it is for moves (and trajectory segments) to be *objectively* good: they are progress-conducive and robust. But objective goodness does not guarantee the manifestation of know-how. An agent might traverse an objectively good segment by luck, or because some external factor just steers her along the good trajectory. To capture a further anti-luck element in know-how manifestation, one that goes beyond the mere robustness of a given move, we therefore need a further condition. Intuitively, when an agent manifests know-how, it is to *her* credit—not merely to fortune or to a guide—that she makes good (robustly progress-making) rather than bad moves in pursuit of the task goal G_ϕ.

I will use *control* as a term of art for the familiar, intention-guided kind of agential control emphasised in action theory and the philosophy of skill (see, e.g., Shepherd 2014; Fridland 2017; Bratman 1987; Frankfurt 1997): a move is, to a first approximation, controlled when it issues from the agent's own goal-directed agency in a non-deviant way. The point is to capture an "in-the-driver's-seat" relation to one's moves (good or otherwise) without building know-how (or its manifestation) into the analysis. In particular, control is neutral between objectively good and objectively bad moves; what matters here is that *good* moves are candidates for manifesting know-how only when they are controlled.

Two minimal constraints on control are salient. First *reasons-responsivenes*. In selecting and executing a move, the agent must be appropriately responsive to reasons.[15] The point is not that she must explicitly deliberate, but that her control system is disposed to adjust what she does in the light of relevant considerations. This rules out cases in which the explanation of why she makes the move bypasses her responsiveness to reasons—for example, cases of compulsion, random guessing, or covert external steering.

Second: non-deviance. The causal route from the agent's intention (or goal-directed aim) to her performance of the move must be non-deviant (Davidson 1971; Kelley 2022; Kearl 2023). The intention must guide the move in the normal way, rather than via bizarre accidents, intervening causal loops, or other deviant chains. Otherwise, the intention may cause the outcome without *controlling* it; the agent succeeds "by mistake".

These constraints yield the following working definitions.

15. See, e.g., Fischer and Ravizza (2000) for a canonical presentation of reasons responsiveness in connection with actions that we 'own'. See also (for an extension of reasons-responsiveness) to epistemology, Carter (2022).

> **Controlled move.** *S* intentionally controls a move *m* at time *t* (under the aim of pursuing G_ϕ or some subgoal) iff: (i) *S* performs *m* under an intention at *t* to advance G_ϕ (or the relevant subgoal); (ii) that intention guides the performance of *m* via a non-deviant causal route; and (iii) *S*'s selection and execution of *m* are reasons-responsive in the sense just sketched.
>
> **Controlled segment.** A trajectory segment $\sigma = \langle m_1, \ldots, m_n \rangle$ is under *S*'s control iff, for each $i \leq n$, *S* intentionally controls m_i at the point of performing it.

Control in this sense can be exercised locally and stepwise. The agent need not plan the whole segment in advance.[16] What matters is that, at each decision point, she governs what to do next under the standing aim, adjusting as the task unfolds. (As throughout, "move" should be understood at a task-relative grain: we do not require control over every micro-movement, but over the moves that are candidates for being objectively good or bad in the sense described.)

5.4 | The Account Stated

We are now in a position to state the positive account of manifested know-how more formally. In what follows, let ϕ be a non-basic telic task with a distinctive goal state G_ϕ, and let *E* be an episode in which an agent *S* undertakes task ϕ (an attempt at ϕ that may or may not succeed). The account says:

> **Good Trajectories (GT).** For an agent *S*, task ϕ, and attempt episode *E*: *E* manifests *S*'s know-how to ϕ in so far as *E* contains some non-trivial trajectory segment (some contiguous stretch of ϕ-directed activity) such that *S* controls a good trajectory segment toward G_ϕ.

In other words, within *E* there must exist a span of action in which S is *in control* of each move, and each move in that span is a robustly progress-making step toward the goal G_ϕ (a step that increases the security of eventual success in a way that would not easily be reversed by minor variations). When such a controlled good segment is present, know-how to ϕ is on display, even if *E* as a whole does not culminate in the successful attainment of G_ϕ.

Several aspects of this GT condition are worth highlighting. First, *locality*. The verdict about manifestation can be determined by looking at an *initial segment* of

16. See, on this point, Bratman (1987, 20) on partial plans.

the episode in isolation. If a controlled good segment occurs in the *prefix* of E, then *that part* of E manifests know-how – regardless of what happens afterward. So if the attempt later derails or is cut short, it (importantly, given criticisms of SOC) does not retroactively cancel the fact that know-how manifested earlier. In this way, GT vindicates the intuitive idea of *prefix invariance*: once genuine know-how has been displayed in the unfolding of the task, it cannot be negated by a subsequent interruption or failure (though of course the agent might fail to *continue* manifesting know-how after that point).

Second, the condition *is fully completion-independent* in the sense that its satisfaction depends *only* on making controlled, robust progress *toward* completion. Thus, even an interrupted, abandoned, or deliberately "thrown" performance can manifest know-how, provided it contains a suitable controlled good segment. And third, GT makes manifested know-how a gradable property (thus accommodating talk of degrees of manifestation), which fits ordinary assessments. For example, maintaining robust progress towards ϕ-ing for a few minor steps involves manifesting some know-how to ϕ, whereas doing so through many challenging steps under pressure reveals a higher degree of know-how (to ϕ) manifestation. GT can thus distinguish *how much* know-how is exhibited, where SOC would treat any non-success as *equally* zero.

It should be emphasized that GT is not meant to function as a mechanical *test* or algorithm for detecting know-how in every case. The account allows that context may supply what counts (in a given task trajectory pursued) as a "non-trivial" segment of that trajectory, for instance; the same goes for which variations are relevant for robustness, and how much stability is "enough", etc. – these can vary with the specifics of the task and the practical purposes of the evaluation, as we should expect. The key idea is the *form* of the condition – what it's tied to; on the GT account manifestation is tied always and everywhere to *controlled, robust progress along a trajectory**, without being tied to the outcome *per se*. With this framework in hand, we can see how GT handles the earlier examples and what explanatory payoff it offers.

6 | DIAGNOSTICS AND EXPLANATORY PAYOFFS

Let's now revisit the illustrative scenarios (interruption, deliberate non-completion, abandonment, etc.) to show how GT yields the intuitively correct verdicts in each where SOC had gotten it wrong. We'll then return to the overarching *scope question*: How much of our behaviour manifests know-how? We will see how and why GT supports a more optimistic answer to this question (as well as the significance of such optimism for our wider theorising about theoretical and practical intelligence norms).

6.1 | The Cases Revisited

Interrupted Risotto, Forgotten Risotto

In **Interrupted Risotto**, Lina (a skilled chef) begins cooking a risotto but is forced to stop midway by an exogenous interruption (e.g., a gas leak).[17] Up to the interruption, suppose she executes, impeccably and intentionally, the characteristic risotto moves: toasting the rice correctly, adding stock in the right increments, stirring at the appropriate cadence, tasting and adjusting. Each such move is plausibly a *robustly progress-making* step toward the distinctive goal state of plating a good risotto: it improves the dish in a way that would remain advantageous under nearby variations, rather than amounting to a lucky coincidence. Moreover, Lina's performance is not merely causally upstream of those good moves; she *controls* them: at each control point she selects and executes the right move under the standing aim of advancing the risotto, in a reasons-responsive way.

On the *Good Trajectories* account, that's enough for manifestation. The episode manifests Lina's know-how to make risotto in virtue of containing a non-trivial *controlled good segment* directed toward the risotto goal. The interruption prevents the segment from continuing, but it does not retroactively cancel what was already on display. By contrast, SOC has no space for this verdict.[18]

Now consider the contrast case **Forgotten Risotto**. Lina begins just as smoothly, but later seems to forget what she is doing: she loses track of where she is in the process and stops in confusion. We saw that SOC cannot discriminate between this case and *Interrupted Risotto*: neither ends in success, so both get the same "no manifestation" verdict. That is already one of SOC's problems. But **Forgotten Risotto** raises an equal and opposite worry for GT: if manifestation depends only on what happens in a prefix, does GT become too *permissive*?

It would be a mistake to treat that worry as arising from a single determinate case. **Forgotten Risotto** admits of two importantly different readings; GT should deliver different verdicts accordingly. On one reading (call it **Forgotten Risotto (systematic)**), Lina's later breakdown is not a freak lapse but suppose a standing limitation she has: at some characteristic point in the task's trajectory—precisely

17. Though any practically forced interruption would serve. The interruption need not undermine the enabling conditions of the performance itself; it may instead arise from independent practical demands, such as an urgent phone call.

18. If manifestation requires success, then no risotto plated means: no manifestation of know-how to make risotto—full stop. SOC can at best re-describe the episode as manifesting some other know-how (boiling rice, stirring, and so on). GT can agree that those more local competences are manifested too; the point is that the broader know-how to make risotto is manifested as well.

where the evolving state of the dish is supposed to *rationally cue* what to do next—she predictably comes up empty. In that setting, the natural diagnosis is that Lina does *not* know how to make risotto *full stop*. What she has, at most, is know-how for initial subroutines (toasting the rice, stirring, beginning the stock-adding routine), or know-how to make something *risotto-like* up to a point.

If that is right, then there is no pressure on GT to say that Lina manifested know-how to make risotto in the prefix. One cannot, after all, manifest what one does not possess. At most, an episode can manifest only the more local know-how the agent actually has. A comparison with *Moonlight Sonata* is apt here. A pianist who reliably plays the opening bars beautifully—a left-hand octave (often $C\sharp$) followed by right-hand triplets outlining the tonic chord ($C\sharp - E - G\sharp$)—but who cannot play the piece through, because the final section defeats them every time, may manifest genuine know-how to play *those bars*, and perhaps to *begin* playing the sonata; but that does not amount to manifesting know-how to play the sonata.

A defender of SOC may object: "But on GT, manifestation is prefix-sensitive; and up to some time t the *systematic* forgetting case can look just like **Interrupted Risotto**. Doesn't prefix invariance force you to treat the prefixes alike—so that you can't deny manifestation here without inconsistency?" The reply is that *prefix invariance is not a principle of superficial behavioural matching*. It says: if two episodes match up to t in all *task-relevant* respects, they should receive the same verdict about what is manifested up to t. But on GT, what is task-relevant includes not only the (mere) kinematics of stirring and ladling, but the *ϕ-directed control* under which those movements are selected and regulated. In the systematic case, the later breakdown is not supposed to be a mere downstream mishap but is rather symptomatic of a standing gap in the control system that is meant to regulate the attempt across control points. In particular, the "tasting and adjusting" that superficially matches a competent performance need not be playing its normal guidance role; it may amount to a scripted routine that does not genuinely integrate feedback into the selection of the next move.[19] If so, then *even up to t* the two episodes differ in a respect that matters on GT: one is an exercise of

19. The point can be put in the terms of the GT machinery. Moves are individuated by their functional role relative to the task goal G_ϕ, and controlled moves require reasons-responsive selection and execution under an intention to advance G_ϕ (or an appropriate subgoal). In a structured task like risotto-making, "tasting and adjusting" is not a detachable flourish: it is supposed to function as feedback that guides what to do next. If Lina's performance predictably ends in genuine confusion at the very point where such feedback should rationally cue the next move, the best diagnosis is that her earlier routine was not operating as risotto-directed route-management in the relevant sense, even if it mimicked the outward profile of a competent start; hence there is no controlled good segment toward plating risotto, and no manifestation of know-how to make risotto.

risotto-directed, reasons-responsive route-management; the other is not.

Consider by analogy: Someone can be excellent at taking exactly one step east and stopping. That bodily movement can coincide physically with the first step of a route to Cincinnati without thereby constituting a move that the agent controls *as part of a Cincinnati-directed trajectory*. Likewise, Lina can perform actions that typically occur early in risotto-making without those actions being controlled as elements of a trajectory *toward plating risotto*, if the relevant goal-directed control is absent. In that sense, the prefixes need not "match perfectly up to t" in the way required to trigger prefix invariance.

On a different reading (call it **Forgotten Risotto (glitchy)**), suppose Lina really does know how to make risotto and is proceeding under normal, risotto-directed control, but then has an unusual, momentary lapse—say, a retrieval failure—so that she cannot bring the next step to mind and loses control. Here GT should give the opposite verdict from the systematic case. Since the episode contains a controlled good segment in its prefix, the episode manifests Lina's know-how to make risotto *in virtue of that segment*. When the lapse occurs, she *ceases* to manifest that know-how: the later part is no longer governed by controlled, reasons-responsive progress toward the risotto goal. But the lapse does not erase the earlier manifestation; it merely marks the point at which manifestation stops. This is exactly the kind of time-sensitive discrimination GT was designed to capture.

Thrown Mate

Thrown Mate was a case with the structure of intentional self-sabotage following a paradigmatically controlled and good segment of a trajectory. In the example case, a chess grandmaster masterfully sets up an inevitable checkmate in a few moves – constructing a *mating net* – but then deliberately refrains from delivering the final checkmate (releasing the opponent from the mating net), out of mercy borne out of a change of heart. According to GT, the grandmaster's *net-construction sequence* is a controlled good segment towards checkmate: each move in the sequence makes non-trivial, robust progress toward the goal of checkmating (e.g. limiting the opponent's king's escape options, tightening the net), and the sequence is controlled (all moves in the trajectory segment are controlled). Thus, that segment of play manifests her know-how to achieve checkmate from the given position. When she later deviates from that path (choosing not to administer the mate), she has then *ceased* to pursue G_ϕ, so the manifestation of *know-how-to-checkmate* ceases at that point. But it *did exist prior to the deviation* – often quite vividly so (as Aron Nimzowitsch memorably put it, "The threat is stronger than the execution"—in chess, mastery is often already on display in the making of the threat). Whereas SOC again we saw forced the wrong verdict (because no check-

mate occurs, SOC would say no know-how was manifested at all), GT explains how the grandmaster's knowledge how to win from the initial position was on display in the lead-up to the mate, even though the opponent was let off the hook.

Tired Climber

Next let's look at abandonment cases which were a plague for SOC. In **Tired Climber**, recall that *S* attempts to summit a mountain but, near the final stage, decides to turn back solely due to fatigue (not because of any mistake or inability - they could have carried on and, had they done so, they'd have made it, easily). Up to the point of turning back, suppose the climber's ascent was executed expertly: she chose an optimal route, maintained a good pace, navigated challenges safely – in other words, her ascent up to the highest point reached was a sequence of *good moves* under intentional control. By GT, that stretch of controlled ascent constitutes a manifestation of *S*'s know-how to summit the mountain, even though she elects not to finish the climb. The later decision to abandon the attempt does not erase the fact that, *before* that decision, *S* was performatively impeccable. In the above way of reading the case, then, we think of Tired Climber as very similar to Thrown Mate: as one of gold standard behaviour up to a point where one intentionally changes policy, one *gives up* controlling of good moves intentionally.

Of course, not every turn-back case is like this. A superficially similar description can be made to support a quite different diagnosis. Suppose instead that the climber turns back because of *mismanagement*: she has exhausted herself early by poor pacing, failed to conserve resources, or otherwise made progress-undermining choices such that, by time *t*, she cannot continue without collapse (and if she tried to continue she would not succeed). Now the later failure is plausibly *explained by* defects in the earlier trajectory. On GT's own machinery, that matters: the relevant early "moves" are not individuated *just by the fact that they are uphill steps*, but by their functional role in advancing G_ϕ. If S's early decisions systematically narrow the margin for success—*by burning through the very resources the standard route requires*—then those decisions thereby fail to qualify as *good* moves toward G_ϕ in the first place. In such a case, GT need not (and should not) treat the early stretch as a controlled good segment directed toward summiting.

However, it would be a mistake to treat "mismanagement" as a single determinate gloss. Here the lesson from **Forgotten Risotto** generalises: one must distinguish at least two readings. Suppose *S regularly* "runs out" at the relevant stage; she does not know how to manage the control points the route demands. On this assumption, then the best diagnosis is that she does not (in the contextually relevant sense) know how to summit this peak at all—she knows at most how to *begin*

such an ascent or to execute subroutines. In that setting, GT is not being permissive; it is simply *out of scope*, because the possession condition for the relevant know-how is not met.

Now, by comparison, (and roughly by analogy with **Forgotten-Risotto-glitchy**) suppose *S* really does know how to summit this peak, and up to some stage she proceeds normally and without a hitch; then an atypical lapse (a rare miscalculation about pace, an unforeseen nutritional problem, a temporary loss of self-monitoring or something to this effect) derails that particular attempt. In that setting, GT should not deny manifestation "all the way up to the abandonment". Rather, it should say: *S manifested* her know-how in the earlier controlled good segment, and then *stopped manifesting* it once control broke down (or once she could no longer sustain summit-directed management). (SOC of course cannot even formulate that mixed verdict, because it has only the endpoint to consult.)

Finally, it's worth noting that there will surely be intermediate cases here: agents who are *just about* above the threshold for counting as knowing how to summit this mountain, but whose route-management is thin, fragile, or intermittently suboptimal. Here it should be unsurprising if there is no sharp "right verdict" to which a theory must conform. The predicate "knows how to summit this peak" itself plausibly admits borderline instances, and in such a region we should expect the corresponding manifestation judgments to be graded or indeterminate rather than all-or-nothing. What matters dialectically is that GT has, unlike SOC, principled resources to reflect that kind of structure: it can say that the episode contains only a short or low-quality controlled good segment (hence a low degree of manifestation), perhaps followed by loss of control.

6.2 | Returning to the Scope Question

We can now return to the sceptical *scope* question that we began with: *How much of our behaviour genuinely manifests know-how?* We are in a position to answer this question, on a principled basis, in a more optimistic way.

SOC implies that all interrupted, abandoned, and "thrown" performances – which make up a large part of our actual real life performative footprint – would automatically count as non-manifestations of know-how (indeed, equally so, with no further distinctions). Since such incomplete or diverted performances happen frequently, SOC implies that genuine manifestations of know-how are actually relatively rare, relative to the baseline of all our performances. It paints a sceptical picture in which much of what we consider skilful, intelligent, competent, etc. activity does *not* officially qualify as displaying know-how at all, unless it reaches a successful conclusion.

By contrast, the Good Trajectories (GT) account yields a far more optimistic verdict. GT allows that *many imperfect or incomplete episodes do manifest know-how* – albeit in a partial or local way. So long as an attempted task contains at least some good trajectory segment of controlled activity, then to that extent the agent's know-how was on display, even if the ultimate goal was never reached. Even a modest "good trajectory" within an unsuccessful performance counts positively. Consequently, under GT the manifestation of know-how is *widespread* in human life, not any kind of exception. Much of our day-to-day intelligent behaviour now qualifies as genuine know-how in motion. This result underscores the *normalcy* of know-how: displaying know-how is not a special rarity but a routine aspect of our conduct.

This anti-skeptical outcome coheres with a broader epistemological principle: that knowledge is the norm of certain activities. Consider the debate over assertion. A widely held view says that knowledge is the norm of assertion (Williamson 1996, 2000; Kelp and Simion 2021; Kelp 2018). This does not mean that every single assertion we make is backed by knowledge; rather, it means that an assertion made without knowledge is improper—defective *qua* the speech act it purports to be. Crucially, the plausibility of this knowledge norm would break down if it were *systematically violated* in practice. If speakers usually asserted claims they did not know (so that the norm was flouted more often than honoured), the norm would hardly function as a norm (see Bicchieri 2005; Bicchieri and Xiao 2009 for more general discussion on this point).[20] And indeed, our conversational practices suggest the opposite. People normally expect assertions to express knowledge. That is why it makes sense to challenge a speaker with *"But you didn't know that!"* (see Williamson 1996, 505) when an assertion turns out to be unwarranted: we treat knowledge as the proper credential for the act of asserting. Such patterns of challenge and retraction are evidence that the knowledge norm is in force in ordinary discourse, and that it is normally respected in our speech.[21]

20. Indeed, there is some doubt as to whether a speech act such as *assertion* would even be made in such circumstances. As Mona Simion and Chris Kelp (2020) note, "one cannot violate too many constitutive rules of an activity too systematically without ceasing to engage in the activity" (2020, 66). The thought is that persistent and wholesale violation of the norm of assertion does not merely yield *bad* assertions; at some point it undermines the very classification of the behaviour as assertion at all. Although it goes beyond what I aim to establish here, a closely parallel point plausibly applies in the case of action. If know-how functions as a norm on action, then sufficiently systematic and flagrant violation of that norm (cf., Maitra 2011) threatens the applicability of the action description itself to the violating activity. At that point, we are no longer dealing with merely incompetent or unlucky action, but with behaviour that falls outside the normal activity of acting for a purpose, and whose gold standard is intentional success. What remains may be (something like) movement, gesturing, or random/exploratory flailing, but not action in the sense that know-how norm would be in the market to regulate.

21. For some criticisms of the knowledge norm of assertion who take knowledge to be too strong an epis-

Parallel considerations apply on the practical side. Many philosophers have argued that knowledge plays a comparably normative role in action and intention (Williamson 1996; Pavese 2024). In the case of propositional knowledge, John Hawthorne and Jason Stanley (2008) have suggested—roughly—that one should *act* only on what one knows. Similarly, Habgood-Coote (2018) holds that know-how is the norm of *intending*: in short, one should intend to ϕ only if one knows how to ϕ. These proposals presume that knowledge (be it knowledge-that or knowledge-how) is the standard case in guiding our behaviour. Acting or intending without the relevant knowledge is seen as a kind of mistake or exception. For such norms to be credible, again, knowledge must be the typical state of the agent, not a rarity.

Here is where the wider theoretical advantage of GT becomes clear. GT's understanding of manifestation vindicates the idea that know-how is (albeit albeit to different degrees) the default across the spectrum of our intelligent conduct. Because GT lets know-how shine through even in many unfinished or unsuccessful attempts, it portrays knowledge-how as pervasively *in motion* throughout what we're doing. Thus it would not be an empty ideal to say that our actions are *normally governed by know-how*. By contrast, SOC's restrictive stance would make a knowledge-how norm on action infeasible: if genuine know-how only in the best cases where we undertake extended (non-basic) tasks, a norm requiring knowledge-how for rational action or intention would be violated left and right. In sum, the Good Trajectories account offers a non-sceptical answer to our scope question. Know-how is *abundantly manifested* in what we do—far more often than the success-only pessimists would allow—and this fits neatly with a general picture of human agents as typically acting in light of what they know.

7 | Objections and Replies

Let's now consider some potential objections and clarify how GT handles them.

Objection 1: What about cases of mere readiness or poised ability, before any action? Sometimes we speak of know-how being manifested *even before the person actually performs an overt action*. For example, a tennis player's ready stance and attentive posture before the serve might themselves seem to display her expertise, *prior* to any particular move. Doesn't GT – with its insistence on a segment of *action* – have trouble accounting for these cases of *poised readiness*, where know-how seems manifest in an agent's preparedness or manner rather than in executed

temic credential, see Lackey (2007), Douven (2006), Weiner (2005). For some worries that it is too weak, see Lackey (2011); cf., Simion (2016).

moves?

Reply: Two points in response. First, the target notion the account is manifestation *in action*. I've thus framed manifestation as an action-involving condition: it requires performance of a controlled good segment. So if truly *no* action is taken, then by definition our account will say know-how (in this strict sense) hasn't been manifested yet. This is by design: we are analysing what it is for know-how to be on display *in doing something*. Thus, it is no surprise (and no defect) that mere disposition or readiness, absent any activity, doesn't meet the condition – that falls outside the scope of "manifesting in action." *Second*, many so-called poised readiness cases arguably do involve *subtle actions* that count as the early moves of the task. For instance, a golfer's shifting of weight, careful looking out for shifts in the wind direction, adjusting of the grip, and so on are all *preparatory actions* that can be construed as part of the larger task's move-space (hitting the ball in the green, say, from 150 yards out). These are minor, sub-intentional moves that "set up" the shot, all of which occur before the player actually swings the club. If such preparatory behaviours are considered part of the task (of hitting it on the green from 150 yards out), then GT can count them as constituting an initial controlled good segment (they are intentional, well selected choices that improve the agent's chances of success once she takes the club back). In short, GT does not require that the *first* move be anything grand – only that *some* controlled, robustly progress-making moves occur. If mere "ready stance" truly involves no action at all, it falls outside our action-focused manifestation notion; but in practice, much of what we interpret as manifested know-how in readiness is already an active shaping of the situation, however small. Those small acts can thus satisfy GT's requirements.

Objection 2: In *Thrown Mate*, the agent intends not to checkmate. How can she manifest know-how to checkmate? The thrown chess game scenario raises a puzzle: the grandmaster never actually *intended to achieve* G_ϕ (mate) – in fact, she intended *not* to deliver mate. How, then, can her activity be said to manifest know-how *to checkmate*? Doesn't manifesting know-how to ϕ require aiming at ϕ? If she was aiming *not to* ϕ, it seems counterintuitive that we credit her with manifesting (to any degree) know-how to do ϕ, no matter how impressive the mating net was.

Reply: The objection trades on an overly strong link between manifestation and intention—one that the Good Trajectories account has independent reason to reject. First, GT is explicitly *local* in that it ties manifestation to the presence, within an episode, of a *non-trivial segment* in which the agent *controls* a *good* trajectory segment *toward* G_ϕ. That local condition is not equivalent to "having a settled intention to achieve G_ϕ". It requires, rather, that in the relevant

segment the agent's control be guided by an intention *to advance* G_ϕ (or an appropriate ϕ-subgoal) at the relevant control points. In a non-basic telic task, intending to secure a ϕ-subgoal—here, for instance, *constructing a mating net*, *restricting escape squares*, *forcing the king onto a terminal file*, *reducing the opponent's defensive resources to make an unpleasant Zugzwang more likely*, etc.—is precisely the kind of intention under which one can intentionally control ϕ-directed moves. Those intentions are *ϕ-directed* in the sense that fixes which moves count as progress-making and which count as blunders or inaccuracies: they are intentions to do things *as steps toward* checkmate, regardless of whether one at a later time intends to take the final step.

Second, there is no incoherence in combining (i) an intention *not* to ϕ with (ii) an intention to perform a sustained stretch of activity that *advances* ϕ. The structure is familiar. One can intend *not* to finish a climb while intending to ascend expertly up to a certain altitude (for training, reconnaissance, or prudence). One can intend *not* to publish a proof while intentionally deriving its key lemmas to test an idea. And likewise, the grandmaster can intend *not* to checkmate their opponent (a hapless friend they take pity on) while intentionally steering play into a position in which checkmate is inevitable *if* one continues with the natural continuation. That is exactly what it is to set up mate while refraining from actually delivering it.

Third, once we register the above distinction, the GT verdict is straightforward. In the net-construction phase of *Thrown Mate*, the grandmaster's moves are (by hypothesis) *robustly progress-making* toward G_ϕ: they systematically tighten the net in a way that would not easily cease to be progress-making under small variations (different, but nearby, defensive replies). And she does not drift into those moves by luck; she *controls* them; she selects and executes them reasons-responsively under the relevant ϕ-directed sub-intentions. So the episode contains a controlled good segment toward G_ϕ, and thus, by GT, *manifestation occurs during that segment*. When she later plays a release move, she stops advancing G_ϕ, and so she stops manifesting know-how to checkmate at that point. But GT's locality blocks the retroactive inference from "she later refuses the mate" to "no know-how to mate was on display earlier".

Finally, notice what the objection would force if it were accepted. It would make manifestation depend on a global intention to complete, so that two agents could match in (i) the quality and robustness of their progress-making moves and (ii) their reasons-responsive control of those moves, yet differ in whether they manifested know-how to checkmate merely because one harboured an overriding intention not to administer the final move. That is the wrong kind of dependence: it treats the exercise of chess mastery in constructing the mate as insufficient for

manifesting the relevant know-how, even when the agent's play is indistinguishable (up to the throw) from paradigmatic checkmating play. GT avoids that mistake entirely by locating manifestation where the real work is getting done: in the controlled, robust progress of the trajectory.

Objection 3: Maybe SOC is true of ability manifestation; why not of know-how? A defender of the success-only view might press: "Perhaps to manifest an *ability* (capacity) it really *is* necessary to succeed (see, for instance, Millar 2009) in doing what the ability is an ability to do; if one has an ability but fails, maybe then that failure means the ability just wasn't manifested (at all). If that is plausible for abilities, why shouldn't the same go for know-how? After all, know-how is closely tied to ability (both for anti-intellectualists, of course, but also for intellectualists for whom knowing how reduces to possessing a kind of propositional knowledge under an ability-entailing practical mode of presentation – see Williamson and Stanley (2001, 429)). So if ability-manifestation is success-bound, know-how manifestation should be too."

Reply: Even if we grant (for the sake of argument[22]) that in some cases manifesting a *basic ability* requires success, it *does not follow* that manifesting know-how does. Know-how, especially for non-basic, extended tasks, is a *richer state* than a bare basic ability – it includes the capacity to intelligently navigate the unfolding task, make decisions, correct for contingencies, etc. That richer capacity can be exercised and displayed *long before* the task is completed. In other words, one can *use* one's know-how effectively throughout an attempt without necessarily getting to the end, as long as the task has an internal structure allowing mid-course skill to show. So when tasks have *structured intermediate stages* at which intelligence can manifest (as non-basic telic tasks do), know-how can be manifested in those stages without final success. So the inference from "perhaps ability manifestation requires success (in simple cases)" to "know-how manifestation requires success (even in complex cases)" is unwarranted. In fact, the argument shows that *in the non-basic telic case, the inference fails**: abilities related to primitive actions might be all-or-nothing, but know-how for extended tasks is *constitutively* something that can shine through in process. Thus, SOC may hold a grain of truth for very simple acts, yet still be false where it matters most – in precisely those extended performances that raise the hardest and most interesting versions of the scope question.

22. Though for criticism of this thesis, see Carter (2021).

8 | Conclusion

How much of what we do counts as manifesting know-how? A *success-only* theorist (SOC) offers a stark answer: only successful performances manifest know-how. We have seen that this answer is implausible, at least for the realm of extended, goal-directed activities. In such tasks, practical intelligence is frequently exhibited (and to differering degrees) in the *management of the route* toward the goal, not merely in crossing the finish line.

The *Good Trajectories* framework developed here preserves what was attractive about SOC – namely, its resistance to attributing know-how across various *bad* cases – but locates manifestation exactly where it belongs in extended action: in the agent's *controlled, robust progress along a trajectory toward the goal*. On this view, many performances that are cut short (by choice or circumstance) or strategically derailed partway *do* often manifest know-how, often in a partial or in ways local to specfiic segments of the overall task. What we get is aa far less sceptical and more realistic answer to the scope question: the footprint of manifested know-how in ordinary life is significantly *wider* than SOC would have us believe. Our know-how is regularly *in motion* well before we cross the finish line – and our theories of practical intelligence should countenance this; when we act (indeed whether we succeed or fail) the manifestation of know-how is entirely normal.

References

Altshuler, Roman. 2025. "Time in Philosophy of Action." In *The Routledge Companion to Philosophy of Time*. Routledge.

Andrada, Gloria. 2024. "Does Know-How Need to Be Autonomous?" *Inquiry* 0 (0): 1–18. https://doi.org/10.1080/0020174X.2024.2309875.

Anscombe, G. 1957. *Intention*. Harvard University Press.

Bengson, John, and Marc A. Moffett. 2011. "Nonpropositional Intellectualism." In *Knowing How*, edited by John Bengson and Marc A. Moffett. Oxford University Press.

Bicchieri, Cristina. 2005. *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge University Press.

Bicchieri, Cristina, and Erte Xiao. 2009. "Do the Right Thing: But Only If Others Do So." *Journal of Behavioral Decision Making* 22 (2): 191–208. https://doi.org/10.1002/bdm.621.

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

Carr, David. 1981. "Knowledge in Practice." *American Philosophical Quarterly* 18 (1): 53–61.

Carter, J Adam. 2022. *Autonomous Knowledge: Radical Enhancement, Autonomy, and the Future of Knowing*. Oxford University Press.

Carter, J. Adam. 2021. "Exercising Abilities." *Synthese* 198 (3): 2495–509. https://doi.org/10.1007/s11229-019-02227-4.

Carter, J. Adam, and Duncan Pritchard. 2015. "Knowledge-How and Cognitive Achievement." *Philosophy and Phenomenological Research* 91 (1): 181–99.
Chase, William G., and Herbert A. Simon. 1973. "Perception in Chess." *Cognitive Psychology* 4 (1): 55–81. https://doi.org/10.1016/0010-0285(73)90004-2.
Chisholm, Roderick M. 1977. *Theory of Knowledge*. Prentice Hall, Englewood Cliffs, NJ,.
Cunningham, J. J. 2022. "Moral Worth and Knowing How to Respond to Reasons." *Philosophy and Phenomenological Research* 105 (2): 385–405. https://doi.org/10.1111/phpr.12825.
Damschen, Gregor. 2009. "Dispositional Knowledge-How Versus Propositional Knowledge-That." In *Debating Dispositions: Issues in Metaphysics, Epistemology and Philosophy of Mind*, edited by Gregor Damschen, Robert Schnepf, and Karsten R. Stüber. De Gruyter.
Danto, Arthur C. 1965. "Basic Actions." *American Philosophical Quarterly* 2 (2): 141–48. https://www.jstor.org/stable/20009161.
Danto, Arthur C. 1973. *Analytical Philosophy of Action*. Cambridge University Press.
Davidson, Donald. 1971. "Agent, Action, and Reason." In *Agent, Action, and Reason*, edited by Robert Williams Binkley, Richard N. Bronaugh, and Ausonio Marras. University of Toronto Press.
Douven, Igor. 2006. "Assertion, Knowledge, and Rational Credibility." *The Philosophical Review* 115 (4): 449–85. https://www.jstor.org/stable/20446923.
Fischer, John Martin, and Mark Ravizza. 2000. *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge university press.
Frankfurt, Harry G. 1997. "The Problem of Action." In *The Philosophy of Action*, edited by Alfred R. Mele. Oxford University Press.
Fridland, Ellen. 2017. "Skill and Motor Control: Intelligence All the Way Down." *Philosophical Studies* 174 (6): 1539–60. https://doi.org/10.1007/s11098-016-0771-7.
Gellner, Ernest. 1951. "Maxims." *Mind* 60 (239): 383–93. https://www.jstor.org/stable/2251324.
Habgood-Coote, Joshua. 2018. "Knowledge-How Is the Norm of Intention." *Philosophical Studies* 175 (7): 1703–27. https://doi.org/10.1007/s11098-017-0931-4.
Hartland-Swann, John. 1956. "The Logical Status of 'Knowing That'." *Analysis* 16 (5): 111–15.
Hawley, Katherine. 2003. "Success and Knowledge-How." *American Philosophical Quarterly* 40 (1): 19–31. https://www.jstor.org/stable/20010094.
Hawthorne, John, and Jason Stanley. 2008. "Knowledge and Action." *The Journal of Philosophy* 105 (10): 571–90.
Hornsby, Jennifer. 2010. "Trying to Act." In *A Companion to the Philosophy of Action*. John Wiley & Sons, Ltd. https://doi.org/10.1002/9781444323528.ch3.
Kearl, Timothy. 2023. "What We Know When We Act." *Philosophical Studies* 180 (9): 2665–83. https://doi.org/10.1007/s11098-023-01997-5.
Kelley, Mikayla. 2022. "How to Perform a Nonbasic Action." *Noûs*, ahead of print. https://doi.org/10.1111/nous.12440.
Kelp, Christoph. 2018. "Assertion: A Function First Account." *Noûs* 52 (2): 411–42.
Kelp, Christoph, and Mona Simion. 2021. *Sharing Knowledge: A Functionalist Account of Assertion*. Cambridge University Press.

Lackey, Jennifer. 2007. "Norms of Assertion." *Noûs* 41 (4): 594–626.
Lackey, Jennifer. 2011. "Assertion and Isolated Second-Hand Knowledge." In *Assertion: New Philosophical Essays*, edited by Jessica Brown and Herman Cappelen. Oxford University Press.
Lavin, Douglas. 2013. "Must There Be Basic Action? *." *Noûs* 47 (2): 273–301.
Maitra, Ishani. 2011. "Assertion, Norms, and Games." In *Assertion: New Philosophical Essays*, edited by Jessica Brown and Herman Cappelen. Oxford University Press.
Millar, Alan. 2009. "What Is It That Cognitive Abilities Are Abilities to Do?" *Acta Analytica* 24 (4): 223.
Miller, George A. 1956. "The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information." *Psychological Review* (US) 63 (2): 81–97. https://doi.org/10.1037/h0043158.
Miracchi, Lisa, and J. Adam Carter. 2022. "Refitting the Mirrors: On Structural Analogies in Epistemology and Action Theory." *Synthese*, ahead of print. https://doi.org/10.1007/s11229-022-03462-y.
Pavese, Carlotta. 2017. "Know-How and Gradability." *The Philosophical Review* 126 (3): 345–83.
Pavese, Carlotta. 2020. "Probabilistic Knowledge in Action." *Analysis* 80 (2): 342–56. https://doi.org/10.1093/analys/anz094.
Pavese, Carlotta. 2024. "Factive Mindreading in the Folk." *Putting Knowledge to Work: New Directions for Knowledge-First Epistemology*, 145.
Pritchard, Duncan. 2005. *Epistemic Luck*. Clarendon Press.
Roland, Jane. 1958. "On" Knowing How" and" Knowing That"." *The Philosophical Review* 67 (3): 379–88.
Ryle, Gilbert. 1949. *The Concept of Mind*. Hutchinson & Co.
Setiya, Kieran. 2014. "The Midlife Crisis." *Philosopher's Imprint* (Ann Arbor, MI) 14 (31). http://hdl.handle.net/2027/spo.3521354.0014.031.
Shepherd, Joshua. 2014. "The Contours of Control." *Philosophical Studies* 170 (3): 395–411. https://doi.org/10.1007/s11098-013-0236-1.
Simion, Mona. 2016. "Assertion: Knowledge Is Enough." *Synthese* 193 (10): 3041–56. https://doi.org/10.1007/s11229-015-0914-y.
Simion, Mona, and Christoph Kelp. 2020. "Assertion: The Constitutive Norms View." In *The Oxford Handbook of Assertion*, edited by Sanford Goldberg. Oxford University Press.
Sosa, Ernest. 1980. "The Raft and the Pyramid: Coherence Versus Foundations in the Theory of Knowledge." *Midwest Studies in Philosophy* 5: 3–25.
Sosa, Ernest. 2015. *Judgment & Agency*. Oxford University Press UK.
Weiner, Matthew. 2005. "Must We Know What We Say?" *The Philosophical Review* 114 (2): 227–51. https://www.jstor.org/stable/30043667.
Williamson, Timothy. 1996. "Knowing and Asserting." *The Philosophical Review* 105 (4): 489–523.
Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford University Press.
Williamson, Timothy. 2017. "Acting on Knowledge." In *Knowledge First: Approaches in Epistemology and Mind*, edited by J. Adam Carter, Emma C. Gordon, and Benjamin W. Jarvis. Oxford University Press Oxford.
Williamson, Timothy, and Jason Stanley. 2001. "Knowing How." *The Journal of Philosophy* 98 (8).