

Good Moves
J. Adam Carter

— Draft. September 2026.

Abstract. The topic of this paper is telic know-how. Telic know-how ranges over tasks defined by a clear terminus. It is argued that telic know-how is a stable, trainable disposition to select and sequence good moves across the live contexts the (telic) task presents, viz., those at which the agent's choice still bears on the outcome. A move is good when it issues from a policy the agent herself implements that robustly improves the prospects of completion over entry-level engagement, and no available alternative clearly betters it. Possession of the disposition is partly a matter of its history. It must have been formed or maintained through sources calibrated to the facts that make moves good. Propositional knowledge and motor skill enter the disposition as components, in proportions that vary with the task, and in particular with its *variance*, i.e., how far what works on one occasion carries over to the next attempt at a task of the same type. So the question intellectualists and anti-intellectualists have contested, which of these know-how really is, has no general answer. Because the disposition is manifested in good moves rather than in just securing the tasks's terminus, an agent can exercise know-how throughout stretches of an unsuccessful attempt. The account accommodates the embedded-question semantics for know-how ascriptions by treating policies as answers to practical questions, and good policies as the correct ones.

1. Consider the following case:

SENATE: It's midnight in the 'war room' for a candidate running in a US Senate race in a swing state. The race is tied in the polls, and a complicated scandal (involving your candidate!) has just broken in the past hour. The normal playbook options here are all rather bad ones. Responding with a statement in the early stages is likely going to just amplify the scandal. There's a risk though that ignoring it effectively concedes the worrying framing that's now spreading on social media. So what should the campaign do? There is no straightforward recipe here to follow. Whatever might work tonight could easily backfire as the news cycle develops and political relationships are tested.

It's in situations like this that highly skilled campaign managers are sought out. We say the very best of them know how to win close races, and what the money buys is the quality of their calls at just these kinds of choice points, the ones no playbook covers. Now, suppose your candidate has gone all in and has retained such a strategist, Rivera, who – through decades of experience – knows how to do what it takes to win even on nights like this one.¹

¹A reader may balk at 'knows how to win'. Surely Rivera knows how to campaign, and winning isn't up to her? §7 argues for the win-directed ascription. (Until then I'll

What have we just attributed to Rivera, in saying she knows how to win a close campaign race? The question of course generalises. We're presumably attributing the same *kind* of thing, "know-how", that we attribute to Albert Pujols catching a fly ball at first base, as well as to a young child who's just learned to hang their coat on a hook.

Note that all these tasks, despite their differences in terms of difficulty, are tasks that proceed in time toward a terminus that is either attained or not (the campaign is won or lost, the child's coat ends up on the hook or it doesn't, etc.). Call tasks like these *telic*: tasks whose completion licenses 'I have ϕ -ed' (for the term, [Comrie, 1976, 44–48](#)). They are unlike other things we know how to do, like "being a compassionate friend" or "keeping up a garden" (see [Setiya, 2014, 12–13](#)), which, being *atelic* (they don't proceed in time towards a terminus and so aren't ever "completed" at any given point), answer instead to open-ended norms of sustainment.² My subject in what follows is what *telic know-how* is (and, hereafter, 'know-how' means telic know-how.)

Telic tasks differ significantly amongst themselves when it comes to how far what works on a given occasion carries over to the next, where we're thinking of "what works" as what brings one closer to the terminus. The SENATE case sits at one pole of that dimension. Compare now a "recipe-style" case at the other.

RISOTTO: Lina is a cook preparing mushroom risotto. She has made it hundreds of times, and it goes much the same way each time: the heat stays where she set it, the stock is as salty as it always is, and the rice absorbs it at the rate rice always does. Small things go wrong and she puts them right on the spot (e.g., the pan gets dry, so she adds a ladle of stock), and she finds out straight away whether she judged it right, because she keeps tasting.

It's true of a lot of people that they know how to make risotto, and a reason for this is that the things you do that make for good risotto today make for good risotto tomorrow, for whoever does them, in whatever kitchen. Notice too that each state the dish can be in has a standard fix (rice still chalky, so give it a few more minutes), so a quick taste needs to just tell her which one it is in, and then the next move is basically settled. What she's left with after learning how to make risotto, then, is a *way* of making it, in the sense that will matter in §2: roughly for now, a procedure that tells her what to do at each state the dish can be in, and that she can keep and reapply because it stays a way of making good risotto from one attempt to the next.

take attribution practices at face value. Donors, candidates, etc. sort prospective campaign managers by whether they know how to win close ones.)

²The account may extend to atelic tasks with adjustments (a sustained pattern in place of a terminus). My official theses are restricted to telic ones, where, as Wiggins puts it of the snooker player choosing a shot, 'there is no problem about the end itself' ([1976, 43](#)).

SENATE is the reverse on essentially every count here. What works shifts pretty quickly with the news cycle. And using a tactic can itself be the very thing that makes it stop working later, since the other side is watching and learns how to answer it. Beyond all that, the feedback is not only delayed (the polls that would tell you whether last night's call worked may lag a week), but also often misleading (a "wicked" environment in Hogarth's terms (2001, 89)). I'll call tasks at Lina's pole *low-variance* and tasks at Rivera's pole *high-variance* (what varies is what works at a given choice point, not how often one succeeds).

High-variance tasks, it's worth noting, aren't at all peculiar to managing Senate races. They actually cover a lot of the know-how we value most, such as emergency medicine, litigation, and indeed much of what we have rigorous processes to certify people for. Medical education, for one, now distinguishes *adaptive* expertise from routine expertise, where testing for adaptive expertise involves testing for the capacity to handle cases the protocol does not cover, and treats it as an explicit target of training rather than something that simply gets left to experience (Cutrer et al., 2017, 70; Mylopoulos and Regehr, 2007).³

2. On the matter of what kind of state telic know-how is, two families of answers have dominated the option space. One says that knowing how to do something is a matter of knowing a truth about a way for one to do it, where this is a proposition grasped in some distinctively practical way (e.g., under a practical mode of presentation) (Pavese, 2015, 2017, 2021, 2026; Stanley, 2011b, 2011a; Williamson and Stanley, 2001). The other says it's fundamentally a matter of possessing an ability, viz., a (refined and suitably embodied) disposition to succeed (Carter and Pritchard, 2015; Dreyfus, 2005; Hornsby, 2012; Löwenstein, 2017; Markie, 2015; Noë, 2005; Ryle, 2009).

This debate, which is the one I want to intervene in, has been active broadly for seventy years, the last twenty-five or so intensely, and is more or less stuck. Asking *why* is to kick a (partly sociological) hornet's nest, but here is one reason among several.

Both the intellectualist and anti-intellectualist explain – albeit very differently – what kind of state constitutes know-how in terms of the agent's relation to a way, or ways, of succeeding at the task, if not always explicitly. What's deeply contested is whether an agent who knows how to ϕ (as on the former approach) thereby *knows* of some way *that* it's a way for her to successfully ϕ , or whether (on the latter approach) she's *able* to ϕ by means of some way she has (well enough) mastered.⁴

³High-variance tasks are also some of Ryle's own paradigms (e.g., take his mountaineer, who 'is ready for emergencies' (2009, 30)). As he puts it, those dispositions he's 'largely concerned' with are those 'the exercises of which are indefinitely heterogeneous' (2009, 32). The literature that followed him has worked instead with the simpler end of his examples (the light bulb, the lifejacket, the bicycle, etc.).

⁴Even though 'ways' are more canonically associated with intellectualism, this kind of talk

(Does Sally know, of a certain fingering, *that* it is a way for her to play the scale, or *is she able* to play the scale by means of such a fingering?) Most tasks leave various ways of succeeding open at a given time, and both camps quantify existentially over them, so nothing here assumes that a task has a single canonical way.

What matters is that, on both views, a way to successfully ϕ must combine two properties if a subject's relation to it, whether it be knowledge of it or dispositional mastery of it, is to be in the market for constituting know-how. Call the properties *guidance* and *durability*:

Guidance. A way w of ϕ -ing is *guiding* for an agent just in case her standing in the relevant relation to w – knowing it (intellectualism), or having mastered it (anti-intellectualism) – puts her in a position to act in the sense that: given where she now stands in the task, w settles, or at least substantially narrows, what she is to do next.

Durability. A way w of ϕ -ing is *durable* just in case w remains a way of succeeding at the task from one occasion of undertaking it to the next, so that the relation to w an agent forms on one occasion is still a relation to a way of succeeding on later ones.

Guidance is a matter of a way being specific enough to act on. An unambiguous, complete recipe sits at one end. Consider that this would settle very nearly everything to do across the task's choice points (if the rice is dry, add stock, etc.). A rule of thumb or a checklist *narrows* what to do next without actually settling it, and narrowing is enough. 'Make something delicious' narrows (nearly) nothing at all. It just names an end and leaves the way to do it open. (Compare the two-step guide to drawing an owl that circulates online: step one, 'draw two circles'; step two, 'draw the rest of the owl'. The second step is pure end and no way.)

Durability is what makes a way retainable. Know-how is a *standing state* after all. It's kept from attempt occasion to attempt occasion (one can have, e.g., known how to land a plane since the 1970s) – so its object (that is, the way that's known or mastered) has to "hold still" long enough to be kept.⁵

Here's why both properties are needed on either view. Suppose knowing how to complete a telic task T consists in a relation to a way of succeeding at T . On the intellectualist's account, knowing of that way *that* it's a way (for you) to succeed at

crops up regularly on the anti-intellectualist's side. Markie's (2015, 3194) special ability view takes know-how to be constituted by "dispositions to behave in certain ways in order to ϕ ... [where] those behaviors are effective ways to ϕ in the circumstances", and Löwenstein's (2017) Rylean competence is exercised in choosing "the best way available" (Löwenstein, 2017, 283). For other examples, see Glick (2012, 126), Rosefeldt (2004, 374) and Hornsby (2017, 89).

⁵Holding still doesn't require that the way tell you the same thing on every occasion. It can be a standing procedure whose instructions depend on the circumstances ('if the polls move, move the ad money' etc.), so long as the procedure *itself* stays fixed. §4 will call such a procedure a *policy*.

T , on the anti-intellectualist's, *being able* to succeed at T by means of it. If the way lacks guidance, then knowing it or having mastered it wouldn't be enough to tell the agent what to do next, so a relation to it wouldn't suffice for know-how. If the way lacks durability, then it stops being a way to succeed suitably cross-situationally from one occasion to the next, so no standing state could be a relation to it. So if a subject's (propositional or dispositional) relation to a way is to constitute know-how, the way must have both the properties of durability and guidance.

An account of know-how therefore faces a choice point. It can look for ways that have both properties, which is the natural reading of both camps, since each explains know-how as a relation (propositional or dispositional) to a way of ϕ -ing. Or it can give up the presumption that know-how to ϕ , in full generality, is to be explained by intellectual or behavioural relations to ways of ϕ -ing at all, and begin somewhere else. §3 lays out why the first option fails once we focus on not just low-variance but also high-variance tasks; the second is the one this paper takes.

3. At the low-variance pole, guidance and durability never come apart. The way one learns to change a light bulb or put on a lifejacket (and so forth – like the rest of the debate's stock examples) both settles what to do next at each choice point and remains a way of doing the task from one occasion to the next (Hawley, 2003, 20; Snowdon, 2004, 12; Stanley, 2011a). This means both camps' proposals are going to predictably fit such examples equally well. The intellectualist can say the agent's know-how is her knowing that way (a way exhibiting both guidance and durability) and the anti-intellectualist that it is her having mastered that same way, and so nothing in a low-variance case really favours one over the other. That's a recipe for stalemate.

In (much) less-discussed high-variance tasks, though, the two properties come apart, and, as we'll see, not even the best versions of intellectualism and anti-intellectualism can hold them together without an unexplained remainder.

3(a). So the intellectualist view, applied to SENATE, says that Rivera, who knows how to win this race, knows some way-proposition that, say, a political television pundit (fluent in analysing strategy but who's unable to execute any of it) doesn't. Now here's a question that generates a dilemma for the intellectualist. What happens as we vary *how detailed* w is?

Suppose w is a full plan, the political analogue of a recipe. Filled out on the night of the scandal, it might run something like this (call Rivera's plan Plan-R):

Plan-R: (a) Say nothing tonight, then (b) put the candidate on local radio at seven with a two-line apology; (c) release the internal timeline at noon and (d) move the rest of the week's ad money onto the opponent's tax record; then (e) use the debate a fortnight out to pivot from the scandal to the economy, and (f) run a closing ad in the final week

on the opponent's record while then (g) spending the last seventy-two hours on turnout in the three counties that decide the state.

Suppose Rivera believes that **Plan-R** is a way for her (qua campaign manager) to win. Does Rivera *know* this is a way for her to win? Only if her belief that "Plan-R is a way (for Rivera) to win" couldn't easily have been false, since knowledge requires *safety* (or some anti-luck condition of that kind).⁶ But it easily could have been.

Had the opponent got the timeline out first, or a second story broken at five, etc., the apology would have landed as damage control and Plan-R would no longer have been even likely to win, so the way-proposition would have been false. And those worlds are very close. So however good her judgement is, then, Rivera *can't* know that a plan this specific, even if (at midnight, as the scandal breaks) it's her very best plan, is a way for her to win. If way-knowledge requires this level of detail, the view is accordingly going to significantly under-ascribe know-how. And although at least *tonight* her specific plan does just fine by the lights of *guidance*, most of the plan, in full detail, will be stale by next week (indeed, likely within a day), so a plan of that specificity is just not the kind of thing a standing state such as know-how could be a relation to. What fails here is *durability*.

So to get around that, an intellectualist will presumably need to coarsen up *w* to something less specific in its substantive detail, like a rule of thumb. Something like:

Plan R*: (a) Read the polls, (b) drop spoiled messages, (c) avoid risky gambits when conditions turn.

Now the proposition that **Plan R*** is a way to win is modally stable, because the rule is context-sensitive. But it's also, as it were, *too easy* to know. An amateur can truly and safely believe that "Read the polls, drop spoiled messages, and avoid risky gambits when conditions turn" is a way to win without thereby knowing how to win. A way this general supplies at most *deontic* knowledge-how (knowing how one *ought* to go about winning, or of what, according to the rules, would suffice for winning). Both sides of the intellectualism/anti-intellectualism divide agree deontic know-how is boring old propositional knowledge and that it's not what's at issue between them. What's at issue is the *non-deontic* reading, knowing how to win (Carter and Poston, 2018, 116; Williamson and Stanley, 2001, 423–425).

The intellectualist's account of what distinguishes Rivera from the pundit, who also knows that the coarse-grained rule is a way to win, has two components that work together. One is the *practical mode of presentation* (Glick, 2015; Pavese, 2015, 2019; Stanley, 2011a; Williamson and Stanley, 2001). The pundit and Rivera know the same rule encapsulated in **Plan R***, but only she knows it under a practical mode (i.e. roughly, she thinks of the way as a way for her to act on, rather than under

⁶See, e.g., Pritchard (2005).

the demonstrative or descriptive guise she would get from watching or reading about it), and such a mode, according to mainstream intellectualism, ‘entails the possession of certain complex dispositions’ (Williamson and Stanley, 2001, 428–429), which Rivera has and the pundit lacks. The other is the success-indexing of the content, as in Pavese:⁷ ψ -ing is ‘a way for one to ϕ just in case it is sufficiently likely for one to succeed at ϕ -ing by ψ -ing’ (Pavese, 2021, 1601), so what Rivera knows is that **Plan R*** is a way for *her* to win, and the pundit’s corresponding belief about himself is false, since attempting to follow **Plan R*** wouldn’t make success (for *him*) likely.⁸

But ask *why* the content is true of Rivera and false of the pundit. The content ‘*w* is a way for me to win’ (where *w* is the way encapsulated in **Plan R***) comes out true in her mouth and false in his because she has the ability that the practical mode entails and he lacks it. Pavese says as much, namely, that the ability ‘is built into the knowledge component *because* it is built into the practical mode of presentation component’ (Pavese, 2017, 375, *my italics*; cf. Pavese, 2026, 212, 276). So building the ability into the content of way-knowledge is what stops the view from crediting the pundit with know-how. But notice that the way-knowledge is now true in virtue of the ability, so if we want to know what separates Rivera from the pundit, it’s the ability we need a theory of. More widely, the situation seems to be that intellectualism secures the durability desideratum on an account of know-how by going for a coarse rule rather than the fine-grained plan, and secures guidance through an ability of the subject’s that it doesn’t analyse, or at any rate not in terms of propositional knowledge.

Here’s a second hitch. The evidence that practical modes are psychologically real comes from tasks with a fixed decomposition into operations the agent’s stored routines execute directly (‘reach, grab, twist’), namely control theories of motor behaviour and procedural-memory models of cognitive skills such as counting and algebra (Pavese, 2015, 13–16; 2019, 792–801). Such routines are all the mode’s ‘complex dispositions’ (Williamson and Stanley, 2001, 429) really come to on that evidence. A high-variance task like winning a Senate race really has no such straightforward decomposition, though, so what does the selecting there isn’t very plausibly going to be a routine of either kind (whether it be motor or procedural)

⁷Pavese makes the success condition explicit, but it is not peculiar to her. On Stanley and Williamson’s semantics the way-proposition is that *w* is a way in which one *could* ϕ , and the modal is read as counterfactual success in the relevant conditions (Stanley, 2011a, 127; Williamson and Stanley, 2001, 426; cf. Hawley, 2003).

⁸Why not? Because a rule this coarse leaves its own *application* open, and applying it is where the pundit is going to fail. After all, something like ‘Drop spoiled messages’ doesn’t say which message has spoiled, etc.

but (again) something more akin to an ability.^{9,10}

3(b). The intellectualist's view (in high-variance cases) started showing signs of problems when we asked how specific the way needs to be. For the anti-intellectualist, we'll start by asking this: when one has mastered a way of ϕ -ing by developing some refined ability, what is the ability an ability to do, exactly? If the task is playing a musical scale, the answer's easy. For a volatile campaign the terminus, winning, is definite (modulo contested elections!), but the means are settled only as events unfold, so the most an ability-to-win can be (when one has this at the start of a campaign) is a tendency to reach that terminus by whatever means the race turns out to require.

A tendency to *succeed* won't do the job here, because it is fixed by outcomes alone. A lucky dilettante strategist (who, unbeknownst to them, is being helped out by a super PAC) and Rivera might share a tendency to win, even though everything that makes Rivera's conduct of the campaign intelligent (e.g., how she avoids falling into traps) plays no part in fixing that tendency, since two agents who win equally often trivially have the same tendency to win, however differently they got their wins. Nor, for that matter, can a bare tendency to succeed register the difference between Rivera and a third character, Bengson and Moffett's (2011) Irina, an ice skater who is seriously mistaken about how a salchow is done (she thinks one takes off from the front outside edge rather than the back inside edge) but whose neurological abnormality makes her perform the correct sequence whenever she

⁹Stanley and Williamson's later account of skill makes the same kind of move in its own terms. 'To be skilled at the action type of ϕ -ing is to be disposed to form knowledge appropriate for guiding tokens of ϕ -ing' (2017, 715). What makes formed knowledge *appropriate for guiding* doesn't get unpacked (and in a high-variance task that is the part that needs analysing, since there the guiding knowledge has to be formed afresh at each choice point.)

¹⁰Note that this second hitch doesn't depend on how specific the known content is. This is so even if the content is the kind of adaptive policy §4 will introduce (Waights Hickman (2019) comes closest to that, with knowing a method in what Ryle (2009, 30) called 'the executive way'). This is because what makes the knowledge "practical" is the *mode*, and, as noted above, the evidence supporting modes (as disposition-entailing) comes from tasks with a fixed decomposition (Pavese, 2015, 13–16; 2019, 792–801). Pavese's move in response to this problem is to maintain that what the agent stores is a *motor schema*, viz., a control program that holds what stays the same across performances of the action and leaves the rest to be filled in on the occasion (2019, 810). This is an improvement, but it doesn't help in a high-variance task. A schema stores what stays the same across task types. (So, in winning a campaign, what stays the same from one night to the next is going to be something thin like a rule of thumb and nothing more.) Now, what makes tonight's move a good one is a matter of tonight's particulars (which message has just gone out of date, etc.), and those are what changes from night to night. So they are not among what's stored in the schema. They are *what the schema leaves to be filled in* when the occasion comes. Bearing this in mind, then, it looks like, for the intellectualist, either the practical mode is given computational substance, albeit, something thin that doesn't separate Rivera from the pundit, or it isn't given such substance, and the work is done by whatever is filling in the schema, which is an ability we don't yet have an account of.

tries to perform a salchow, so that she reliably lands the jump without knowing how to do it (Bengson and Moffett, 2011, 171–172; Bengson, Moffett and Wright, 2009, 392–393). What separates Rivera from the dilettante and from Irina (as well as from a campaign-strategist variation on Irina), to put it roughly for now, looks like it has to be something about *how she chooses*, not how often she succeeds when she tries.

The strongest current versions of anti-intellectualism are built against such worries. Löwenstein's (2017) answer is that know-how is an ability one has in virtue of *understanding* the activity, where the understanding is itself a competence – specifically, a capacity to assess one's options and performances with reference to what doing well in that activity requires (2017, 276, 282–283). That's enough to deal with Irina. She's just reliable when she tries, whereas someone with bona fide know-how chooses the best option 'because one understands that this is what one ought to do' (283), something Irina doesn't do. Let's grant this line for the sake of argument and then consider what the understanding comes to. In RISOTTO, what doing well requires is settled and widely known (the stock goes in gradually, the rice should end up creamy), so the understanding Löwenstein appeals to comes cheap. Anyone who has learned the recipe has it. In high-variance cases like SENATE it's the reverse. What doing well requires tonight isn't settled or widely known, so the competence Löwenstein calls understanding is doing all the explaining. His account says what that competence is for, assessing options against the standards of the activity, but not what it consists in when those standards have to be worked out on the night, and that is what a theory of Rivera's know-how would have to give us.

And, zooming out, guidance and durability come apart for the anti-intellectualist much as they did for the intellectualist. That is to say, the way Rivera has mastered here is a general one (read the polls, drop spoiled messages, etc.), which is why it can be kept from race to race (indeed, as Rivera continues to get hired by new candidates each election cycle), and the guidance comes from something in Rivera (for Löwenstein, her understanding), which, in a high-variance task like SENATE, is what separates her from the super PAC-backed dilettante and from an Irina-style strategist, and is what the account leaves effectively unexplained.

3(c). Now a quick recap. We saw that, for a task like SENATE, no way either camp offers exhibits both guidance and durability at once. We also saw that the sophisticated versions of both views make effectively the same move in response (in trying to have both). They (i) keep the relevant way sufficiently general (which is needed for achieving the durability desideratum) and end up (ii) locating guidance in some further capacity of the agent that settles what to do next where the general way doesn't. The two camps describe this capacity differently, but they agree on its role, which is *to keep finding what to do next as the situation changes*, as we saw Rivera

does and the pundit, the dilettante, and Irina don't. But neither camp says what actually makes a call a good one (whether or not it leads to success), or what it takes to have the state that reliably produces good calls, and those are the two things a theory of the capacity would have to say.

So the question that needs answering isn't how to characterise the relation to some way of succeeding that constitutes the agent's know-how (viz., as knowing it or as having dispositionally mastered it), but rather (and to a first approximation that §4 will sharpen) what it takes to have a capacity to keep finding what to do next as the situation changes across the contexts the task presents, and, relatedly, whether having that capacity consists in knowing something or in being able to do something. §§4-6 answer the first question. The answer to the second, on both counts, is already in hand from 3(a) and 3(b), and a note in 3(a) shows that the intellectualist's side of it survives whatever content is chosen for the known way, the adaptive policy of §2's note included. §7 gives the semantics.

4. So what on earth is the strategist doing all night long (along with drinking coffee)? At any given point, she's *selecting what to do next* (we'll say more shortly), which, for now, means only picking what to do from the options the situation leaves open. An expert like Rivera is someone who goes on selecting well as the race unfolds. That is, she keeps picking good options *even as what counts as a good option changes from one choice point to the next*.

In light of this, two questions present themselves:

The goodness question. Take a telic task, a context *c* that an agent is in while engaged in it, and the moves available to her at *c*. What makes one of those moves a *good* move at *c*, given that no move guarantees success and that which moves are good varies from context to context?

The possession question. What is it for an agent to *have* a state that reliably yields good moves across the contexts the task presents, rather than to produce good moves by luck or with outside help?

Know-how, I'm going to argue, is the state that answers the second – to a first approximation, a stable, trainable disposition to *select and sequence good moves*. This section answers the goodness question. §§5 and 6 answer the possession question.

4(a). Answering the goodness question takes some terminology. A *task*, as I use the term, is a goal-directed activity, which is not the same thing as its end (e.g., in SENATE, winning is the end whereas the campaign itself is the task). Tasks are typically pursued by way of sub-goals. They also come with constraints on what counts as an admissible way of pursuing the end (e.g., you can't make illegal moves to win at chess, and not running anyone over in your car is a commonsense constraint on the task of getting to the store.) A telic task (as opposed to atelic tasks)

carries a *terminal* success condition, *G*. These again are our focus. I'll use *context* to refer to the total situation an agent can be in while engaged in the task.

The context includes her own physical and mental condition, the relevant state of the environment (and of other agents) and as much history as bears on her options. The goodness question asks what makes a single selection of what to do next good, so it needs a unit of selection, namely the thing an agent picks, from the options available to her, at a choice point. Whole actions like 'run the campaign' are too big for that, since they are made up of many choices spread over time, and there is no single point at which one picks the whole thing (one can decide at the outset to run the campaign, but that decision settles none of the choices within it). Bodily motions are too small, since the agent does not pick them at all (e.g., nobody chooses which muscle fibres to contract in order to make a phone call). The unit I need is in between, and I will call it the *move*.

Move. A move is a unit of choice, meeting three conditions. It is (i) *available*, viz., within the agent's own repertoire ('dunk the basketball' is a move for LeBron, not for me). It is (ii) *at the level of choice*, the finest level the agent's psychology can directly select ('hit a cross-court forehand' is a move; 'contract muscle fibre 42' is not, being sub-personal). And it is (iii) *individuated by its function in the task*: raising a hand is one move in a classroom and another at an auction.

Two brief comments. The relevant level of choice that matters in (ii) is fixed by the agent's actual trained psychology as opposed to what's discriminable in principle. A concert pianist, for example, registers alternatives (voicings, fingerings) that a beginner can't treat as options at all (see Pacherie (2008) and Christensen, Sutton, and McIlwain (2016) on the layering of skilled control). And, further, notice that the definition of a 'move' I'm using doesn't say anything about intention or control, which is deliberate. A definition of 'move' in terms of intentional control would make the account of good moves presuppose the intelligent control it is meant to explain.

4(b). The move as a unit of assessment has a precedent in Ryle (2009). Ryle's drunkard at the chessboard makes one move that wrecks his opponent's plan, and (in his example) the spectators put it down to luck because his other moves show 'no tactical connection with the position of the game' (2009, 34). Ryle leaves the point there.

So what does make a move *good*, asking this now not just for chess but in full generality, for any telic task? It can't be that it produces success by securing the task's *terminal* success condition, *G*. In SENATE even the best move of the night does not such thing. My proposal (put roughly now, and made precise as we go) is that a good move is one that comes from a good way of choosing moves, and a

good way of choosing moves is one that robustly lowers the risk of failure. Each part of this needs spelling out, and in particular: (i) lowers risk *relative to what*, and (ii) robustly *in what sense*?

Let's start with (i); (ii), the sense in which the lowering must be robust, is taken up in 4(c). First the way of choosing needs a name:

Policy. A policy is a way of choosing what to do next: a function that, for each context the task can present, selects a move from the options available there.

The floor defined below is itself a policy (namely, the entry-level one). In chess, Magnus Carlsen's policy for winning a Sicilian as white is (much) better than the entry-level one. Policies also compose. Running a campaign breaks into sub-tasks. Since at any point in the campaign she is engaged in one of these sub-tasks, the sub-task policies between them settle what she does at every context the campaign presents, which is to say that they make up a single policy for the whole campaign.

With that said, to say a policy *lowers the risk* of failure at its task involves comparing two chances: the chance of success (i.e., securing the telic terminus) if the task is continued from here as the policy selects, and the chance if it were continued some other way. So the 'other way' has to be fixed.

What's needed then is a *baseline continuation* of the task for the policy to be compared with. It's tempting to just compare it with doing nothing. But that's wrong. In most tasks nearly anything beats doing nothing (castling, when it is a good move, is not good in virtue of beating the alternative of sitting there frozen, or meditating for infinity, etc.).

The baseline we need is what I will call the task's *floor*:

Floor. A telic task's *floor* is entry-level engagement with it, that is, continuation of the task by an agent who grasps its end, its constitutive rules, and its admissible operations, but brings no special training or experience beyond that.

In chess, for instance, the floor is (something like) the novice's game – viz., legal play, aimed at winning, alert to the obvious capture and the obvious threat, but with none of the refinements that proper training adds. So, no opening repertoire, no tactical eye, etc. The floor is what supplies the baseline continuation we were after (the 'other way' the task might be continued, against which a policy's lowering of the risk of failure gets measured). It's the least an agent can bring and still count as undertaking the task, so it is fixed by the task itself (e.g., by its constitutive norms) rather than by anyone's choice. (Anything less is not engaging with the task – as Christoph Kelp (2021, 366) puts it, "throwing chess pieces around does not qualify as playing chess", and anything more is already some of what training adds, which

is effectively the thing here to be measured.)¹¹

Three further bits of terminology help make this more precise. The first (a ‘prospect’) is a quantity the other two are built from:

Prospect. The *prospect* at a context c under a policy π , $P(G \text{ given } c, \pi)$, is the probability that the task ends in success if play from c continues as π selects. Since the floor is a policy, $P(G \text{ given } c, \text{floor})$ is the prospect under entry-level continuation; and $P(G \text{ given } c, m, \pi)$ is the prospect if move m is made at c and play then continues as π selects.

So, from a given chess position, the prospect under Carlsen’s policy is the chance that the game is won if it is played on from there as Carlsen would play it, and the prospect under the floor is the chance that it is won if a novice plays on instead.

Improvement. A policy’s *improvement on the floor* at c is its prospect at c minus the floor’s, that is, $P(G \text{ given } c, \pi)$ minus $P(G \text{ given } c, \text{floor})$.

Hold the position and the opponent fixed. If the game would be won three times in ten from there with a novice playing on for our side, and eight times in ten with Carlsen playing on, the improvement of Carlsen’s policy on the floor at that position is 0.5 ($P(G \text{ given } c, \pi_{\text{Carlsen}})$ minus $P(G \text{ given } c, \text{floor})$), which is $0.8 - 0.3$. Both prospects depend on the position, and so does the gap between them. From a lost position both are low and the gap is small; from a level position the gap is at its widest.

Clearly bettered. A move m is *clearly bettered* at c , under π , when some alternative m' available there has a prospect, $P(G \text{ given } c, m', \pi)$, that exceeds m ’s by a clear margin.

Call a difference *clear* when it exceeds the margin the practice treats as significant.¹² Carlsen’s 0.5 is clear by any standard chess could set. The margin matters where improvements are smaller and the practice has had to decide what counts. So, say, a strategist whose policy raises a candidate’s chance of holding a swing seat from 40 to 48 per cent has done something the profession would notice and pay (a lot) for. One whose policy raises it from 40 to 41 has not.

Floor-level play loses often, of course, but that doesn’t matter vis-à-vis its status as a baseline. What improvement on the floor is concerned with isn’t how often

¹¹The floor just defined is the *naïve* floor, which I use throughout; ascriptions can set the bar higher (‘knows how to make risotto *well*’, ‘knows how to make risotto *extraordinarily*’, etc.), and §7 returns to this.

¹²The practice’s standards needn’t be entirely the last word. They are themselves answerable to the facts that make moves good. The idea here is that, whenever they are *not* actually calibrated to those facts as they perhaps purport to be (as the practice of bloodletting’s were not, as we now know), clearing the practice’s bar counts for naught (§6’s calibration condition will explain more).

the floor wins but how much *more* often the policy wins than the floor does. Take a position in which things have already gone badly for our side, so that winning from there is unlikely whoever plays on. A novice would win one game in twenty from it, and Carlsen would win one in four. Neither prospect is high, but the second exceeds the first by 0.2, and that number is what Carlsen's policy contributes from that position. Importantly (nipping in the bud over-intellectualisation worries) the agent needn't actually represent any of this at all. Risk-reduction is just a property her choices have (compare: safety is a property a belief has whether or not the believer knows it.)

Why the floor, rather than a baseline drawn from the agent herself? Two reasons. Take Carlsen. If the baseline were how Carlsen himself plays (i.e., best in the world!), then the prospect under his policy and the prospect under the baseline would be the same number, and his improvement on the baseline would always be zero. Also consider that a baseline that differs from agent to agent gives no common scale, and the account needs one, because its verdicts compare agents. That Rivera knows how to win and the pundit doesn't, or that she knows how *better* than a rival does (§7 takes up such comparatives), is a claim about whose policy improves *more*, and it makes sense only if both improvements are measured against the same thing. The floor by contrast is the same for everyone who undertakes the task.¹³

4(c). 4(b) asked two things about the proposal that a good move comes from a way of choosing that robustly lowers the risk of failure, namely (i) *relative to what*, and (ii) robust in what sense. The floor answered the first. This subsection answers the second. Two points about the measures we just gave show why a clear answer is needed to deal with certain kinds of cases. First, take Ryle's drunkard a step further. Suppose someone making moves at random stumbles into a mate in one and plays it. In terms of prospect (recall that this is the chance that the game ends in a win if play continues from the move as the mover's policy selects), obviously no move could do better (it's checkmate!). It's still not a good move though in the sense that matters here. And this is because of *how it was selected*. Here we can recall a familiar point about belief and luck from the literature on (basis-relative) safety and necessary truths (Pritchard, 2009, 34). Do the sum $2 + 2$, write '4', and believe it. Your belief is necessarily true. Is it safe? Even though there are no worlds where $2+2$ isn't four, your belief might still be unsafe. If you got there by the rule 'put the 2s together to make 22, then subtract 18', your calculating was atrocious, since that way of calculating goes wrong nearly everywhere else. So just as a necessary truth can be believed unsafely, and thus fall short of knowledge, a

¹³Why the floor rather than *typical* play? Consider that an ordinarily competent risotto-maker's moves *are* typical. So they would show no improvement over a 'typical-play' baseline, and she would, absurdly, count as making no good moves despite succeeding nearly every time.

(trivially) risk-eliminating move can be selected badly (thrown up by dumb luck from a policy that goes wrong nearly everywhere else) in a way that manifests no know-how at all. In both cases the goodness that matters attaches to the *way the result was produced*, which in our case is the *way of choosing*. Second, the definition of a move’s prospect made a choice that needs defending, namely that after the move it continues play as *her* policy selects, not as the floor would. Consider why.

SACRIFICE. A grandmaster, deep in a middlegame, gives up her queen. The sacrifice is a brilliancy, in that it begins a forcing sequence that she will conduct, through a dozen further choices, to a won endgame. But in the nearest worlds where the resulting position gets “handed off” to baseline (floor-level) continuation, the sacrifice is a disaster. This is because the floor-level novice who inherits it won’t find the forcing sequence and will simply be a queen down and thus *worse off* than under a normal (non-sacrificing) floor-level continuation.




Figure 1: The same move, two continuations. A move’s prospect is computed with the agent’s own policy continuing play, not the floor.

Had a move’s prospect been computed with floor-level play taking over after it, the sacrifice would score as a blunder. For many moves the choice of continuation doesn’t make a difference, given that what the move contributes isn’t going to depend on who plays on afterwards (a move that wins a queen helps whoever inherits the position). But expert moves are often like the sacrifice. And so is Rivera’s decision to send the candidate out tonight for an unscripted press conference (this, we can suppose, is her best move *only because she will manage what follows*, and might be a disaster from an amateur). So a move’s contribution then gets assessed *as she would carry it through*, and the goodness that matters attaches in the first instance to her whole way of choosing, with individual moves inheriting their goodness from it.

Two more definitions unpack that idea. The first says when a policy is good, where ‘good’, here and throughout, means good for reaching the task’s terminus:

Good policy. A policy is good at a context c just in case (i) its improvement on the floor at c is clear, and (ii) the improvement is robust.

Robustness in (ii) is really just the safety idea again, transposed from belief to the policy, or rather to the *policy’s edge over the floor*, which is what it is that has

to hold up across nearby situations as (by parity) a belief's truth does in the case of a safe belief. So a safe belief is one that couldn't easily have been false (in the sense of [Pritchard, 2005, 156](#)), whereas a policy's improvement on the floor is robust when that improvement *couldn't easily have been absent*. That is, we hold the agent's identity, her policy, and the task fixed, vary the situation in the ways it easily could have varied (viz., the *nearby situations* – the variations the practice trains for and examines against, including those in her own passing condition, plus the ways the task's circumstances in fact vary often enough to matter), and the edge continues to be there. The policy's "edge" over the floor is then no fluke of this particular situation (in the way the random mover's checkmate plainly was).

The second definition says when a move is good, which is what §4 set out to answer. Since a move inherits its goodness from the way of choosing, a move is good by being what a good policy selects (though with two provisos the definition builds in):

Good move. A move made at *c* is a good move just in case (i) it is what a good policy selects at *c*, and that policy is the one the agent's own choosing put into effect, not merely one whose output her move happens to match; and (ii) it is not clearly bettered at *c*, under that policy, by the alternatives available there.

Clause (i), which references **Good policy**, builds a de facto anti-luck component into what counts as a good move (the Ryle-style drunken checkmate above fails it, since no good policy of the mover's (just a drunk random policy) selected the move).¹⁴

The second part of clause (i) handles "friendly helper" cases, such as this variant of Pavese and Beddor's archer Archie ([2023, 610](#)). Suppose Archie's archery instructor always holds his bow to help him hit the bullseye, and Archie releases. Is his assisted draw and release a good move here? Well, maybe in some colloquial sense, e.g. "Great shot! Too bad he needed help", but not in the sense that has anything to do with whether *Archie* possesses and manifests knowledge how to hit the bullseye. A good policy was indeed in effect, but here it was the helper's choosing, not Archie's, that put it there.¹⁵

Consider now clause (ii) in **Good move**, viz., the non-betterment clause. Clause (i) grades a move only against the floor. Beating the floor is a low bar. A policy can clearly and robustly improve on entry-level play while routinely selecting moves

¹⁴The same clause covers Pavese's Sylvia ([2026, 262–263](#)), who sends a signal on time by trusting a clock that, entirely unknown to her, stopped exactly twenty-four hours earlier. Her policy 'trust the clock' has no robust edge over the floor, so her move succeeds by luck rather than through a good policy.

¹⁵'Own choosing' requires only that the selection be made by something in the agent rather than by someone or something outside her. It doesn't require intentional control. (A chess engine's selections are its own in this sense.)

that are clearly worse than the alternatives available at *c*. ‘Available’ here is in the sense **Move** already fixed, so the comparison class is just the other moves in the agent’s own repertoire (and at her level of choice, so no selecting of muscle fibres, etc.) that the context leaves open. Clause (ii) then says that to be good, a move must not only issue from a clearly floor-beating policy but also hold up against what else she had available in her repertoire at *c*. Take an experienced physician whose policy clearly beats entry-level practice. With the culture results on the chart identifying the infection, she prescribes the broad-spectrum antibiotic she would have used before the results came in, when the targeted drug those results call for is within her repertoire and on the ward. She has not made a good move. Her policy clearly beats the floor, so clause (i) is met, but the move is clearly bettered by an alternative available to her, so clause (ii) isn’t.

4(d). Before turning to the disposition itself, let’s zero in on three independent dimensions of a task along which a policy can fail the robust-improvement condition on a good policy. Each is a way in which a policy’s edge over the floor can be fragile. (§5 will divide the disposition into parts that answer to each of them.)

A policy’s edge over the floor can be fragile, first, because of the variance of §1, viz., the extent to which the good options shift from one context to the next. Where variance is high, the nearby situations that robustness ranges over include ones in which the good options have moved, so an edge that holds tonight can be absent in them. Take the task of winning a tennis match. A player whose policy is to serve wide to the backhand has a clear edge over the floor against an opponent who can’t return that serve. But the opponent will study the pattern. They’ll then start standing wider, and against that adjustment the very same policy adds nothing to the chance of winning the match.

The other two dimensions threaten a policy’s edge differently. Whereas variance moves the good options themselves in such a way that what the policy selects can stop being good vis-à-vis the task at hand, the other two leave the good options where they are and instead make the good option, on a given occasion, either (a) hard to *spot* or (b) hard to *execute*. Call these task dimensions *alignment* and *precision*:

Alignment. A task’s *alignment* is how far the options that naturally occur to an agent line up with the good ones.

Where alignment is low the obvious move is a trap, and much of choosing well consists in not making it. SENATE’s playbook options were all bad.

Precision. A task’s *precision* is how narrow the band of execution is within which a chosen move stays good.

Where precision is high, a move that is selected well can still fail because it is carried out a little wrong. In removing a tumour, the neurosurgeon’s decision about

where to cut can be the right one, and the cut can still cost the patient speech if the blade travels a millimetre past where it should stop. Chopping carrots for a stew, on the other hand, is a low-precision task. Errors are forgiven as long as the carrots come out roughly the right size.

Combinations of the three (high to low variance, alignment, and precision) give recognisable kinds of task. Some of these combinations already have familiar names. A low variance, high alignment, and low precision combination lines up with easy routine tasks (risotto, §1's low-variance pole, or playing a scale). Combine low variance and high alignment with medium to high precision and you get typical *crafts* (carving, the concert stage). This is where Aristotle's *technē* sits – viz., productive knowledge with a settled and teachable account of how the product comes about (Crisp, 2014, VI.4, 1140a). The high variance, low alignment, and low precision combination gives what we might call “wicked-strategic” tasks (campaigns, day trading, etc. – this is §1's high-variance pole, which is hostile on two dimensions at once). The high variance, low alignment, and high precision give wicked-complete tasks (a special-forces operation behind enemy lines, where the situation keeps changing, the obvious move is often a costly trap, and a shot or a step a little wrong is fatal). The debate's usual battery of examples come mostly from the first two kinds (routine tasks and crafts), and the cases that drive this paper come from the wicked-strategic and wicked-complete kinds.

5. Now the view itself:

Know-how/good moves (KH-GM): To know how to complete a telic task *T* is to have a *stable, trainable disposition to select and sequence good moves* across the live contexts *T* presents.

Let's unpack it, using the terminology we've now got on the table from §4. The disposition KH-GM identifies with know-how is a standing state of the agent (one that's trainable, and so improvable – more on this shortly), and like any disposition of an agent it can be possessed without being exercised. Brahms, for example, *knew how* to play a fugue even when he was busy cooking, etc.

What makes that disposition *know-how*, on KH-GM, is a matter of *what it would do in engagement*.¹⁶ Specifically, someone who knows how to complete a telic task *T* is such that, were she engaged with *T* at any of its *live* contexts (i.e., contexts at which selection still bears on the risk of failure),¹⁷ the moves she made would characteristically be good ones. That is, her own choosing would put into effect

¹⁶To be engaged with *T* is to be pursuing *T*'s end, whatever that end is. So a grandmaster who plays to let a child win isn't engaged with the task of winning the game but rather with a different telic task, namely, that of (say) losing it gracefully, and may know how to do that too.

¹⁷In 4(b)'s terms, a context is live when some policy's improvement on the floor there is clear, so that the task is neither already lost whatever is done *nor* already secured by entry-level play.

a policy good at that context (good in that it clearly and robustly improves on the floor), and what it selects would, as a rule, hold up against the alternatives available there. This is the *selecting and sequencing* the KH-GM definition speaks of. The policy selects at every context the task presents. And since a policy is good at a context only as assessed as the agent herself, not just anyone, would carry it through (SACRIFICE's lesson), each move is selected with the moves to follow in view. Sequencing good moves is then a matter of selecting each of them in that way and actually putting each into play, as a step in a course of play she will herself carry on.

Regarding *stable*, as it features in KH-GM: when someone knows how to ϕ , for some telic task, her moves would characteristically be good not merely at this or that context but at any live one. The restriction to live contexts matters given that, in a hopelessly lost context (one where failure is effectively guaranteed, and no move that could be selected would change that more than negligibly), or in one already secured (where entry-level continuation reaches the terminus with near certainty, as when all that remains is to press the clearly marked button at point blank range), *no policy clearly improves on the floor*, so there aren't any candidate good moves left to demand of anyone.^{18,19} *Trainable* names the kind of state the disposition is. It's a state that can be built and kept in tune through feedback.²⁰

What it takes to *have* such a disposition, the one word in KH-GM still unexplained, is taken up in §6. (A further feature of the view, that the disposition is manifested *in the good moves themselves* and not in success, will be covered in 5(c).)

5(a). So what's the structure of the disposition the possession of which KH-GM identifies with being in a state of (telic) know-how? The answer (to which 'selecting and sequencing' gives a hint) relates to the intellectualism/anti-intellectualism

¹⁸But aren't we inclined to praise some attempts as being still better than others even in a hopelessly lost position? I think the temptation for talking like this comes from not holding the terminus fixed. Hold it fixed (winning *this* game, *this* race) and nothing on offer alters the risk of failing at *it*, even if it would alter the risk of failing at similar tasks in less hopeless circumstances.

¹⁹Here's a worry about routine steps. In the middle of a task there are contexts where the expert and the novice would make the same move (adding stock when the pan runs dry, say). At such a context the expert's move isn't any better than the novice's. So it might look as though no policy clearly improves on the floor there. Such a conclusion would imply that the context is not live. Here's why that's wrong. A policy's improvement at a context gets measured over the *whole continuation from that context*, and so not just over the single move made there. Start from the dry pan and let the expert carry on cooking, and the risotto is more likely to come out well than if the novice carries on, even though both add stock now (a clear improvement on the floor at this context). The context thus counts as live even though the two of them do the same thing there. It would not count as live if, beyond making the same move now, the expert's continuation from there did no better than the novice's either.

²⁰On improbability through practice as what marks skill off from reflexes and habits, see Ryle (2009, 29–31), Fridland (2014, 2731–2732) and Stanley and Williamson (2017, 713, 716).

divide. It falls out of the KH-GM view that there are two components of the disposition (to *select* and *sequence* good moves), playing complementary functional roles. Call these *underwriting components* and *manifestation components*.

Underwriting components are the standing features of an agent that fix which move her choosing selects at a context (i.e., which options occur to her as candidates, when they occur, and which of them survives to be made).

Manifestation components are the standing features of an agent that fix how well a selected move is carried out (i.e., whether the move made is, within the practice’s tolerances, the move chosen).

Paradigm underwriters include things like standing propositional knowledge of the domain (e.g., for someone who knows how to get from London to Edinburgh, knowledge of when the LNER services leave King’s Cross, etc.), including knowledge-wh (Stanley, 2011c; Stanley and Williamson, 2017, 714–715), as well as such things as internalised heuristics (e.g., the strategist’s rule of thumb to hold a response until the first news cycle has run), trained perception that makes the relevant features salient (Bengson, 2016), and, where the evidence supports positing them, practical representations in Pavese’s sense, that is, stored rules that break a task into commands the agent can directly carry out (Pavese, 2015). Paradigm *manifesters* include the stored motor programs that turn a chosen move into a sequence of bodily movements (motor schemas (e.g., Schmidt, 1975)), the fineness of sensory and motor discrimination that practice sharpens over time (acuity (Stanley and Krakauer, 2013)), as well as the kinds of fast, unreflective adjustments by which the motor system corrects a movement in flight when it starts to go wrong (Fridland, 2017; Shadmehr, Smith and Krakauer, 2010), as when an outfielder runs so as to keep the ball’s image rising at a steady rate rather than computing where it will land (Clark, 2016, 190–191; McBeath, Shaffer and Kaiser, 1995).

Recall the three dimensions of 4(d) along which a task can threaten a policy’s edge (variance, alignment, precision). It’s going to be one or other of the two components of know-how (underwriting components and manifestation components) that keeps the edge in place against each threat when one manifests the know-how one has. Against high variance and low alignment, it’s going to be the underwriting components, since both make it harder to select the good move (high variance because the good options keep moving, low alignment because the obvious option is often not one of them), and against high precision it’s the manifestation components, given that when there’s a narrow band of acceptable execution, this makes it harder to carry out the move selected, even if the move itself is easy to select.

The distinction, to be clear, divides (functional) roles, not brain or any other physical regions. And for certain tasks, it might be that a single capacity serves both (see, e.g., Shepherd on the metacognitive monitoring (2026, 13–14)). What

matters is that both roles have to be filled, given that the disposition to select and sequence good moves is absent if the agent cannot select the good move or cannot carry out the move she selects. That's why the distinction explains two familiar ways of falling short. Take first a case adapted from Snowdon (2004):²¹

ARMLESS PIANIST: Pia, a concert pianist, is in an accident and loses her arms. While she investigates prosthetics, she can't play what used to be her signature concerto.

Pia's musical understanding is untouched by the accident. What the accident took from her (and prosthetics have not yet returned) is motor capacity. So her underwriting is superb (the candidate moves that occur to her would all be good if played – indeed, she can even 'play' it in her head). But there is no manifestation for Pia anymore. And so no integrated disposition of the kind KH-GM identifies with know-how.

Can she still teach the piano? Of course! The underwriters she keeps ground the ascriptions that are felicitous even after the injury ('she can teach it wonderfully'; 'she knows what to do at the midway point', etc). If she's asked to perform tonight, though, a know-how ascription comes out (rightly) false.

Now consider the lucky safe cracker:

LUCKY SAFE CRACKER: Lucy has picked thousands of safes and has superbly coordinated hands. But she has no idea what to do with this particular safe. All she can do is guess, jiggling a pin this way and that. Amazingly, it opens.

Pia lacks the disposition to select and sequence good moves given that she's lacking the needed manifestation components. Lucy falls short when it comes to the underwriting ones. KH-GM rightly predicts that *neither* Pia nor Lucy has the relevant know-how (Pia to play that concerto, at that context; Lucy to crack that safe, at hers).

5(b). The two roles (underwriting and manifestation) also explain two grades of credit our practices already draw. I'll mark these with a pair of labels adapted from Watson (1996, 229, 233–237): *attributability-grade* and *accountability-grade* credit. In their home debate about moral responsibility, conduct is, in Watson's terms, *attributable* to an agent when it expresses who she is. That is, it expresses her own commitments at work in what she does (unlike, say, a sneeze or a stumble). An agent is for Watson *accountable* for conduct that she can properly be asked to justify and answer for.

²¹This is a variation on some cases from Snowdon (2004, 8, 10). The safe-cracking case below is likewise built from Snowdon's examples of capacity without know-how (2004, 11–12).

Attributability-grade and accountability-grade credit for know-how are these same two kinds of appraisal, applied by analogy to performances in a task rather than (as per in Watson's original setting) to moral conduct. On the attributability side, what a performance expresses is the agent's trained skill at the task rather than her character. We often say things like 'Her hands knew what to do' when talking about a surgeon or an improvising pianist. That's *attributability-grade credit* we're giving here, viz., credit for a *manifestation well brought off*, viz., for the move being *carried out* as it was chosen to be. What such credit primarily tracks are the subject's manifestation components more so than her underwriters.

Our practices are a bit different when the know-how manifested relies more on underwriting components. Think of attributions like 'She made the right tactical call' (rather than 'her hands knew what to do'). That is *accountability-grade credit*, credit for the *call* itself. This kind of credit is available only where the practice traffics in reasons, where 'why did you do that?' is a question in order, and it brings a standing the first grade doesn't, since the agent can be asked to defend the call and thus take responsibility for it. She can be asked because an underwriting-heavy call is made on considerations she registered and weighed, so a demand for reasons has something to attach to, as it doesn't with a surgeon's hands.

A good move earns credit distinctive of each grade (attributability vs accountability) in proportion to how much the corresponding component contributed to making it good, and those proportions are sensitive to where the task sits on 4(d)'s three dimensions.²² In a high-precision task the hard part isn't picking the move at a given live context so much as carrying it out. That means execution is going to contribute more of what makes the move good, and by extension the larger share of its credit is attributability-grade. In a high-variance or low-alignment task, on the other hand, the hard part is *finding* the move in the first place, so selection contributes more, and the larger share of its credit is thus going to be accountability-grade.

Practitioners already talk this way when they distinguish players with great shots from players with game sense. Tiger Woods's safe nine iron to the middle of the twelfth green in the final round of the 2019 Masters, while his rivals attacked the pin and found the water, was nearly *all call*, and its credit is accountability-grade. John Daly's longest drives are nearly all execution ('grip it and rip it'), and their credit is principally attributability-grade.

5(c). The view also non-trivially separates know-how from success, which is an advantage over a natural rival view on which knowledge how to ϕ is *manifested* only in successful instances of ϕ -ing. On KH-GM, even when the disposition is

²²This is what separates KH-GM from a hybrid view that says know-how involves both knowledge and ability and leaves it there. The proportions are fixed by the task, so the question which of the two know-how is has no general answer, and a middling answer would be a general answer too.

being exercised, its manifestation and the task's success come apart, given that the disposition is manifested in the selected good moves themselves and success (attaining the terminus) is a further matter. Consider:

EVACUATED RISOTTO. Lina is making risotto impeccably. She's toasting the rice, ladling stock in the right increments, stirring at the proper cadence, tasting and adjusting and so forth. A gas leak forces an evacuation. No risotto is plated.

Lina manifested her know-how in every move she made, and the task failed anyway. A notable rival view about know-how manifestation – the success-constraint view – must deny this, which is a bad mark for extensional adequacy for such a proposal. According to Cunningham's know-how success constraint, knowing how to F 'can manifest itself only if its manifestation is a successful instance of F'. Know-how is 'essentially successful', as knowing-that is factive, so an agent whose F-ing fails at most attempted to display her know-how (Cunningham, 2022, 4.3).

Cunningham motivates this with a pianist who plays the fugue with many mistakes. His point is that it sounds odd to say that she manifested her know-how in the failed performance. But KH-GM not only gets EVACUATED RISOTTO right; it also predicts the oddness of the case Cunningham uses to motivate the constraint. The pianist's mistakes are bad moves, issuing from a policy that doesn't clearly beat the floor. Keep every move good and add a gas leak instead, as with Lina, and the oddness in attributing know-how is gone. So the success constraint is right that know-how can't be manifested in a bad performance, but wrong about what has to be good for it to be manifested.

6. One word in KH-GM is still a promissory note, *have*. §4 said what good policies and good moves are, and nothing yet says what it is to *have* the disposition that would select them. On the ordinary view of dispositions the answer looks trivial. To have one is just to be in a state from which its manifestations would issue, and how the state came about is neither here nor there, given that intrinsic duplicates are disposed alike (Lewis, 1997, 147–148). This section argues that, at least for *this* disposition, that's not going to be enough, and that its history matters to whether it is had.

With that in mind, consider the following case, due originally to Stanley and Williamson (2001, 435).

FLIGHT SIMULATOR. Bob wants to learn to fly using a flight simulator. His apparent instructor, Henry, is a malicious imposter who has inserted a randomising device in the simulator's controls and intends to give incorrect advice. By sheer luck, the randomiser outputs exactly what the real controls would have done, and Henry's advice matches

what a real instructor would have said. Bob passes the course with flying colours. He has never flown a real plane.

Take the telic task the simulator course trains for and the flight test examines, landing the plane. Nothing Bob did in the simulator manifested know-how to land a plane, and KH-GM can explain why. Bob's way of choosing moves at each live context in the simulator was something like "trust the instruments and the instructor." Across nearby situations, ones where the randomiser outputs something else or the advice goes wrong, that policy has no edge at all over the floor. This is because Bob cannot tell good inputs from corrupted ones. Instrument failure and bad instruction are among the things flight training drills for, so those situations are in the neighbourhood, and thus Bob's policy fails robustness. So no good policy of Bob's selected his moves, and by clause (i) of the definition of a good move none of them was good. (A trainee whose instruments compute real aerodynamics is in a different position, since her instruments could not easily have said otherwise, so no such situations are nearby for her.)

Here is a trickier version of Stanley and Williamson's case.

SIMULATOR GRADUATE. By the case's own stipulation, the randomiser output exactly what the real controls would have done. So the standing control-responses Bob takes away from the course are the true ones. Put him next to Carla, a graduate of the certified course next door, and suppose their standing dispositions are, at this moment, exactly alike: same responses, same profiles, across every approach and landing either could fly.

If having the disposition were fixed entirely by how you are now constituted, Bob and Carla could not differ. But they clearly do differ! (An examiner and an insurer would agree). Carla knows how to land a plane. Bob does not, not yet at least, as he's only just begun training. So having the disposition can't be fixed on present constitution alone. Something about where it came from matters. The task is then to say which feature of its history that is, without overfitting the condition to this one case.

Consider how each of them came to have the policy she has. The facts that make a move good, for the task of landing, are facts about how planes behave, since those settle whether a given move at a given point in the approach (easing the throttle back, say, or lowering the flaps, and other similar kinds of things you get from pilot training) improves the chance of a safe landing. Carla's policy was built by a simulator that computes real aerodynamics and by instructors whose corrections were driven by real flying. What those sources told her depended on how planes behave, so had her policy been wrong, they would have corrected it. Bob's policy was built by a randomiser, and what it told him depended on nothing about flying at all – it depended on dumb luck. Had the randomiser been wired differently, it

would have just chucked out rubbish, Bob would have ingested the rubbish, and nothing in his training would have caught the error.

Call a source of training or feedback *calibrated* to the facts that make moves good in the domain, viz., the facts that give a policy whatever edge it has over the floor, if its outputs depend on those facts. Carla's sources were calibrated *ex hypothesi*. Bob's just happened to match. That is the feature of history the account needs, and it can be stated as a condition on having the disposition:

Possession. An agent *has* the know-how disposition (per KH-GM) only if her implemented policy is *keyed* to the facts that make moves good. An implemented policy is *keyed* to the facts that make moves good just in case it was built or is maintained through *calibrated sources*, that is, just in case (i) it was built up through calibrated sources, or has since been held under them long enough that an error in it would by now have been caught, and (ii) nothing that is not a calibrated source now steers it, and whatever maintenance it receives comes from calibrated sources.

The first disjunct of clause (i) is what separates Carla from graduate-Bob. Their constitutions are identical *ex hypothesi*, but her policy was built up through calibrated sources and his was not. The second disjunct in clause (i) is why Bob's verdict is 'not yet' rather than 'never'. If he takes lessons long enough, he becomes like Carla, and his policy becomes keyed. (This is what supervised hours add for a graduate whose policy already, by luck, selects the right moves.)

Clause (ii) excludes, rightly, a know-how variation on Pritchard's (2012) case of "Temp", for whom what's at issue is telling the temperature.

TEMP-KH: Unbeknownst to Temp, a benevolent angel matches the room's temperature to a randomly fluctuating thermometer each time Temp reads it. So the policy Temp implements for telling the temperature, consult the thermometer and trust what it says, always works (cf. Pritchard, 2012, 260).

Temp no more knows how to take the temperature than Bob knows how to land a plane, but he fails **Possession** at a different clause.²³ Temp may well have learned how to consult thermometers in the normal way, through reliable feedback. But his policy for telling the temperature of this room clears the floor only because the angel keeps making it clear the floor. The angel knows the temperature, but is not a calibrated source for Temp's policy. That is, the angel gives Temp no training

²³Clause (ii) also helps to rule in the retired grandmaster. A player who no longer trains or plays (Kasparov, say, years after retirement) receives no maintenance and nothing steers his policy, so clause (ii) is met trivially, and his policy remains keyed on the strength of clause (i) alone.

or feedback, and the policy would be no different had the angel been absent. The angel adjusts the world so that the policy comes out right whatever it does. So what now sustains the policy's edge isn't the policy's connection to the facts but something outside it, and thus clause (ii) fails.

Now that we've laid out the possession condition on know-how, we've completed the good moves account laid out in KH-GM. §5 said what the disposition to select and sequence good moves would do in engagement, and this section has said what it takes to have it. Together they give the account in full. An agent knows how to complete a telic task *T* just in case she is disposed to select and sequence good moves across the live contexts *T* presents, and the policy behind that disposition is keyed.

7. The intellectualist's strongest card has always been the semantics. On the standard compositional treatment, '*S* knows how to ϕ ' embeds a question, and is true just in case *S* knows an answer to the question how to ϕ (Williamson and Stanley, 2001, 420–425).

That semantic treatment, as Stanley (2011a, see esp. 32–39, 70–75, 111–113) has argued at length, is well supported by the linguistic evidence. §3's arguments against intellectualist proposals don't challenge the semantic case specifically, that is, the intellectualist's mainline argument from premises about how know-how ascriptions pattern with the other knowledge-wh ascriptions to the conclusion that the truth conditions of '*S* knows how to ϕ ' are given by *S*'s knowing an answer to the embedded question. So it would seem to be at least a defeasible strike against KH-GM if it had no way of respecting this embedded-question semantics.²⁴

KH-GM can respect it, though it does so by keeping the schema at the core of the semantics while filling that schema in differently from the way intellectualists have. The schema is this:

Embedded question semantics for KH (schema): '*S* knows how to ϕ ' is true just in case *S* knows an answer to the embedded question 'How (for *S*) to ϕ ?'.²⁵

The schema leaves two things open, substantively speaking. It doesn't say *what kind of thing* an answer to the question 'How (for *S*) to ϕ ? is. And it doesn't say anything special about what knowing such an answer requires. As Bengson and Moffett (2011) capture this point in other words, a know-how ascription relates a subject to *something*. So there are two salient things to ask about it, which are (i) what the *relation* is and (ii) what the *thing* is that is its object (2011, 163). Any account

²⁴For responses on behalf of ability theorists, see Noë (2005, 286–287), Devitt (2011) and Hornsby (2012).

²⁵The '(for *S*)' marks the unpronounced subject of the infinitive (what linguists call 'PRO'), so that the embedded question is 'how PRO to ϕ '. In the reading that matters for know-how, PRO is controlled by the subject of 'knows', so the question is how *for S* to ϕ rather than how for one to ϕ in general (Stanley, 2011a, 70–75; Williamson and Stanley, 2001, 422–424).

of the state of know-how itself, as opposed to the semantics of its ascription, should really say what each of them is if it accepts the schema.

S&W answer the two questions by carrying over the classical semantics of theoretical questions. On Karttunen's (1977) version, which they adopt, an embedded question denotes the set of its true answers (Hamblin, 1973). Those answers are propositions, and 'know' relates the subject to one of them as it relates her to any proposition (Williamson and Stanley, 2001, 420–421, 425). Groenendijk and Stokhof's (1982) treatment differs in the details but agrees on the point that matters here, which is just that the answers are propositions. Applied to 'how to ϕ ', the answers are propositions of the form ' w is a way for S to ϕ ', and what makes knowing one *practical* (as we want in the case of non-deontic, rather than mere deontic, know-how) is the mode of presentation under which it is known.

§3 already gave reasons to doubt both halves of that answer. What I want to show is that KH-GM is able to keep the above core embedded question semantics for know-how at the schematic level and then answer the two questions differently.

Let's start with what an answer is. The question embedded in 'Rivera knows how to win this race' is a practical one. It asks what (for Rivera, to win) to *do*, and a complete answer would say what to do *at each context the race can present*, at all of its live contexts. KH-GM has a name for something that does that, viz., a *policy*. A policy is a function from contexts to moves, so a policy is a *complete answer to a practical question*.

Now, a plan like Plan-R is (often) only going to be a partial answer to any corresponding practical question. That's because (in many high variance tasks at least) it says what to do along one course of events but then *de facto* goes silent once events leave that course. And a rule of thumb like Plan R* is an answer so incomplete that (as we saw) the pundit can give it too (and thus, it's not an answer to the embedded question the possession of which distinguishes those who truly know-how to do the relevant thing from those who don't).

So, then, the thing at the other end of the relation, on KH-GM, is a policy. Compare the two things that orthodox views put there, as Bengson and Moffett set them out (2011, 163). The propositionalist puts a proposition there, and the relation to it is knowing. The dispositionalist puts the action-type ϕ -ing there, and the relation to it is being able. The second option is what leaves ability views at odds with the schema, since ϕ -ing isn't an answer to anything. KH-GM avoids that trouble, because a policy, unlike an action-type, is an answer to the very question the ascription embeds.

Answers can be correct or incorrect. For a theoretical question, the correct answers are the (direct) true ones, which is why the classical semantics has the question denote its true answers. For a practical question of the form 'How (for S) to ϕ ?' the correct answers are the *good ones* (in §4's sense.) A good policy – viz., one whose improvement on the floor is clear and robust – is a correct answer to

the practical question ‘how to win?’. A *bad policy*, on the other hand, isn’t a correct answer to that practical question, any more than ‘in the cellar’ correctly answers ‘where are the keys?’ when they are on the hook. So ‘How (for *S*) to ϕ ?’ can be given the classical treatment as it stands, such that it denotes the set of its correct answers, with goodness in the place that truth occupies for a theoretical question.

Now to the remaining question of what knowing an answer requires. Let’s take the theoretical case first. To know where the keys are, it isn’t enough that a true answer be simply available to one. One has to hold it, by believing it (occasionally or dispositionally), and one’s holding a true answer must be no accident, which is the anti-luck condition 3(a) invoked in the form of safety (or aptness, or what have you). So knowing an answer has two ingredients. There’s an attitude that makes the answer the subject’s own, and then there’s a condition that ties the answer’s correctness to the subject non-accidentally. KH-GM supplies both for the practical case (and much of the work of supplying them has already been done in §6.)

The attitude is *having the policy*, in the sense the good-move definition fixed. The policy is one the subject’s own choosing puts into effect as opposed to one whose outputs her moves happen to match. The anti-luck condition pertinent to knowing the answer to a practical question by having a policy, is simply *that the policy be keyed*, i.e., that it was built or is maintained through sources calibrated to the facts that make moves good, *so that its goodness is no accident*, (as Bob’s and Temp’s clearly were). Having a *keyed good policy* thus stands to *practical questions* as *safely believing a true proposition* stands to *theoretical questions*.

The verb ‘know’ accordingly expresses the same relation in ‘knows where the keys are’ and in ‘knows how to win’. In both, the subject has a correct answer to the embedded question through a connection to what makes it correct. What differs is specifically the *kind of answer* the question calls for, and with it, what *holding the answer* and *being connected to its correctness* come to.²⁶

An objection presents itself at this point. If holding an answer and being connected to its correctness come to one thing for factual questions and another for practical ones, hasn’t ‘know’ just been rendered ambiguous, with one sense for ‘knows where’ and another for ‘knows how’? It hasn’t. In ‘knows where the keys are’ and in ‘knows how to win’, the verb contributes the same relation, to wit, that of *having a correct answer to the embedded question through a connection to what makes it correct*. What differs between the two ascriptions is the question each embeds. One

²⁶Habgood-Coote (2019, 89–92) also keeps the embedded question schema, on which knowing how to ϕ is knowing an answer to ‘How (for *S*) to ϕ ?’, while parting with the way Stanley and Williamson fill it in. But he parts with them at a different place than I do. Habgood-Coote leaves the answers as propositions, as Stanley and Williamson do, and changes what constitutes knowing one (for him it’s an ability). On the KH-GM semantics for ‘knows how’ it is the *other* place in the schema that’s revised. The *answers* to the relevant ‘How (for *S*) to ϕ ?’ questions are no longer propositions but *policies* (functions from the contexts a task presents to the moves to be made there).

calls for a proposition and the other for a policy. It's *that* difference in the answers (as opposed to any difference in the verb) that makes holding the answer and being connected to its correctness come to different things in the two cases.

Nor, it's worth noting, is the intellectualist in any position to object to variation of that kind. After all, the intellectualist's appeal to modes of presentation already makes what having an answer takes differ from case to case. As the intellectualist tells us, you can grasp a way demonstratively just by watching someone act on it, as Hannah does in Stanley and Williamson's example when Susan points to John on his bicycle (2001, 428–429), and you can grasp it descriptively just by reading about it in a book. But to grasp it under the practical mode you need, on the intellectualist's own account, to have certain dispositions (Pavese, 2017, 375; Williamson and Stanley, 2001, 429). None of that is supposed to make 'know' somehow objectionably ambiguous on the intellectualist's view. And it doesn't on the good moves account either. The difference between us is *where the practical character of know-how is located*. The intellectualist puts it in a mode of grasping a proposition. I don't (for, among other reasons, the problems noted in 3(a)). I put it in the question, which calls for a policy rather than a proposition, and in what having a policy is.

With the account and its semantics now both in place, let's return to the promissory note where we began. In particular, let's go back to the ascription to Rivera of knowing how to win. 'Rivera knows how to win this race' embeds the question how to win this race, and by the schema it's true just in case Rivera knows an answer to it. On the good moves account, an answer to that question is a policy, a way of selecting and sequencing moves across the contexts the race presents. A correct answer is a *good policy*, one whose moves are good moves, since its improvement on the floor is clear and robust. And *knowing a correct answer* is having a good policy that is keyed, in §6's sense, to the facts that make moves good. Rivera has such a policy. She thereby knows an answer to the question how to win this race, and so 'Rivera knows how to win this race' is true. The pundit by contrast has no such policy, despite whatever he can recite about rules of thumb for campaign strategy.

A *prima facie* worry briefly floated in §1's opening note was that 'Rivera knows how to win' can't be true, since winning isn't entirely up to her. On the account given here, though, the ascription doesn't require that be the case.²⁷ And indeed that's just as it should be. In any contested task the outcome depends in part on what the other side does, so if knowing how to win required a policy that *guaranteed winning* no matter what the other side does, nobody would ever have known how to win anything contested. So when campaign donors sort prospective managers by whether they know how to win close races (those are the ones they're seeking

²⁷What it requires is that Rivera have a correct answer to the question how to win, which means having a good policy that, by §4's definition, clearly and robustly improves the chances of winning over the floor. It doesn't need to guarantee a win.

out after all), they are saying something literally true of the ones they hire.

The rest of the data go the same way. Take ‘She knows how to make risotto, but not well’. The relevant floor that a policy has to beat is entry-level engagement, but an ascription can with qualifiers raise the baseline. And ‘well’ is a paradigmatic such baseline raiser. It moves the baseline up from the entry-level cook’s policy to the highly competent one’s. So the first conjunct is true because her policy clearly beats the entry-level floor, and the second is true because it *doesn’t* clearly beat the competent cook’s. ‘She knows how to run a campaign better than he does’ compares how much the two policies improve on the floor across the task’s live contexts, which is why as we saw 4(b) needed a baseline that’s common to everyone.

8. The account given here is of telic know-how only. A natural next step would be to extend it to atelic activities, and thus to *atelic know-how*. It’s not hard, with the tools we’ve already got in place, to envision the rough shape such an account might take (though developing it in substantive detail will need to wait for another occasion.)

Here’s the rough, high-altitude idea. An atelic activity like maintaining a pleasant home or being a good friend has no terminus. But it *does* have a success condition, namely, a *state of affairs that counts as the activity going well* (Setiya, 2014, 12–13; cf. Comrie, 1976, 44–45). Whereas a telic task’s terminus is reached once (and for all), after which the task is over, an atelic success condition either obtains or fails to *at each moment the activity goes on*. That is, at any given time while you’re engaged in the activity, you are either, e.g., maintaining a pleasant home at that time or you’re failing to.

To extend the good moves framework to atelic know-how, it will need to be the case that a move, at a live context in the (atelic) activity, can make its success condition more or less likely to obtain, as, by analogy, a move in a telic task can raise or lower the prospect of reaching the terminus. And indeed that much seems straightforward. If you (say) betray a friend or refuse to apologise after a mistake, this is going to make it less likely that the friendship stays in good order, whereas lending a sympathetic ear or a repaying a debt, etc., more likely. Now, with that granted, §4’s definitions can be restated (at least schematically, as we sketch this out) for the atelic case with the atelic activity’s success condition swapped in the terminus’s place. The atelic floor will still be entry-level engagement (here we can think of the conduct of someone who grasps what the activity is for but brings no more to it than that). An atelic policy is still a function from contexts to moves. And the policy’s improvement on the floor is still just the difference between two chances (to wit, the chance that the success condition obtains if the activity is carried on as the policy selects, and the chance that it obtains under entry-level continuation.) None of these schematic definitions, note, smuggled in any talk of a terminus as such. What they mentioned was a chance of success. It’s a further fact about telic

tasks that success there just is reaching a terminus once and for all. For an atelic activity, though, the chance *of its going well* (as opposed to the reaching of a terminus) is what an atelic prospect would measure and thus what an improvement is going to be an improvement in. What has to change though is the range over which the disposition is assessed, since there's no completion anymore to actually bound it. That does complicate things. The matter of how to go about setting that range will have to wait.

References

- Bengson, John 2016. 'Practical perception and intelligent action'. *Philosophical Issues*, 26, 1: 25–58.
- Bengson, John and Moffett, Marc A. 2011. 'Nonpropositional Intellectualism'. In Bengson, John and Moffett, Marc A. (eds.) *Knowing How*. Oxford University Press, pp. 161–195.
- Bengson, John, Moffett, Marc A. and Wright, Jennifer C. 2009. 'The folk on knowing how'. *Philosophical Studies*, 142, 3: 387–401.
- Carter, J. Adam and Poston, Ted 2018. *A critical introduction to knowledge-how*. Bloomsbury Publishing.
- Carter, J. Adam and Pritchard, Duncan 2015. 'Knowledge-How and Cognitive Achievement'. *Philosophy and Phenomenological Research*, 91, 1: 181–199.
- Christensen, Wayne, Sutton, John and McIlwain, Doris J.F. 2016. 'Cognition in Skilled Action: Meshed Control and the Varieties of Skill Experience'. *Mind & Language*, 31, 1: 37–66. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/mila.12094>.
- Clark, Andy 2016. 'Surfing Uncertainty'.
- Comrie, Bernard 1976. *Aspect: An introduction to the study of verbal aspect and related problems*. Cambridge university press.
- Crisp, Roger 2014. *Aristotle: Nicomachean Ethics*. Cambridge University Press.
- Cunningham, J. J. 2022. 'Moral Worth and Knowing How to Respond to Reasons'. *Philosophy and Phenomenological Research*, 105, 2: 385–405.
- Cutrer, William B., Miller, Bonnie, Pusic, Martin V., Mejicano, George, Mangrulkar, Rajesh S., Gruppen, Larry D., Hawkins, Richard E., Skochelak, Susan E. and Moore, Donald E. 2017. 'Fostering the Development of Master Adaptive Learners: A Conceptual Model to Guide Skill Acquisition in Medical Education'. *Academic Medicine*, 92, 1: 70–75.
- Devitt, Michael 2011. 'Methodology and the nature of knowing how'. *The Journal of Philosophy*, 108, 4: 205–218.
- Dreyfus, Hubert L. 2005. 'Overcoming the myth of the mental: How philosophers can profit from the phenomenology of everyday expertise'. In *Proceedings and addresses of the American Philosophical Association*, vol. 79. JSTOR, pp. 47–65.
- Fridland, Ellen 2014. 'They've lost control: reflections on skill'. *Synthese*, 191, 12:

- 2729–2750.
- 2017. ‘[Skill and motor control: intelligence all the way down](#)’. *Philosophical Studies*, 174, 6: 1539–1560.
- Glick, Ephraim** 2012. ‘Abilities and know-how attributions’. In Brown, Jessica and Gerken, Mikkel (eds.) *Knowledge ascriptions*. Oxford University Press, pp. 120–139.
- 2015. ‘Practical modes of presentation’. *Noûs*, 49, 3: 538–559.
- Groenendijk, Jeroen and Stokhof, Martin** 1982. ‘[Semantic analysis of wh-complements](#)’. *Linguistics and Philosophy*, 5, 2: 175–233.
- Habgood-Coote, Joshua** 2019. ‘[Knowledge-How, Abilities, and Questions](#)’. *Australasian Journal of Philosophy*, 97, 1: 86–104.
- Hamblin, Charles L.** 1973. ‘Questions in Montague English’. *Foundations of Language*, 10, 1: 41–53.
- Hawley, Katherine** 2003. ‘[Success and Knowledge-How](#)’. *American Philosophical Quarterly*, 40, 1: 19–31.
- Hogarth, Robin M.** 2001. *Educating Intuition*. Chicago: University of Chicago Press.
- Hornsby, Jennifer** 2012. ‘3 Ryle’s Knowing-How, and Knowing How to Act’. *Philosophical Inquiry*, 36, 1: 80.
- 2017. ‘Knowledge how in philosophy of action’. *Royal Institute of Philosophy Supplement*, 80: 87–104.
- Karttunen, Lauri** 1977. ‘Syntax and semantics of questions’. *Linguistics and philosophy*, 1, 1: 3–44.
- Kelp, Christoph** 2021. ‘[Theory of inquiry](#)’. *Philosophy and Phenomenological Research*, 103, 2: 359–384.
- Lewis, David** 1997. ‘Finkish dispositions’. *The Philosophical Quarterly*, 47, 187: 143–158.
- Löwenstein, David** 2017. [Know-How as Competence. A Rylean Responsibilist Account](#). Frankfurt am Main: Vittorio Klostermann.
- Markie, Peter J.** 2015. ‘The special ability view of knowledge-how’. *Philosophical Studies*, 172, 12: 3191–3209.
- McBeath, Michael K., Shaffer, Dennis M. and Kaiser, Mary K.** 1995. ‘How baseball outfielders determine where to run to catch fly balls’. *Science*, 268, 5210: 569–573.
- Mylopoulos, Maria and Regehr, Glenn** 2007. ‘[Cognitive metaphors of expertise and knowledge: prospects and limitations for medical education](#)’. *Medical Education*, 41, 12: 1159–1165.
- Noë, Alva** 2005. ‘Against intellectualism’. *Analysis*, 65, 4: 278–290.
- Pacherie, Elisabeth** 2008. ‘[The phenomenology of action: A conceptual framework](#)’. *Cognition*, 107, 1: 179–217.
- Pavese, Carlotta** 2015. ‘Practical Senses’. *Philosophers’ Imprint*, 15.
- 2017. ‘Know-how and gradability’. *The Philosophical Review*, 126, 3: 345–383.
- 2019. ‘[The psychological reality of practical representation](#)’. *Philosophical*

- Psychology*, 32, 5: 784–821.
- 2021. 'Know-how, action, and luck'. *Synthese*, 198, 7: 1595–1617.
- 2026. *The practical mind: Skill, knowledge, and intelligence*. Cambridge: Cambridge University Press.
- Pavese, Carlotta and Beddor, Bob** 2023. 'Skills as Knowledge'. *Australasian Journal of Philosophy*, 101, 3: 609–624.
- Pritchard, Duncan** 2005. *Epistemic luck*. Clarendon Press.
- 2009. 'Safety-Based Epistemology: Whither Now?'.
 ———2012. 'Anti-luck virtue epistemology'. *The Journal of Philosophy*, 109, 3: 247–279.
- Rosefeldt, Tobias** 2004. 'Is Knowing-how Simply a Case of Knowing-that?'. *Philosophical Investigations*, 27, 4: 370–379.
- Ryle, Gilbert** 2009. *The Concept of Mind: 60th Anniversary Edition*. New York: Routledge.
- Schmidt, Richard A.** 1975. 'A schema theory of discrete motor skill learning'. *Psychological Review*, 82, 4: 225–260.
- Setiya, Kieran** 2014. 'The Midlife Crisis'. *Philosopher's Imprint*, 14, 31.
- Shadmehr, Reza, Smith, Maurice A. and Krakauer, John W.** 2010. 'Error Correction, Sensory Prediction, and Adaptation in Motor Control'. *Annual Review of Neuroscience*, 33, Volume 33, 2010: 89–108.
- Shepherd, Joshua** 2026. 'Salvaging the 'Sense of Agency': Metacognitive Feelings for Flexible Behavioral Control'. *Journal of Philosophy*.
- Snowdon, Paul** 2004. 'Knowing How and Knowing That: A Distinction Reconsidered'. *Proceedings of the Aristotelian Society*, 104, 1: 1–29.
- Stanley, Jason** 2011a. *Know how*. OUP Oxford.
- 2011b. 'Knowing (how)'. *Noûs*, 45, 2: 207–238.
- 2011c. 'Knowledge-wh'. In Stanley, Jason (ed.) *Know How*. Oxford University Press, p. 0.
- Stanley, Jason and Krakauer, John** 2013. 'Motor skill depends on knowledge of facts'. *Frontiers in Human Neuroscience*, 7.
- Stanley, Jason and Williamson, Timothy** 2017. 'Skill'.
- Waights Hickman, N.** 2019. 'Knowing in the 'Executive Way': Knowing How, Rules, Methods, Principles and Criteria'. *Philosophy and Phenomenological Research*, 99, 2: 311–335. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/phpr.12488>.
- Watson, Gary** 1996. 'Two faces of responsibility'. *Philosophical topics*, 24, 2: 227–248.
- Wiggins, David** 1976. 'II—Deliberation and Practical Reason'. *Proceedings of the Aristotelian Society*, 76, 1: 29–52.
- Williamson, Timothy and Stanley, Jason** 2001. 'Knowing how'. *The Journal of Philosophy*, 98, 8.