Practical Knowledge and the Structure of Action

Will Small
University of Chicago

January 28, 2011

[T]he will is by nature active; cognitive phenomena, since they must bend to the condition of the world, may be described as passive in relation the world they represent.

COLIN MCGINN
The Character of Mind, p. 117

[F]rom what glimpses I have caught of it I gather that this so-called practical knowledge is as queer a bird as any that ever flew in Plato's famous aviary. Evasive as it is, I must try to lay hands on it though I doubt if it is a genuine bird at all.

O.R. JONES
'Things Known Without Observation', p. 137

Can it be that there is something that modern philosophy has blankly misunderstood: namely what ancient and medieval philosophers meant by *practical knowledge*? Certainly in modern philosophy we have an incorrigibly contemplative conception of knowledge. Knowledge must be something that is judged as such by being in accordance with the facts. The facts, reality, are prior, and dictate what is to be said, if it is knowledge.

G.E.M. ANSCOMBE
Intention, §32

Introduction

If practical knowledge has any significance for epistemology, it must be because it is a distinctive *form of knowledge*, or is acquired via a distinctive *way of knowing*, and not merely because it is knowledge of a distinctive *topic* (practice, action) in virtue of which it has significance for those branches of philosophy with a prior interest in such subject matter (philosophy of action, ethics). Its prospects, by these lights, appear limited; practical knowledge surely does have such a distinctive and restricted topic, especially when compared with perceptual, inferential, and testimonial knowledge. Yet renewed interest in G. E. M. Anscombe's *Intention* has provided some encouragement: Anscombe argues that if an agent is doing something intentionally, he knows "without observation" that he is doing it. This *agent's knowledge* (as I'll call it) is, Anscombe claims, practical knowledge, the formal character of which she characterizes positively with a slogan borrowed from Aquinas: practical knowledge is "the cause of what it understands" (1963: §48).

But the idea that causing, or acting, or intending, could be a way of knowing somehow like perceiving, inferring, or being told runs counter to deep philosophical intuitions about the nature of knowledge. Knowl-

edge and belief, in Colin McGinn's phrase, "must bend to the condition of the world" (1996: 117); attaining a suitably non-accidental grip on the way things are anyway is what characterizes the kind of cognitive achievement that is knowledge. By contrast, causing, acting, and intending involve changing the way the world is, or representing it as to be so changed; so-called "conative" representations like intention and desire exhibit the opposite "direction of fit" from their "cognitive" brethren. Such intuitions and commitments can suggest that the expression "practical knowledge" could not but fail to denote a form of knowledge: if what it refers to is really practical, then it cannot really be knowledge, and if what it refers to is really knowledge, then it cannot really be practical.

Anscombe offers a diagnosis: it is but a dogma of modern philosophy to conceive of knowledge as everywhere speculative, contemplative, beholden to the way things are anyway; this is to project features characteristic of a species of knowledge onto the genus. "Speculative" knowledge *is* "derived from the objects known" (1963: §48), but practical knowledge is not: it is the cause of what it understands, and supposedly no less knowledge for that. Yet the remedy is far from clear: J. David Velleman—who is more sympathetic than most to this strand in Anscombe's thought—expresses a common reaction when he writes that her account leaves agent's knowledge "looking not just causally perverse but epistemically mysterious" (1989: 103). A defence of the idea that "practical knowledge" denotes a distinctive form of knowledge therefore must tread a thin line, between advancing a conception of something genuinely cognitive, but only dubiously practical (or practical merely in topic), and advancing a conception of something genuinely practical that is only dubiously cognitive, knowledge only so-called.

In this essay, I try to tread this line. Following Anscombe, I defend the idea that there is a *cognition condition* on intention and intentional action: an agent who is doing *A* intentionally has practical knowledge that he is doing *A*, and an agent who intends to do *A* has practical knowledge that he is going to do *A*. However, well-known counterexamples, devised by Donald Davidson and Michael Bratman, are typically thought to pose a serious problem for this kind of view. Consequently, recent attempts to preserve a defensible kernel of the Anscombean thought have typically proposed weakening the kind of cognition involved in the condition, and/or weakening the condition's content or scope. By contrast, I argue that it is a mistake to weaken the cognition condition in any such way, and that the felt need to do so is symptomatic of a deep failure to understand the possibility of practical knowledge. The true ground and meaning of the cognition condition, properly understood, lies in the (i) *calculative* and (ii) *temporal* structure of intentional action, some brief remarks about each of which are in order.

(i) Many of the things we intentionally do, we do by intentionally doing other things. Many of these other things we do by doing yet further things, and so on. The agent's practical knowledge is the source of this calculative structure; without this knowledge, what happens is not intentional. Thus, it will emerge, an agent intentionally doing something knows not just what he is doing, but how and why he is doing it. This practical knowledge is, in Anscombe's phrase, knowledge "in intention" (1963: §32), not knowledge that inheres in a belief that somehow accompanies intention. Accordingly, an expression of intention is itself an expression of knowledge, of what the agent is actually doing (if it expresses an intention in action, e.g. 'I'm walking to school') or of what the agent is actually going to do (if it expresses an intention for the future, e.g. 'I'm going to wash the car tomorrow'). And just as intentional action is calculatively articulated, so is intention for the

future: an agent who intends to do something knows not just what he is going to do, but how and why he is going to do it.

(ii) An intentional action, like any event, unfolds in time. If what is happening falls under an event concept ϕ , then it will be no accident if that which is ϕ -ing ends up having ϕ -ed. If what is happening falls under an intentional action concept *do A*, then it will be no accident if the agent who is doing *A* intentionally ends up having intentionally done *A*. In the specific case of intentional action, the agent's practical knowledge is the source of this temporal structure; without this knowledge, what is happening is not intentional. Thus, it will emerge, an agent intentionally doing something knows not just what he is doing, but that it will be no accident if he succeeds. Similarly with intentional action in prospect: an agent who intends to do something knows not just what he is going to do, but that it will be no accident if he succeeds.

For one to know that it will be no accident if one ends up having intentionally done what one is doing intentionally (or intends to do), it is necessary, I argue, that one knows how to do what one is doing (or intends to do). In many cases, this demand is supplied by *procedural knowledge*: one knows that one can do *C* by doing *B*. But if one exercises this knowledge, in intentionally doing *C* by doing *B*, then one must know how to do *B*, and though one might know that one can do *B* by doing *A*, it seems obvious that procedural knowledge cannot be the whole story. As Jennifer Hornsby puts it, “[a]mong the things a person knows how to do, some of them he must know how to do ‘just like that’, on pain of needing to ascribe to him indefinitely many distinct pieces of knowledge to account for his ability” (1980b: 88). Hornsby’s suggestion about the structure of practical knowledge mirrors the standard suggestion made in response to a corresponding regress that threatens our thought about the structure of action: though many of the things we do we do by doing other things, and so on, this picture appears to depend on there being basic actions, things we can do directly, without doing them by doing anything else.

However, Michael Thompson and Sebastian Rödl have recently argued that reflection on the calculative and temporal structure of intentional action—the very structure that, I argue, explains and validates the cognition condition, properly understood—reveals that there can be no such thing as basic action. This creates two problems, one metaphysical, one epistemological. The metaphysical problem is to explain how intentional action is so much as possible, if it can appear compelling that basic intentional actions are necessary, as naïve reflection on the idea of doing one thing by doing another suggests, and yet impossible, as deeper reflection allegedly reveals. The epistemological problem is to explain how agent’s knowledge is possible: If there is no such thing as doing something without doing it by doing something else, it seems there can be no such thing as having non-procedural knowledge of how to do something. But in the absence of such knowledge there can be no such thing as knowing that it will be no accident that one will end up having done intentionally what one is intentionally doing (or intends to do), which knowledge is internal to knowing what one is intentionally doing (or intends to do). The resolution of both problems is the same (for they are but the same problem, considered, as it were, from different sides): we need to recognize the way in which one form of practical knowledge—*knowledge in intention*—depends, metaphysically and epistemically (but not inferentially), on another—*skill*.

Thus, in §1 I introduce the idea of a cognition condition on intention and intentional action via Anscombe’s claim and the purported counterexamples to it. In §2, I argue that the true ground and meaning of the cogni-

tion condition begins to come into focus through reflection on the calculative structure of intentional action in progress, and in §3 I show how and why this structure belongs to intention for the future, which is intentional action in prospect. I then argue that the cognition condition can be articulated adequately only by taking into account the temporality of intentional action (in progress and in prospect), and the way in which the temporal structure of intentional action depends on its calculative structure (§4). These reflections enable me to show, in §5, why the purported counterexamples from §1 do not get a grip on the cognition condition, properly understood, and why we need and ought not weaken it in the ways that some have recently suggested. However, the very same reflections manage to suggest both that there must be basic action (§6) and yet that there cannot be (§7), a paradox to whose resolution I point in §8.

1 The cognition condition

1.1 Anscombe's claim

G. E. M. Anscombe famously claims that, if an agent is doing *A* intentionally, he knows “without observation” that he is doing it (1963: e.g. §28). By contrast, if an agent is asked why he is doing *A* and he (sincerely) responds “I didn’t know I was”, then this suffices, she thinks, to show that his doing *A* is not an intentional action. Moreover, even if the agent does know that he is doing *A*, but he knows this only on the basis of observation, then this too suffices to disqualify his doing *A* as an intentional action.

The basic idea here, that when you’re up to something intentionally, *you* don’t have to observe or theorize yourself to know what it is that you’re up to, whereas *someone else* has to find or figure out what it is that you’re up to, is quite intuitive. A hum-drum example: I find you in the kitchen, apparently making breakfast. You take the eggs out of the fridge, and I suppose you’re going to cook them (I know you’re not the type to crack a raw one in a glass, whisk it up, and drink it). But how are you going to cook them? I rule out boiling and poaching when you take out the frying pan: you’ll either fry them or scramble them. And then, when I watch you beat them as the oil heats up, I know you’ll scramble them. But wait! After you pour the eggs into the pan, you don’t start scrambling; and it’s only now that I spot a bowl of grated cheese on the side—you were making an omelette the whole time! Of course, instead of observing your actions, forming hypotheses, drawing inferences, and so on, I could have just *asked* you what you were doing: not only were you making an omelette the whole time, you *knew* you were making an omelette the whole time—you weren’t gradually coming to that conclusion on the basis of the evidence as I was. Anscombe’s point is just this: the concurrence of these two facts—that you were making an omelette the whole time, and that you knew without observation (indeed, without *any* kind of evidence) that this is what you were doing—is no accident.¹

Now, while Anscombe’s claim is usually thought of as applying just to intentional action, it in fact has a wider scope. Readers of *Intention* will recall that the book opens by introducing the topic of intention “under

¹ Anscombe’s illustrations (1963: §28) of the relevant contrast cases—things done in the *absence* of the relevant non-observational knowledge, which absence thereby reveals them to be things done *unintentionally*—ring true too: someone is asked “Why are you ringing that bell?” and replies “Good heavens! I didn’t know *I* was ringing it” (ignorance of the action); and someone, opening a window and thereby focusing a spot of light on the wall, is asked “What are you doing making that light come on the wall?” and replies “Ah yes, it’s opening the window that does it,” or “That always happens when one opens that window at midday when the sun is shining” (knowledge of the action only by observation).

three heads: expression of intention for the future, intentional action, and intention in acting” (1963: iii), and almost immediately Anscombe insists that, temptation to proliferate senses notwithstanding, “it is implausible to say that the word is equivocal as it occurs in these different cases” (1963: §1).² And so her claim, that if someone is acting intentionally, he knows what he is intentionally doing, can presumably be extended to the other two “heads”: an expression of intention for the future (e.g. ‘I am going to do A’) will be an expression of knowledge about the future (i.e. ‘I am going to do A’); and if someone is doing A with the intention of doing B, then he will know, not only that he is doing A, but that he is doing A with the intention of doing B, and thus that he is thereby doing, or going to do, B. In both these cases, too, the knowledge will be “non-observational”, and though it is less clear here what the corresponding observational knowledge would be, our examples have brought out that the kind of knowledge of action in progress or in prospect that Anscombe means to exclude is not observational knowledge narrowly construed (that is, perceptual knowledge), but knowledge *based on any kind of evidence whatsoever*.

The idea of a kind of knowledge not based on any evidence—indeed, *incompatible* with being based on evidence—might already sound suspicious. Perhaps one might cautiously admit the possibility of “groundless self-knowledge”; Anscombe’s claim, however, pertains not to knowledge of a mental state or event, but to the very transpiring of a public, material event: for, if you are making an omelette, then presumably an omelette is being made, and so knowing that you are making an omelette presumably includes knowing that an omelette is being made. It’s one thing to think that there might be groundless knowledge of the mind—but groundless knowledge of *omelettes*?

However, to worry about the details and consequences of Anscombe’s claim at this point would be to get ahead of ourselves. We said that her claim—that if an agent is doing something intentionally, he knows without observation that he is doing it—is intuitive. We did not yet say it is true.

1.2 Davidson’s objection

Donald Davidson thinks that it is not. He thinks that if an agent is doing A intentionally, he need not even *believe*—let alone *know*, let alone know *without observation*—that he is doing it. His well-known counterexample to Anscombe’s claim is this: “A man may be making ten carbon copies as he writes, and this may be intentional; yet he may not know that he is; all he knows is that he is trying” (1971: 50).

The core of the case—the idea of someone trying to do something that he doubts he will pull off and yet succeeding—can easily be adapted to target the corresponding Anscombean claims about the intention with which someone acts, and intention for the future. Davidson later deployed a slight variation of his carbon-copier case that targets the former: “in writing heavily on this page I may be intending to produce ten legible carbon copies. I do not know, or believe with any confidence, that I am succeeding” (1978: 92). And Michael Bratman presents the following case against the latter: “I might intend now to stop at the bookstore on the way home while knowing of my tendency towards absentmindedness.... If I were to reflect on the matter I would be agnostic about my stopping there, for I know I may well forget” (1987: 37).

²On the question of the univocality of ‘intention’, philosophers disagree: while Kieran Setiya claims that “The principal task of the philosophy of intention is to uncover and describe the unity of the three forms” (2009a), David Velleman “giv[es] up, once and for all, on the project of finding a unified analysis for the expressions in which [the words ‘intention’ and ‘intend’] appear (1989: 112–3).

Philosophers disagree as to the force of carbon-copier style counterexamples. Bratman is very impressed by them; he concludes that “such examples are worrisome enough that [one] would do well to develop [one’s] account of intentions and plans in a way that does not require the strong assumption that to intend to *A* I must believe I will *A*” (1987: 38). Kieran Setiya too thinks that these examples pose a difficulty for Anscombe’s claim, though he is less impressed by them than Bratman; he thinks they are “exceptions to [the] rule [that] serve only to confirm the basic insight,” and thus that “the challenge to Anscombe is limited” (2007: 25). Davidson occupies an interesting position. While he holds that intentions are not, and do not include or entail, beliefs, he clearly does not think that agent’s knowledge is irrelevant to action: shortly after presenting the original carbon-copier case in ‘Agency’ he goes on to add that “[a]ction does require that what the agent does is intentional under some description, and this in turn requires...that what the agent does is known to him under *some* description” (1971: 50, emphasis added). But despite these differences, there is some consensus: that, as Setiya puts it, “Anscombe’s doctrine is false, at least when it is unqualified” (2007: 25).³

The idea, then, is that while Anscombe has suggestively pointed towards a putative requirement on intention and intentional action, she seems at best to have overstated it. Let’s call this requirement *the cognition condition*. Now, we could weaken Anscombe’s cognition condition by weakening either the *kind*, the *content*, or the *scope* of the relevant cognition (or some combination of the three). With respect to the *kind* of cognition, we might—out of “modesty,” as Setiya puts it—“[set] aside the claim to knowledge as well as the claim to knowledge without observation” (2007: 25), and cast the cognition condition in terms of what the agent *believes*. This proposal has the advantage of postponing the inevitable difficulties that face accounting for the truth of the agent’s belief that he will in fact do what he intends to do, the justification with which the agent believes that he is doing, or will do, what he intends, and anything else that might be required for “agent’s belief” to amount to agent’s knowledge.

In the light of Davidson’s and Bratman’s counterexamples, however, we might want to modify the *content* of the cognition. One possibility, suggested by Davidson’s initial presentation of the carbon-copier case, would be that if one is intentionally doing *A*, one knows (or believes) that one is *trying* to do *A*.⁴ Setiya makes a different suggestion about the cognition’s content: if one is doing *A* intentionally, “there must be something he is doing intentionally, not merely trying to do, in the belief that he is doing it” (2007: 26). The “something” need be *A* itself only if one’s doing *A* is a “basic action”, which *making ten carbon copies* does not appear to be—the copier may be making the copies intentionally while believing only e.g. that he is writing firmly on the top sheet.

An alternative possibility for accommodating the counterexamples would be to limit the *scope* of the cognition condition. John Gibbons thinks that the cognition condition need not be a *necessary* condition on

³ While Velleman thinks that carbon-copier cases fail in their aim, he too thinks Anscombe’s view must be qualified—indeed, in a way that inverts Davidson’s position: Velleman thinks that intentions are beliefs about the future, but that intentional action needn’t involve spontaneous agent’s knowledge (see e.g. 1989: 112–7; 2007: xv, n.7).

⁴ I hesitate to attribute the view to Davidson himself; in ‘Freedom to Act’, he writes: “[Armstrong] asks the question what we must add to ‘*A* tried to do *x*’ in order to have necessary and sufficient conditions for ‘*A* did *x* intentionally’. Most of us would say—I certainly would—that trying itself isn’t necessary in many cases...” (1973: 78). It is clear that Davidson thinks an agent must know some, but not all, of the descriptions under which his action is intention; it is less clear how to figure out which descriptions, if any, Davidson thinks are ones the agent *must* know, or if there is a principle that accounts for this.

intentional action. This is not simply because “we shouldn’t expect to find such a thing, given the history of [such conditions’] failure,” but rather because the important questions about “privileged access”—whether to our intentional actions or to anything else—are “what kind of access, how it works, and why we have it” (2010: 84). Crucially, “we shouldn’t expect an account of how you know when things go smoothly to deliver a non-trivial answer to the question of when things will” (84)—things don’t go smoothly for the carbon copier, but such cases are outliers that don’t impugn the account, providing it is properly understood. Similar permutations (kind, content, scope) arise when considering whether there is a cognition condition on intention for the future, and what form it takes if there is one.⁵

According to the account developed in this essay, all such weakening and finessing is a mistake. Once we have uncovered the real ground and meaning of the cognition condition, which lies in the calculative and temporal structure of intentional action, we will see that the carbon-copier example poses no threat, because it trades on confusion as to how to understand the possibility of failing to act as one had intended.

1.3 The contemporary debate and beyond

Before proceeding to my defence of the cognition condition, it is worth mentioning the way in which contemporary accounts of action and intention deal with the issue. *Cognitivists about intention*, like Velleman and Setiya, are primarily impressed by the fact that, like belief, an expression of intention, such as ‘I am walking to school’ or ‘I am going to walk to school’, makes a claim not merely about the speaker’s state of mind, but about what is actually happening, or going to happen, in the world. Finding the similarity overwhelming, cognitivism takes the bold course of *identifying* intention with belief—albeit belief of a special kind. What is special about an intention is that, unlike a “regular” belief about the future, it (or its expression)⁶ represents, and in the good case represents correctly, the future as causally dependent on the belief.

By contrast, *non-cognitivists about intention*, like Davidson and Bratman, reject the assimilation of intention to belief, which they consider to be a theoretical attitude. Instead, they insist that intention is a special kind of “pro-attitude”, a distinctively practical attitude that differs from mere desire by representing one’s course of action as “settled”; whereas mere desires *only potentially* influence conduct, intentions *actually control* it. As we saw above, Bratman takes carbon-copier-style counterexamples to decide the issue in favour of non-cognitivism. But in any case, if intention is a pro-attitude, then as such it has a different direction of fit from cognition: it must represent the world not as it is, but as it would in some sense be good to be, or it must at least consist in its bearer’s being disposed to act so as to bring the world into alignment with the representation. Thus what is practical cannot be knowledge, and what is knowledge cannot be practical.⁷

⁵It is worth noting, however, that while even Davidson accepts a weak version of the cognition condition on intentional action, intention for the future can appear to be yet further removed from cognition: Bratman’s bookstore case supposedly exemplifies intention not merely in the absence of belief in *success*, but in the absence of the belief that one will even *try*. But this is not necessary; it is easy to construct an example that simply transposes the features of Davidson’s carbon-copier example to intention for the future (indeed, Bratman himself does so). See Bratman 1987: 35–42.

⁶I add this qualification to leave open the possibility that, as surface grammar suggests, the content of an intention is predicative (e.g. *to walk to school*), whereas the expression of intention is propositional (e.g. *I intend to walk to school*, *I’m going to walk to school*, *I am walking to school*, etc.).

⁷See Bratman 1987: 15–20; cf. the closing pages of Davidson 1978, where Davidson distinguishes between desires, conceived of as expressed by “prima facie judgements”, and intentions, conceived of as expressed by “all-out judgements”. While Davidson

From the cognitivist point of view, however, this stricture reflects only a lack of imagination. According to Velleman, intention has “the same direction of fit as belief but the same direction of guidance as desire,” where direction of guidance “consists in whether the attitude causes or is caused by what it represents”; in virtue of combining the cognitive quality of belief with the practicality—the causal efficacy—of desire, intention is “a case of practical cognition” (2000: 25).⁸ By contrast, the account of practical knowledge advanced here is based not on considerations pertaining to the admissible combinations of directions of fit and guidance that a given type of mental state might possess, but rather on figuring out how intentional action could have the structure that it has.

The contemporary debate between these accounts proceeds against the backdrop of “epistemological modesty” sketched above: the idea that, if there is a cognition condition at all, it must be construed, in the first instance, in terms of what the agent believes. (If such a construal is successful, the questions whether and how “agent’s belief” amounts to agent’s knowledge can be addressed.) Against this backdrop, it is easy to picture Anscombe as holding an “immodestly” strong version of cognitivism. But we might well wonder whether this is what she has in mind, given her ultimate characterization of the cognition condition in terms of practical knowledge, by which she is widely acknowledged to have meant more than: a mental state that is the efficient cause of intentional action.⁹ On one traditional conception of things, theoretical reason is the power of theoretical knowledge, practical reason the power of practical knowledge.¹⁰ And we might say, adopting the old terminology, that, on the one hand, a belief is an act of theoretical reason, the power of theoretical knowledge, such that if a belief is a proper, unimpeded exercise of that power, it is a case of theoretical knowledge; whereas on the other hand, an intention is an act of practical reason, the power of practical knowledge, such that if an intention is a proper, unimpeded exercise of that power, it is a case of practical knowledge. Anscombe says that practical knowledge is knowledge “in intention” (1963: §32); we

situates both desires and intentions within “the genus of pro attitudes expressed by value judgements” (1978: 102), it is worth noting that Bratman explicitly disavows the “tendency to see intention as either a special kind of evaluation or a special kind of belief” (1987: 110). I am inclined to think that the deepest motivation for a non-cognitivist view lies the traditional thought that practical reasoning is distinguished from theoretical reasoning in taking as its formal object the good rather than the true. Unfortunately, a discussion of such a non-cognitivism, which in any case is not Bratman’s, lies beyond the scope of this essay.

⁸On Velleman’s view it is crucial not only that an intention is the *cause* of what it represents, but also that it *represents itself* as that cause (see Velleman 1989: 102, 2000: 25). In recent work, Setiya has distanced himself from the idea that the kind of causally-efficacious belief in which, he still believes, an intention consists *represents itself* as the cause of the intended action; [...practical reasons]] [CITE old and new].

As our earlier discussion indicated, weaker versions of cognitivism about intention are possible: for example, the view that the intention to do *A* is the belief, of the special kind, that one will try to do *A*, or that one will do something with the aim of doing *A*, or try to do something with the aim of doing *A*, etc.

⁹The thesis that Anscombe is rightly understood as claiming that practical knowledge is the *formal* cause of what it understands, rather than merely its efficient cause, is elaborated and defended at length in Moran 2004; Velleman “doubt[s] whether Anscombe takes the word ‘cause’ seriously. The connection that she posits between the agent’s knowledge and his action seems conceptual rather than causal” (1989: 102 n. 24). Anscombe claims that we cannot make sense of practical knowledge without making sense of practical reasoning (1963: §33). Moreover, in her caustic remarks on the “mince-pie syllogism” she disparages the idea that practical reasoning is just theoretical reasoning about a special topic (i.e. action) or involving special concepts (e.g. ethical concepts): her examination of Aristotle’s practical syllogism leads her to conclude that there is “a difference of form between reasoning leading to action and reasoning for the truth of a conclusion. ...[T]he conclusion [of practical reasoning] is an action whose point is shewn by the premises” (1963: §33).

¹⁰This traditional conception is found as late as Kant, according to Stephen Engstrom (2009), who emphasizes the continuity of the view as it appears in Kant with Aristotle and medieval philosophy; these latter, of course, are those to whom Anscombe credits, in the remark that appears as the epigraph to this essay, the concept of practical knowledge—“blankly misunderstood” by modern philosophy—which she seeks to rehabilitate (1963: §32).

might say, following this line of thought, that theoretical knowledge is knowledge “in belief”. Thus, in locating the perhaps-defensible kernel of her claim in the idea that a necessary condition on intention and intentional action is that an agent has certain *beliefs*, we run the risk of fundamentally misconstruing her view with an unconsciously prejudiced application of the principle of charity.

Of course, a hasty sketch of long-forgotten faculty psychology and the quite possibly dubious invocation of some scholastic terminology proves nothing; I bring it up only to raise suspicion of the quick and ready assimilation of Anscombe’s view to contemporary cognitivism about intention and to suggest an alternative strategy, both for interpreting Anscombe and for offering an account of agent’s knowledge. According to the broadly Anscombean view I develop in §§2–4, there is a cognition condition on intention and intentional action that is to be explicated not in terms of belief, but rather in terms of the cognitive character of intention itself—in terms, that is, of practical knowledge. This explains why I speak of a *cognition condition* (rather than a *belief condition*): the label is intended to be neutral not merely between knowledge and something less than that (belief, on the standard conception), but between the different types of “acts of reason” in which knowledge might “inhere”.

The idea that practical knowledge is practical in virtue of its form, where this means more than merely having certain kinds of effect, is far from clear. I propose to approach its clarification obliquely, via an initially obscure distinction Anscombe draws between an intention’s being contradicted and its “falling to the ground”:

...the contradiction of ‘I’m replenishing the house water-supply’ is not ‘You aren’t, since there is a hole in the pipe’, but ‘Oh, no, you aren’t’ said by someone who thereupon sets out e.g. to make a hole in the pipe with a pick-axe. And similarly, if a person says ‘I am going to bed at midnight’ the contradiction of this is not: ‘You won’t, for you never keep such resolutions’ but ‘You won’t, for I am going to stop you.’ (1963: §31)

Now suppose what he says [i.e. ‘I’m replenishing the house water-supply’] is not true. It may be untrue because, unknown to the agent, something is not the case which would have to be the case in order for his statement to be true; as when, unknown to the man pumping, there was a hole in the pipe round the corner. ...[If this is the case], we may say that in face of it *his statement falls to the ground*, ...but it is not a direct contradiction. (1963: §32, emphasis added)

One would think, from these passages, that ‘You aren’t, since there is a hole in the pipe’ and ‘You won’t, for you never keep such resolutions’ are on a par with one another, and thus that if in the face of the former the intention (in action) to replenish the house water-supply falls to the ground, then so in the face of the latter the intention (for the future) to go to bed at midnight will likewise fall to the ground. But this seems faintly absurd. Maybe I do never keep such resolutions, but there’s a first time for everything, and so long as you’re not going to stop me, why shouldn’t I intend to go to bed at midnight? And for that matter, if I’m employing you to replenish the house water-supply and find, at the end of the day, that you gave up your task, and even stopped intending to carry it out, in the face of a simple hole in the pipe—well, I’ve every right to be annoyed: ‘Why didn’t you just plug it?’, I might say. Perhaps, though, the pipe is well and truly broken; it can’t be repaired. If it’s *impossible* to replenish the water supply, surely this makes sense of the idea that the intention to do so “falls to the ground”. But on reflection, this is just as much of a dodge as before, and you won’t get the afternoon off work so easily! The pipe can’t be repaired, so the intention to *replenish the water-supply by pumping water through the pipe* cannot be executed, and thus falls to the ground; but the inhabitants need

running water and there are two buckets over there—get to it! There may be many different ways to replenish the water supply; if so, there will be many ways of filling in one's intention to do so. After all, if we gave up on everything at the first sign of trouble, we'd never get anything done.

And we should also note something else. If you spent the morning pumping fruitlessly, as it turned out, and the afternoon lugging buckets of water up to the house, then what you were doing all day was: *replenishing the water supply*. Of course, there are other ways of describing what you did with your day, notably *pumping fruitlessly and then lugging*, but *replenishing the water supply* applies as well. ('Why didn't you cut the grass as you were supposed to?'—'Because it took all day to replenish the water supply'.) After all, as Michael Thompson points out, "we happily affirm, of someone who is napping, that she is organizing the peasantry; of someone who is sitting reading the paper, that she is baking a loaf of bread; and of someone who is playing a hand of poker, that she is building a house. If confusion arises, we...concede that our agents aren't baking or building or organizing *at the moment* or *right now*, but rather reading, playing poker or napping" (2008: 141). Thompson's examples highlight an important feature of progressive judgments—judgments, that is, of the form 'S is (was) doing B'—namely, their "broadness".¹¹ It might take me ages to get to your house for our rendezvous because I take a wrong turn, which eventually leads me back to my house, whence this time I take the correct route to your house. It would be mere pedantry on your part to insist that it *didn't* take me ages to get to your house, that I just *left late*.

So it might well be that the pumper's intention to replenish the water supply falls to the ground in the face of the hole in the pipe. But he can pick it back up again. And if he does, his statement, 'I am replenishing the water supply', will not be untrue.

Notice the difference between the following two relationships: on the one hand, that between the intention to replenish the water supply and the fact that there's a hole in the pipe, and on the other, that between the belief that the pipe has no holes and the fact that there's a hole in the pipe. If our pumper finds out that there's a hole in the pipe, he must give up, on pain of irrationality, the belief that it has no holes; but he can keep the intention to replenish the water supply, as long as he has, or can come up with, a "Plan B" for its execution.

But now compare the case of *opposed intentions*: X says 'I'm replenishing the water supply' and Y says 'No you're not—I'm going to stop you'. X's intention to replenish the water supply cannot be retained as easily in the face of Y's threat as it could in the face of the hole in the pipe. The problem with the hole in the pipe is that it is preventing X from replenishing of the water supply by means of operating the pump. The hole in the pipe is not going to prevent X from trudging up to the house with buckets of water from the well. By contrast, Y's intention to stop X from replenishing the water supply *does* track X's intention through Plans B, C, and so on. Y's intention to prevent the water supply's replenishment is an intention to *act*, to *do various things*—make a hole in the pipe, throw the buckets down the well, tie X up, and so on—in order to stop X from doing the various things he needs to do in order to replenish the water supply.¹²

X's intention to replenish the water supply can be made consistent with the fact that there is a hole in the pipe, whereas X's belief that the pipe has no holes cannot be made consistent with that fact, and X's and Y's

¹¹ On the broadness of the progressive, see also Galton 1984 and Falvey 2000.

¹² Cf. Moran and Stone 2009: 155–6.

intentions cannot both be successfully executed. These differences illuminate a famous but cryptic remark of Anscombe's, which appears as an epigraph to this essay, in which she complains of modern philosophy's "incurably contemplative conception of knowledge," according to which "the facts, reality, are prior" to the knowledge, which is "judged as such by being in accordance with" those facts (1963: §32). This is how things are with the belief that the pipe has no holes and the hole in the pipe. The belief is false because of how things stand with the pipe, which is prior to the belief. But X's expression of intention "I am replenishing the water supply" is not falsified by the hole in the pipe. In face of the hole in the pipe, the intention falls to the ground. Whether it is picked up or not depends on *what X does*, on whether he seeks out alternative means or gives up. If he seeks out alternative means, his expression of intention is true. If he has given up, then presumably he no longer has the intention, and thus wouldn't express it. But suppose he were to continue to give (linguistic) expression to his intention, while doing nothing whatsoever to execute it; here we can say that his expression of intention is false—he is *not* replenishing the water supply. However, the assessment of his expression as false does not depend on a failure of fit between what he says and "the facts" (i.e. the hole in the pipe), but rather on a failure of fit between what he says and *what he does*: he says he is replenishing the water supply, but he is doing nothing to execute his intention (and he's not just taking lunch).¹³ For his expression of intention to amount to knowledge, it must be true that he is replenishing the water supply. (No knowledge without truth.) But whether or not he is replenishing the water supply is not something that can be ascertained independently of his knowledge of what he is doing and by what means he is doing it.

The foregoing considerations, the purpose of which was to elucidate, preliminarily, some ways in which intention and practical knowledge might differ from belief and theoretical knowledge, have uncovered, albeit haphazardly, the two (interdependent) structures of intentional action of which we need a more detailed and systematic investigation: (i) the *calculative structure* of action—the structure involved in, intentionally, doing one thing by doing another thing, and (ii) the *temporal structure* of action—the relations that hold between what an agent is doing, what he is going to do, and what he ends up having done. The grounds for thinking that there is a cognition condition on intention and intentional action go deeper than noticing that agents can say what they're doing without having to look, and that when they say what they're doing or going to do, their utterances are apparently indicative; the grounds lie in the structure of intentional action itself.

2 Doing one thing by doing another

I will argue that the source of the cognition condition lies in the calculative structure of intentional action—the structure of means and ends represented as such by the agent that is internal to intentional actions. By contrast, Davidson and his followers understand the structure of intentional action in terms of the causal relations that an action, which itself is internally structureless, bears to other events, in virtue of which relations there is structure in action *descriptions*. If the source of the cognition condition lies in the inner

¹³It would be interesting to explicitly compare the conception of "direction of fit" in play here with the conception that is usually employed in the contemporary literature (the conception according to which if S desires that *p*, but *p* is not the case, then somehow the world is to blame, or in the wrong). Unfortunately I cannot pursue this topic here; see Moran and Stone 2009 and Boyle and Lavin 2010 for subtle and important discussions.

structure of intentional action, as I will argue that it does, and if this structure is missing in the Davidsonian account, then it will be no surprise that the cognition condition cannot be adequately comprehended within the Davidsonian framework that still dominates contemporary philosophical approaches to action. Thus it is crucial for my subsequent account of the cognition condition that the Anscombean and Davidsonian conceptions of the structure of action are accurately distinguished.

I begin with some common ground between Anscombe and Davidson, concerning the individuation of actions. Suppose a man raised a paddle, and thereby lodged a bid and cast a shadow on his neighbour. He did three things; let us suppose he did the first two intentionally, the third unintentionally. But how many actions did he perform? Anscombe, Davidson, and the many who have followed them, say *one*: in saying he did three things, we mean that there are three descriptions that truly apply to his action. (The other answer, which was advanced by Alvin Goldman (1970) but has few adherents these days, is of course *three*.)¹⁴ Obviously, the example I provided is grossly simplified: there are indefinitely many descriptions that truly apply to any action, most of which will be quite *recherché* and of no interest to anyone. But among the chaos of these indefinitely many descriptions, there is order to be found, and Anscombe and Davidson each try to find it. Or rather, I should say, there are *orders* to be found.

2.1 Davidson and the accordion-effect

It follows from the Anscombe–Davidson account of action-individuation that while the concept *action* can delimit a class of events, the concept *intentional* cannot delimit a class of actions. The action in our example was intentional under the descriptions “raised a paddle” and “lodged a bid” and unintentional under the description “cast a shadow on his neighbour”. An event belongs to the class of actions if it is intentional under *some* description, and we have at least one description under which it is intentional, so the event described in the example is an action. But, if we were to ask whether “the action itself” belonged to the class of intentional actions, we would have to say both yes and no, which is absurd, showing that the question makes no sense and that there can be no such class.

In his well-known paper ‘Agency’, Davidson concludes from the fact that the same action—the same event—can be intentional under one description and unintentional under another that “although the *criterion* of agency [i.e. being intentional under a description] is [intensional], the *expression* of agency is itself purely extensional.” This suggests to his Quinean side that “the concept of agency is simpler or more basic than that of intention”; thus he looks for “a mark of agency that does not use the concept of intention,” hoping that “[t]he notion of cause may provide the clue” (1971: 47). This is the context in which Davidson provides his most in-depth analysis of the orders that can be discerned in the welter of descriptions of an action.

The clue Davidson finds in the notion of cause is this: actions can be described in terms of their effects, and there will be an order to these descriptions determined by the causal chains running out from the action.

¹⁴The Anscombe–Davidson line on this is, I think, obviously correct; but it is easy to overstate: suppose a man slices some bread, then butters it, then slices some cheese, assembles these ingredients, and thereby makes a cheese sandwich. His slicing of the bread isn’t the same action as his slicing of the cheese, and neither slicing is identical with his making of a cheese sandwich. The two slicings, the buttering, and the assembling amount to four actions; but the making of a sandwich is not a fifth (as Goldman would have it, wrongly again): rather, the making of the sandwich is identical with the sum of the four sub-actions or parts. ‘...and then...’ and ‘...and thereby...’ denote quite different relations; the Anscombe–Davidson line is appropriate for ‘...and thereby...’, but we must take care when thinking about ‘...and then...’. See my discussion of the part–whole relation in §2.2 below.

Roughly speaking, he thinks that what it means to say that *S* did *B* by doing *A* is that *S* did *A*, and *S*'s doing *A* caused an event in virtue of which *S*'s doing *A* may be legitimately redescribed as *S*'s doing *B*.¹⁵ Now, *S*'s doing *A* might have caused the event that licensed that redescription by way of causing any number of intermediate events. According to what Davidson (following Joel Feinberg) calls “the accordion effect”, we can leave out links of the causal chain; as he puts it, “[t]here are...a great many tunes we can play on the accordion” (1971: 58). Thus if the queen killed the king by poisoning him by emptying a vial of poison into his ear by moving her hand, then it would be true—if opaque—to say that she killed him by moving her hand.

The chain must have a beginning: “[n]ot every event we attribute to an agent can be explained as caused by another event of which he is the agent” (1971: 49). Thus there must be *basic actions*, of which Davidson famously said: “our primitive actions, the ones we do not do by doing something else, mere movements of the body—these are all the actions there are. We never do more than our bodies: the rest is up to nature” (1971: 59). On Davidson’s view, then, *all* actions are basic actions, bodily movements with no intrinsic structure *qua* actions. It follows that an *action-description* may very well be complex—e.g. ‘killing the king’ means ‘doing something which caused (something which caused something which caused something which caused...) the king’s death’. But the *action it describes* is simple; it’s the first event in the chain. Complex descriptions *apply* to actions, but only in virtue of causal relations that hold between an action and other events it causes—relations *extrinsic* to the action itself (cf. Davidson, ‘Adverbs of Action’, 301). The action itself, under its most primitive description, is a bodily movement, which goes unanalyzed by Davidson as a simple element of his theory.

The causal interpretation of the ‘by’-relation cannot, however, be exhaustive, for it may denote a constitutive connection, too.¹⁶ In such cases, *S*'s doing of *A* is his doing of *B* not in virtue of causing another event that licenses redescribing his doing of *A* as his doing of *B*, but rather in virtue of conventional, circumstantial, or metaphysical facts concerning the *A*—*B* connection that license the redescription. For instance, a man may vote by raising his hand (given certain conventions), take a giant leap for mankind by taking a small step for man (given special circumstances), or get some Vitamin X in him by eating pig’s tripe (in virtue of the latter containing the former). The ‘by’-relation is determinable, and can be determined in these ways.¹⁷ Like the causal determination of the ‘by’-relation, each of these kinds of determination too can be instantiated by an agent intentionally, unintentionally, or even unknowingly. So even if we enrich Davidson’s account by adding to it these other kinds of ‘by’-relation, it seems that the structure of action—the order of actions and descriptions of action—still concerns things done (and connections between them) that may be unintentional and unknown, and thus that investigating the structure of action has no bearing on the question of whether there is a cognition condition on intentional action; such is the price of an extensional analysis.¹⁸

¹⁵For example, Davidson says that ‘The queen killed the king by pouring poison in his ear’ and ‘The queen poured poison in the king’s ear thus causing his death’ “are equivalent” (1971: 59).

¹⁶As Davidson concedes, his analysis “works for some cases only” (1971: 55).

¹⁷Other versions of the metaphysical constitutive connection: one may tie a knot by tying a reef knot, which is a *kind* of knot, or cook fish by frying it, which is a *way* of cooking fish. Anton Ford (ms.) calls this determination of the ‘by’-relation a “species–genus” relation. On the variety of determinations of the ‘by’-relation, see also Baier 1970: 648–9, Baier 1971, and Müller 1979: 106.

¹⁸I should note that Davidson himself considers his analysis of agency successful only with heavy qualifications. Yet he thinks that it leads “to a vast simplification of the problem of agency, for it shows that there is a relation between a person and an event, when it is his action, that [because all actions are primitive] is independent of how the terms of the relation are described” (1971: 61). Of course, the analysis of the relation between an agent and his action will depend on intention, but the concept of intention has been

2.2 The calculative structure of intentional action

Though both Davidson and Anscombe investigate what structure there is to be discerned in the descriptions that apply to an action, Davidson does so by locating action in a structure, while Anscombe locates structure in an action. The structure that she finds there is *calculative*, one of means and ends represented as such by the agent. The calculative ‘by’-relation may be instantiated by any of the connections we have encountered; however, it is not instantiated by those connections insofar as they are not represented as means to end by the agent—that is, insofar as they are not intentional. To our list of determinations of the ‘by’-relation we may now also add the part–whole relation: one may do *C* by doing *B'* and then doing *B''* and then.... Unlike our other relations, there is more to doing *C* than doing *B*—more, that is, for the agent to *do*. Davidson’s system is essentially hostile to this relation, in a way in which it is not to the other additions: the part–whole relation brings out that actions may be *composed of other actions*. The relationship between an action and a proper part of it (itself an action) cannot be cast in terms of two descriptions of the very same action. Thus this relation captures what Davidson is blind to: the fact that actions themselves—not just their descriptions—are structured.¹⁹

We can illustrate the calculative structure of intentional action with Anscombe’s famous example, a portion of which I exploited earlier, in which a man is moving his arm (doing *A*), operating a pump (doing *B*), replenishing the house water-supply (doing *C*), and poisoning the inhabitants (doing *D*). These things are connected as a series of means to his end, poisoning the inhabitants: he is poisoning the inhabitants by replenishing the water supply, which he’s doing by operating a pump by moving his arm. Thus we can move down the ‘by’-chain from *D* to *A* by asking ‘How?’, where the sense of the question asks for the agent’s means: ‘How are you doing *D*?’—‘By doing *C*’; ‘How are you doing *C*?’—‘By doing *B*’; ‘How are you doing *B*?’—‘By doing *A*’. Conversely, we can move up the chain from *A* to *D* by asking the corresponding question ‘Why?’, which asks for the agent’s reason for acting: ‘Why are you doing *A*?’—‘In order to do *B*’; ‘Why are you doing *B*?’—‘In order to do *C*’; ‘Why are you doing *C*?’—‘In order to do *D*’.

It may be surprising to learn that Anscombe’s question ‘Why?’, which she spends so much of *Intention* elucidating, corresponds to the question ‘How?’, which receives comparatively little overt attention.²⁰ In fact, however, it is simply another way of putting one of her central points, which is not explicitly made until towards the end of the book, that the order of practical reasoning—the order of calculation from “the thing wanted” to “the immediate action” (1963: §41), the order, that is, of *how* to achieve one’s end—and the order of action explanation—the order defined by the question ‘Why?’—are “the same order” (1963: §42). In light of this, if we say that the expression ‘In order to do...’ gives the intention with which the agent acts, we can say that the man is moving his arm (doing *A*) intentionally with the intention of operating the pump (doing

successfully expelled from the analysis of the structure of the action itself. Consequently, a case for a cognition condition would have to be made on other grounds than considerations arising from the structure of action; and the most promising grounds would seem to be psychological. This would be a reason to turn away from the question of a cognition condition on intentional action, towards the question whether intention involves belief.

¹⁹Note that the part–whole ‘by’-relation too has non-intentional instances; only its intentional instances, of course, instantiate the calculative relation. In his later paper ‘Intending’, Davidson *does* consider an example of an action (writing the word ‘action’) “[s]ome temporal segments of [which] are themselves actions: for example, first I write the letter ‘a’” (1978: 88). But the incompatibility of such cases with the doctrine of ‘Agency’ passes without mention, and a positive account of how the parts are unified in the whole is not provided.

²⁰Anscombe notes the correspondence briefly at 1963: §26.

B). And indeed, reminiscently of the accordion-effect, we can say, if obscurely, that he is moving his arm with the intention of poisoning the inhabitants of the house, or that he is poisoning the inhabitants by moving his arm (Anscombe 1963: §26).

Now, it might be thought that the “operations” performed on the A—D order by the questions ‘How?’ and ‘Why?’ are not symmetrical, because whereas both the question ‘How are you doing B?’ and the answer ‘By doing A’ mention *doings*, when the question ‘Why?’ is asked of doing A, the response ‘In order to do B’ apparently does not mention a doing—in other words, it doesn’t say that B is *being done*. However, Anscombe points out that, if the agent is doing A intentionally with the intention of doing B, then he *is doing B intentionally*, provided **two conditions** are met: (i) he is in fact doing B, and (ii) his doing B is not an accidental consequence of his doing A (1963: §23); and thus the symmetry of ‘Why?’ and ‘How?’ is secured. That is, the same calculative nexus of doing A and doing B is captured, from different angles, so to speak, by the following exchanges: ‘Why are you doing A?’—‘Because I’m doing B’, and, ‘How are you doing B?’—‘By doing A’. Employing the calculative senses of ‘by’ and ‘because’ just elucidated, we can say that ‘I am doing A because I am doing B’ and ‘I am doing B by doing A’ are but different expressions of the same thought, as are ‘At Plataea the Greeks defeated the Persians’ and ‘At Plataea the Persians were defeated by the Greeks’.

2.3 Doing A in order to do B, but not doing B

Anscombe claims that if an agent is doing A intentionally with the intention of doing B, he is doing B intentionally, as long as two conditions are met. **Condition (i)** appears obvious: the agent must actually be doing B. Now, one way in which an agent can be doing A in order to do B, but not be doing B, is if he is not doing B *yet* (buying eggs in order to make an omelette, for example). However, Anscombe claims that the distinction between, on the one hand, cases of which we will happily grant that, in doing A, an agent is (actually) *doing B*, and, on the other, cases of which we will say only that, in doing A, he (merely) *wants* or *intends* or is *going to do B*, is not a sharp one:

is there much to choose between ‘She is making tea [by putting on the kettle]’ and ‘She is putting on the kettle in order to make tea’—i.e. ‘She is going to make tea’? Obviously not. And hence the common use of the present to describe a future action which is by no means just a later stage in activity which has a name as a single whole. E.g. ‘I am seeing my dentist’, ‘He is demonstrating in Trafalgar Square’ (either might be said when someone is at the moment e.g. travelling in a train). (1963: §23)

I will discuss the questions whether and how the achievement of the objective (doing B) can be, and be known to be, somehow “present” in what is happening (doing A) in some detail in §4 below.²¹ For now, I will consider cases in which an agent is doing A in order to do B, but not doing B, where this is not simply because he cannot be said to be doing B *yet*, but rather because something is going wrong with his attempt to be doing B *now*.

Earlier, we considered the case in which the pumper is intentionally pumping in order to replenish the water supply, but he is not replenishing the water supply because of a hole in the pipe. We said that his intention to replenish the water supply by pumping thus “falls to the ground”; pumping, *in these circumstances*, is in fact not a way in which to replenish the water supply—it’s not a *part* or *species* of replenishing the water

²¹See in particular n. 74 and the surrounding discussion.

supply, it doesn't *constitute* replenishing the water supply in virtue of any special circumstances or conventions, and it's not a *cause* of the water supply's being replenished. We also said that the pumper could "pick up" his intention to replenish the water supply: either by doing something in order to restore the possibility of pumping-in-order-to-replenish-the-water-supply (e.g. fixing the hole in the pipe), or by taking alternative means to replenishing the water supply.

However, might not condition (i) go unsatisfied in a different way? Suppose there is no hole in the pipe; there is no problem with the pumper's intention to replenish the water supply by means of pumping. And indeed, he is pumping (operating the pump) intentionally. But suppose that his stroke is too shallow; it doesn't draw the water, and in fact he is only pushing air through the pipe. In this case too, he is not replenishing the water supply—not right now, anyway. But here there is nothing wrong with his *intention* (to replenish the water supply by pumping); it is his *execution* (his pumping) that is at fault: "the mistake is in the performance," Anscombe famously says (1963: §45), following Theophrastus.²² However, it seems as if we can push on this case to bring it closer to the previous one (the hole in the pipe). For, we can ask the pumper *how* he is pumping, and thereby get him to further articulate his intention: to replenish the water supply by pumping, by making *these* movements with his arm. And now, just as pumping is not a way of replenishing the water supply if there's a hole in the pipe, so making *those* movements with his arm is not a way of pumping; again, he needs to respecify his intention. In both cases, what he's doing *isn't working*, though there remains a difference: in the one case, the pump-pipe system isn't working, in the other, the pumper—his pumping—isn't working.²³

Examples such as these illustrate what is invariably the case: that there is a kind of constant correction that goes on throughout the course of an intentional action, as the agent responds to the miniature successes and failures, obstacles and alternative possibilities both foreseen and unforeseen, that he encounters in what he's doing and what he's acting on and with. What I am bringing out is that this constant correction amounts to *rationally respecifying one's intention*: the answer to the question 'How?' is constantly being finessed. This should be familiar to all of us who produce things like arguments, the construction of which involves many false starts and pursuits of different avenues, the argument only gradually taking on a determinate shape. Now, practical deliberation just is rationally specifying an intention (for the future or, as emphasized here, in action) in the light of one's ends, powers, and circumstances. Thus we can see that practical reasoning continues for the duration of the action; it is not over and done with once one, say, decides that one will replenish the water supply by operating this pump.²⁴

²²Anscombe cites *Magna Moralia*, 1189b 22. In saying that there is something wrong with the pumper's intention in the case where he is pumping fruitlessly (because of the hole in the pipe), I do not mean to imply that the pumper is to blame (though he might be, if e.g. the maintenance of the pipe was his responsibility, or if the pipes hadn't been used for some time, etc.). His intention just needs fixing ("picking up").

²³Thanks to Anton Ford for this way of putting things, and for helpful discussion here.

²⁴In saying that practical reasoning is the rational specification of an intention, I am not denying that it concludes in intentional action; rather, I am elucidating a conception of practical reasoning and intentional action according to which whenever a piece of practical reasoning concludes in an intentional action, it does so by specifying an intention, and respecifying that intention in action (as necessary) throughout the action's duration. The connection between the calculative structure of an intentional action and the action's duration will become important in my investigation into the temporal structure of intentional action in §4 below.

I am suggesting that when someone, for instance, finds, as he is stirring the milk into his coffee, that he has not added quite enough and so adds in a little bit more, this “on-the-fly” correction amounts to a rational respecification of his intention in action and thereby to a piece of practical reasoning. One way in which someone might try to explicate this suggestion is with the idea that, running through the agent’s stream of consciousness, alongside his bodily activity, is some such train of thought as “This coffee isn’t reaching the desired colour; it would reach the desired colour if I added in a little bit more milk; so let me add in a little bit more milk”. It is not phenomenologically implausible to suppose that, on some occasions, agents do accompany these kinds of corrections with such dialogues, in their heads or even out loud. But it would be wildly implausible to suppose that they always do this; our agent might have a bit of song running through his head as he adds a little more milk in order to make it milky enough, or he might be busily engaged in a piece of deliberation pertaining to something quite removed from the business of making coffee. (In denying that practical reasoning must be separate or separable from the intentional action whose structure it explicates I am not, of course, denying that it *can* be separate, that one can come up with a plan of action well before one begins to act.) Moreover, explicating the suggestion in this way would make it mysterious why anyone could think that practical reasoning concludes with intentions such as “to make a cup of coffee” at all.

This might incline one to suppose instead that the practical reasoning in which this respecification consists need not be conscious. But it is not as if, were we to ask the agent why he was putting some milk into coffee into which he had already added milk, he would not be able to simply say something like “Because it wasn’t yet milky enough for my taste” and thereby express, if perhaps only partially, his intention. It is plausible to think that an “unconscious intention”, if there could be such a thing, could not be expressed by simply avowing it in this way, and would need to be discovered by its subject, as precisely our man does not need to do.²⁵ That his intention to add a little more milk in order to make it the right colour is conscious is *shown* by the fact that he can simply say, without having to find out, that this is what he is doing.²⁶ This in turn shows that what we need to do is re-think what it would be for his intention in action, and the practical reasoning that respecifies it as his intentional action progresses, to be conscious. The sensitivity of an intentional action to its own progress is a mark of its intelligence, yet there is no need to picture this sensitivity—this rational respecification of the intention in action—as something standing outside the intentional action itself, as Gilbert Ryle points out: “When I do something intelligently, i.e. thinking what I am doing, I am doing one thing and not two. My performance has a special procedure or manner, not special antecedents” (1949: 32).²⁷ The “special procedure or manner” that belongs to intentional action as such is that of being prosecuted in accordance with the form of Anscombe’s A—D order: a means is a means relative to an end just in case it is sufficient to produce that

²⁵ As will be clear from what is to come, I do not think there can be unconscious intentions. I do not, of course, deny that there are unconscious mental states; but intention is, as e.g. desire is not, essentially a rational and self-conscious act of the will.

²⁶ On the connections between expression, avowal, and consciousness, see Finkelstein 2003. Finkelstein argues, plausibly, that the ability to express a mental state by avowing it is a necessary mark of the state’s being conscious.

²⁷ Ryle objects to the idea that the rational respecification of the intention in action must be conceived of as a special *antecedent* to the action that it is in; his remark is naturally taken as rejecting a picture on which the intention (or intentions, if that is how the respecification is to be understood) in action efficiently cause the bodily movements that constitute (either in virtue of being so caused or together with that cause) the intentional action. (Think, for instance, of John Searle’s theory of action (1983: ch. 3), according to which an intentional action is the causally-complex combination of an intention in action and the bodily movement that is its effect.) But Ryle’s objection—“I am doing one thing not two”—would just as well apply to a conception of intentional action according to which intention (in action) “guides” action, but does not constitute it.

end (perhaps in concert with other means), so if one is doing *A* in order to do *B*, but doing *A* (like *this*) isn't working, one must modify what one is doing in order for what one is doing to be a means to doing *B*.

2.4 Calculative holism

With **condition (ii)**, Anscombe reminds us that it is possible for one to be doing *B* by doing *A*, and doing *A* in order to do *B*, and yet not to be doing *B* intentionally—if one's doing *B* is an accidental consequence of one's doing *A*. (Suppose, for example, that to operate the pump, which unbeknownst to our man does not work, he has to stand on a platform, and that in doing so he has activated the automatic pump, of which he is not aware, and is thereby replenishing the water supply for as long as he stands on the platform.) The point here is that to get intentional action in view *we must consider it in its calculative context*: that is, we must consider the chain *D-by-C-by-B-by-A* (equivalently, *A-because-B-because-C-because-D*) as a unit, one whose complexity is not to be understood by conceiving of it as composed from independently intelligible atoms *A*, *B*, *C*, *D*, and the 'by'/'because'-relation.

On the one hand, *A*, *B*, and *C* show the texture in practical reality, so to speak, of *D*—the particular way in which its possibility is being realized, the particular shape that doing *D* is taking on this occasion. The shape that poisoning the inhabitants of the house is taking on this occasion is *via the water supply*; but if the cook weren't eagle-eyed and incorruptible, it might have been *via the food*. Only with an understanding of the practical texture that a deed takes can we make certain judgments about it, for example whether it was done quickly or slowly, as in an example of Davidson's: "Susan says, 'I crossed the Channel in fifteen hours.' 'Good grief, that was slow.' ... Now Susan adds, 'But I swam.' 'Good grief, that was fast.' We do not withdraw the claim that it was a slow crossing; this is consistent with it being a fast swimming" (1967: 106).

On the other hand, *D* gives the *point* or *use* or *good* of *A*, *B*, and *C*—it sets the standard for their success, for what counts as doing or having done them *well*, *badly*, or even *at all* (that is, for what counts as having *finished doing them*). The standard of success for replenishing the house water-supply depends on whether or not one is replenishing it for the sake of poisoning the inhabitants. Suppose our pumper replenishes the supply with unpoisoned water. If he is replenishing it in order to poison the inhabitants, this constitutes a *mistake*; if not, it doesn't. And if he is replenishing the water supply because he's poisoning the inhabitants, he needs to pump *enough* poison to kill them (supposing *that* to be his intention in poisoning them) and also e.g. pump enough water to avoid detection (at least before the poison does its work). Otherwise the replenishing, and all the pumping and arm-moving that went into it, will have been in vain.²⁸

These considerations suggest that, just as intentional action is structurally complex, so must be its representation. This has an obvious implication for the cognition condition. If an agent is doing *B* intentionally, he is intentionally doing-*B*-by-doing-*A* and doing-*B*-because-he-is-doing-*C*. His doing *B* intentionally cannot be understood apart from its calculative context, and this goes just as much for the *agent's* representation of

²⁸Note that this does not imply that it is not possible to do something intentionally for no reason. According to the conception of the cognition condition I am developing, if an agent is doing *A* for no reason, he knows why he is doing *A*: for no reason. Thus there are no further ends that can set normative or temporal constraints on his doing *A*. If he can be said—and I think in the normal cases of acting for no reason (doodling, tapping one's fingers, whiling away the time, etc.) it would be a stretch—to be doing *A* well or badly, or to be half-way through doing *A*, then such judgments would have to depend on considerations intrinsic to doing *A*. But this, if possible at all, would of necessity be the outlying case.

it enjoined by the cognition condition: if he is doing *B* intentionally, then he knows that he is doing *B*, and he knows that he is doing *B* by doing *A* and doing *B* because he is doing *C*. In short, knowing what you are doing involves knowing how and why you are doing it.²⁹

Knowing why and how are *internal* to knowing what one is doing, not something *external* that might be known “in addition”. Only when situated in its calculative context can an intentional action can be assessed as a success or failure in various ways, and—crucially—the agent needs to be in a position to make these evaluative judgments about his ongoing performance so that he can shape and re-shape it in the ways necessary to achieve his end: the pumper needs to know when to stop pumping, but he cannot know whether he’s pumped too little or gone on too long unless he knows that he’s pumping because he’s replenishing the water supply. It can only be in a useless and etiolated sense that anyone, and in particular the agent, can “know what the action is” independently of knowing why it was, or is being, done.

2.5 The cognition condition, reformulated

The calculative structure of intentional action is a case of *self-conscious teleology*, differing from the teleological structure found in the behaviour of plants, animals, and machines in that it is constituted by the agent’s representation of it. This is why, while no one is going to suggest that the teleological structure of the behaviour of plants, animals, and machines implies a cognition condition, we have reason to think things will be different in the case of intentional action. We understand the functional relations that hold between the different processes that constitute the workings of a machine, the distinction between function and malfunction, and the distinction between a part of the machine (including the improper part) “fulfilling” its function accidentally or non-accidentally, all in terms of the machine’s *design*. An explanation of the teleology of living things will not appeal to an external agency of design, but rather to the nature of the kind of creature under consideration—its *life-form*. In the case of intentional action, the teleology internal to the structure of action is explained by the agent’s *practical thought*. That is, the reason why doing *D* sets the standard on this occasion for the agent’s doing *A*, *B*, and *C* is that the agent himself represents doing *D* as his end, and doing *A*, *B*, and *C* as his means. My heart is beating because it is circulating blood around my body, and thus circulating blood by beating, and it does this whether I think the connection between beating and circulating or not. The form of explanation is biological, and as such it has nothing to do with thought. If I am thinking that my heart is beating in order to circulate blood around my body while it is in fact doing so, this is simply a coincidence.³⁰ But it is not like this in intentional action: the representation of the teleology is constitutive of it.³¹

According to the account developed via the foregoing reflections on the A—D order, the source of the cognition condition on intention action lies in the structure of intentional action itself. But this structure

²⁹It is not uncommon for philosophers inspired by Anscombe to recognize that knowing *what* one is doing involves knowing *why* one is doing it, but that knowing what one is doing equally involves knowing *how* one is doing it typically goes unthematized. See, e.g. Velleman (1989: chs. 1–2), Setiya 2007: 40, Setiya 2010: 89, and Moran and Stone (2009: mspp. 17, 25–6)[CITE].

³⁰Obviously it is no coincidence in the following sense: if my heart weren’t beating, I wouldn’t be thinking. But if my heart is beating, then that it is beating in order to circulate my blood has nothing to do with any representation I might form of this connection. An important point is this: part of the explanation of why *my* heart is circulating blood around *my* body by pumping is that *the* heart circulates blood around *the* body by pumping. But it is no part of the explanation of why *this* pumper is poisoning the inhabitants of *this* house that “the” pumper “poisons” the inhabitants of “the” house.

³¹Cf. Rödl 2007: 47–8.

has been revealed to be the structure of practical reasoning: of the rational specification and dynamic re-specification of means to ends. An agent engaged in a stretch of calculatively-articulated intentional activity knows what he is up to because he understands how and why the elements of his action, which are themselves intentional actions, combine such as to amount to him intentionally producing the whole thing. His understanding is practical: it is arrived at through practical reasoning, and abides in the calculative articulation of his intention, throughout the inevitable respecifications that he makes to it in the course of its execution. This is what it means, I think, to speak of agent's knowledge as *knowledge in intention*, not a knowledge that somehow precedes or accompanies intention or intentional action while inhering in a belief merely contingently connected with intention or action. Thus, my investigation into Anscombe's account of the calculative structure of intentional action has revealed that this structure is the source of the cognition condition, of which a more detailed and robust articulation is now available: if you are acting intentionally, you have knowledge in intention of:

- (i) *what* you are doing (e.g. replenishing the house water-supply),
- (ii) *why* you are doing it (e.g. because you're poisoning the inhabitants), and
- (iii) *how* you are doing it (e.g. by operating this pump).

This formulation strongly suggests that the account just given of the structure of intentional action is, strictly speaking, incomplete. If the answer to the question 'Why?', asked of an intentional action cites just another intentional action, and this can keep on going by extending the A—D order as far as we want, then it seems that we don't understand *any* of the answers. Likewise, if the answer to the question 'How?' asked of an intentional action cites just another intentional action, and this can keep going on *ad infinitum* by extending the A—D order in the other direction, then it seems we don't understand how anyone can act intentionally at all. The first worry concerns whether there are *non-calculative reasons for acting* and what they're like if there are any. These questions lie beyond the scope of this essay.³² The second worry concerns whether there are *basic actions* and what they're like if there are any, questions I address in §§7–9 below.

3 From intentional action to intention

3.1 Intention, aspiration, and knowledge how

Anscombe claims that if an agent is doing something intentionally, he has knowledge in intention that he is doing it. In §1.1, I suggested that, given what she says about the unity of the three headings under which the concept of intention appears—intentional action, intention with which, and expression of intention for the future—we should think of the cognition condition as applying to intention with which and the expression of intention for the future as well. My basic strategy for “extending” the account given of the cognition condition

³²It is enough to point out that any realistic account of acting for non-calculative reasons, whether they be lofty or mundane, will presuppose the account of the calculative inner structure of intentional action that I have sketched—after all, sometimes even acting from the motive of duty might involve crossing the street, and thus looking both ways, taking one step and then another, in order to cross the street, in order to perform one's sublime act. Candace Vogler gives a more detailed argument, along these lines, for the claim that acting for calculative reasons is the fundamental form of intentional action in ch. 6 of her 2002.

in the last section from intentional action to intention for the future is straightforward, if contentious (a direct defence of it cannot, for reasons of space, be given here): intention for the future is intentional action *in prospect*, and thus falls under the same account. We simply shift from considering the calculative structure of intentional action *in progress* to that of intentional action *in prospect*, where we find the applicability the same trifecta of questions ('What?', 'Why?', and 'How?'), and for the same reasons.³³

That is, if someone announces an intention to do something next Tuesday, we get a better idea of *what* he intends to do by asking him *how* he intends to do it and *why*. In finding out the answers to these questions, not only do we have a better idea of what he intends to do, we also become convinced *that* he has the announced intention. Here the question 'How?' is crucial, as Annette Baier points out:

whenever we announce an intention in a vague way, we must be prepared to back it up with a more precise specification or a demonstration of how we will do what we intend doing, if our hearers are to be satisfied that we can do it, and so be satisfied that what we have is an intention and not just an aspiration. ... Answers given to the question "How?" or "What's your method?" may themselves invite the same "how" question.... If the question can continue to come up, then it will remain a mystery how the man will do what he says he will do. (1970: 653–4)

Earlier on (§2.4), we imagined a scenario in which the pumper was doing what he intended (replenishing the water supply), but not in the way he intended (operating the pump with his arm); in fact, he was replenishing the water supply by standing on the platform next to the old hand pump, which activated an automatic pump. Because he didn't know *how* he was doing what he was doing, he wasn't doing it intentionally. That is, the connection between his means and end was accidental; it was not contained in his intention, and thus not something known to him *in intention*.³⁴ He was doing what he intended, but not intentionally; he was realizing his goal not through practical reasoning and knowledge, but simply through luck. Now, Baier's remarks suggest the following parallel. Just as someone who is bringing off his end *D* without knowing how he is doing it is not doing *D* intentionally, so someone who "intends" to do *D* without knowing how he is going to do it does not really *intend* to do *D*: he merely *aspires* to do *D* (cf. Anscombe 2005: PAGE).

Now, aspiring to do something differs from merely wishing for some state of affairs or event to obtain or transpire. As Anscombe notes, "[a] chief mark of an idle wish is that a man does nothing—whether he could or no—towards the fulfilment of the wish" (1963: §36). Even though a mere or idle wish represents its object as in some sense good, it does not amount to practical thought because it has no inner tendency towards action, which is what practical thought is ultimately for.³⁵ By contrast, just as an intention has the inner tendency to further articulate and re-articulate itself in action, so an aspiration has the inner tendency to

³³By focusing, (in the last section) on intentional action and (in this section) on intention for the future and its expression, it might seem as if I have ignored Anscombe's third heading. However, in examining the structure of intentional action, we have seen that it *contains* intention with which: if someone is doing *A* intentionally because he is doing *B* intentionally, he is doing *A* with the intention of doing *B*, and he knows that he is doing *B*. If, on the other hand, he is doing *A* with the intention of doing *B*, but he is not *yet* doing *B*, then he is doing *A* because he intends to do *B*, and an expression of his knowledge in intention will at the same time be an expression of intention for the future. This, then, is why I have not treated the idea that there is a cognition condition on intention with which separately.

³⁴Of course, he could come to know about it by observation.

³⁵Cf. Anscombe: "The primitive sign of wanting is *trying to get*" (1963: §36). The kind of wanting of which Anscombe speaks—which she distinguishes from idle wish, the "prick of desire at the thought or sight of an object, even though a man then does nothing towards getting the object" (67), hope, and "the 'I want' of a child who screams for something (76)—is, as I understand it, the most primitive kind of practical thought.

articulate itself such that it becomes an intention. This articulation too is the work of practical reasoning: an agent's aspiration can become an intention, by his coming up with a plan—by reasoning from the end which is at a distance, to an immediate means that he can take. This inner tendency towards such articulation—an inner tendency towards perfecting itself in action and practical knowledge—is what marks practical thought out as rationally efficacious; that is, *as practical*. It is in the nature of the thought 'I want (aspire, intend) to do *D*' that it ends up articulated 'I intend to do *D* by doing *C* by doing *B* by doing *A*' and realizes that articulation in action and practical knowledge: 'I *am doing D* by doing *C* by doing *B* by doing *A*'. Such inner efficacy is absent in mere wish.³⁶

To illustrate, consider the following attitudes towards becoming rich. Someone who *idly wishes* he were to become rich does nothing towards getting rich (save perhaps purchasing the occasional lottery ticket, though this is by no means necessary); his wish is primarily manifested in fantasy—in imagining that he wins the lottery, or that he has a rich aunt who makes him the heir to her fortune, in imagining where he would live, what he would buy, and so on. Someone who *aspires* to become rich would be getting rich if he knew how, but he does not; in contrast to the idle wisher, his aspiration is manifested in deliberation and action, not merely in imagination—not in getting rich, however, but in figuring out how to get rich (he spends his time at the library, reading up on real estate, the stock market, spread betting, or bank robbery).³⁷ Someone who *intends* to become rich, however, knows how to get rich: he has worked out a plan, which he is putting, or going to put, into action. And if he is getting rich by executing his plan, he is getting rich intentionally.

The distinction between intention and aspiration is easy to overstate, however. Surely it is possible for you to intend to go shopping next weekend, despite not having decided where exactly you'll go or whether you'll drive or take the bus; it would be decidedly off-key to lump you together with the man who wants to get rich but has no idea how to do so, or to insist that you provide us with "a more precise specification or a demonstration of how you will do what you intend doing" on pain of refusing to credit you with an intention. We do not want the implication that an intention is somehow fully determinate (fully calculatively-articulated), and that anything less than that must relegate it to the status of a mere aspiration. The idea that an intention for the future could be fully determinate is a philosophical fantasy; it would preclude what is always possible and usually necessary, the rational re-specification of one's intention in action. But then doesn't this suggest that while someone who intends to do something knows what he intends to and why he intends to do it, he need not know how he intends to do it, and thus that the intention/aspiration distinction is spurious?

I think not. An intention need not be fully articulated into subordinate intended means, but, as practical, it is such as to so articulate itself through practical reasoning. It may be that the question of the means of transportation to the shops need not be settled until the morning of departure. But one who intends to go

³⁶Stephen Engstrom (2009: 68–88) makes some similar remarks, in a Kantian framework, about the way in which deliberation serves to transform "wish"—a "problematic" practical judgment specifying what the agent *would* do that, though merely "practical...*in potentia*" (68), is nevertheless by that very feature to be distinguished from a mere or idle wish—into "choice"—an "assertoric" practical judgment specifying what the agent *will* do by specifying how (by what means) the object of wish is to be realized. According to Engstrom, "the inner act of deliberation is the characteristic form of the efficacy proper to wish itself, an efficacy that can be described as practicality's own self-actualization, and even (given practicality's self-awareness) as its self-conscious self-actualization" (87).

³⁷Such a man might be walking to the library with the intention of reading a book about spread betting with the intention of figuring out how to get rich with the aspiration of getting rich. He would not thereby be getting rich intentionally.

shopping next weekend knows that the means of transportation needs to be decided upon at some point. There is a significant difference between leaving it unsettled which of several options one might take, one shall take, and not having any idea of what the options are, of having any idea how to do what one intends to do. This reveals that we need to be more precise in drawing the distinction between intention and aspiration in terms of knowing how to achieve one's end. One can intend to do *A* without knowing how one *shall* do *A* so long as one knows how one *can* do *A*. But if one doesn't know how one *can* do *A*, then one can merely aspire to do *A* (where the inner tendency of this aspiration is towards the acquisition of knowledge how to do *A*, and thus its own transformation from aspiration into intention). Thus, I may know how I *can* get to the shops without yet knowing how I *shall* get to the shops. This is sufficient for me to intend to go to the shops, in the knowledge that my intention, as such, must at some point become articulated in some such way as 'I intend to go to the shops by taking the bus' if I am to actually get there. But if I do not know how I can get to the shops, then all I can do is aspire to go shopping.

This distinction between two kinds of knowing how—*general knowledge how* to do *A*, and *particular knowledge how* one is going to do *A* on some particular occasion—is of importance, and applies to intention in action, too. If one is intentionally doing *B* by doing *A*, one knows how one is, here and now, doing *B*: by doing *A*. This particular knowledge how is practical knowledge, knowledge in intention. It is not derived from observing that the way one happens to be doing *B* is by doing *A*. Rather, it is possible only because one already has the corresponding general knowledge how, that one can do *B* by doing *A*, or, equivalently, that doing *A* is a way of doing *B*; one puts this general knowledge into practice on a particular occasion in doing *B* by doing *A*.³⁸

The distinction between intention and aspiration is not binary, all-or-nothing. It is not that either I definitely have general knowledge how to do something, and can thereby intend to do it, or I have no idea how to do it, and can thereby harbour mere aspiration. I might well reason: 'I want to do *B*; doing *A* might work; so I'll do *A*'. Is my thought about doing *B* an intention or an aspiration? And if I did *A* and thereby did *B*, did I do *B* intentionally? Well, there need be no easy answer; but (as Dr Johnson reputedly said) the fact of twilight does not mean that we cannot tell between night and day. It is not always easy to tell whether an intended action was performed intentionally or brought about merely through luck, or where exactly it fell between those two poles, but it is easy enough in many cases. This suggests that the inner tendency to calculative articulation that belongs to practical thought as such has two aspects: on the one hand, the specification of means through calculation (the development from the desire to do *D* to the intention to do *D* by doing *C* by doing *B* by doing *A*), and on the other hand, the strengthening of the calculative links (from '*maybe* doing *B* is a way to do *C*, *who knows?*' to '*doing B might* be a way to do *C*' to '*doing B is a way to do C*').³⁹

According to the picture I have developed, aspiration is a kind of practical thought that is not itself intention, but which has an inner tendency to develop itself into intention. The primary way in which this happens is through practical reasoning itself: I aspire to do *B*, but don't know how to do *B*, so I figure out that I can do *B* by doing *A*, and thereby come to intend to do *B* by doing *A*. But it may be that reasoning is not enough.

³⁸Cf. Rödl, 'Two Forms of Practical Knowledge' [CITE]; Setiya 2009b: 135. That general knowledge how may be indexed to certain agents, conditions, and circumstances does not speak against, but rather specifies, its generality.

³⁹This paragraph is based on a conversation with Sebastian Rödl.

Kieran Setiya (2008: 406f.) asks whether he can form the intention to dance the tango at his wedding, though as yet he doesn't know how to dance the tango. The answer, on this view, is that he cannot; so far dancing the tango may figure only as the object of his aspiration.⁴⁰ But not knowing how to dance the tango is not like not knowing how to get home if the subway isn't running. If the subway isn't running, you can come up with an alternative plan by apprising yourself of the bus schedules, looking up a phone number for a taxi, consulting a map to find a route you can walk, and so on; that is, you can engage in practical reasoning that takes you from your end to means that you can in fact take. By contrast, if you aspire to dance the tango at your wedding, your aspiration can become an intention only if you learn how to dance the tango, and there is more to learning how to dance than tango than engaging in practical reasoning; all you can intend, for the time being, is to learn. The difference between the two kinds of case is this: on the one hand, it is a matter of figuring out how to use one's practical abilities in concert with one another given the practical circumstances (you want to do *B*, you know how to do *A*, you just need to figure out that you can do *B* by doing *A*), on the other, it is a matter of acquiring a new practical ability.

Baier says that the "mystery" of how the agent will do what he says he will do "is dispelled once he tells us his intentions in terms of activities we understand, activities where we don't need to be told how *he* will do them, because we know how they are done" (1970: 654). As I understand her, she has something like the following in mind: when we ask 'How?' of a man who announces an intention to woo someone, we are asking him by what means, of the many that are available to him, will he woo her; he tells us he will woo her by serenading her; when we ask him how—by which of the available means—he will do this, he says he will sing *Hello* outside her window and accompany himself on the guitar. He has a plan to woo her: by serenading her with *Hello* from outside her window. Now, he needs to know how to play and sing *Hello*, which means that he needs to know how to play and sing well enough to play and sing *Hello*. But this seems to require not more *planning*, but rather practical abilities or, better, *skill*: though he might play and sing *Hello* in a different *style* from another serenading woo-er on the other side of town—an ill-conceived reggae version, perhaps, or in the style of Frank Sinatra—this will typically constitute "the way in which *he* will play and sing *Hello*" in a sense other than the calculative one in which we are interested when we ask and re-ask the question 'How?' in order to establish that an agent has a genuine intention. In the relevant sense, the way in which he is playing and singing *Hello* is the way in which *anyone* would do so: (Am/G/F/G/F) "I've been alone with you inside my mind / And in my dreams I've kissed your lips a thousand times / Sometimes I see you pass outside my door / Hello...is it me you're looking for?". We know how it is done; if we are still in doubt as to whether he

⁴⁰Sarah K. Paul claims that this "get[s] things backward." She asks: "Does not the rationality of my deciding to learn how to dance the tango depend in part on my having an intention to dance it at some future point, rather than the other way around?" [CITE on Setiya msp.12]. There is no need to suppose that one must intend to dance the tango in order to explain why one has reason to learn; wanting or aspiring to dance the tango gives one reason enough. Paul appears committed to unsatisfyingly representing an agent's practical powers as conditions external to his practical thought, circumstances to be dealt with in the course of executing one's intention and not conditions internal to having an intention. But there is a qualitative difference between, on the one hand, the relation between an intention to dance the tango and the discovery of a circumstance that necessitates re-articulating—so as to "pick up"—one's intention (the discovery, say, that one's shoes are unsuitable for dancing), and, on the other, the relation between an ostensible intention to dance the tango and the discovery that one doesn't know how to dance it. By contrast, Setiya thinks that "it is possible, if inadvisable, to form the intention of doing something I do not know how to do" (2009b: 136). It would be inadvisable, he thinks, because the belief in which the intention would consist on his cognitivist view would be unjustified. Thus he presupposes an account of intention that is normatively constrained, but constitutively unconstrained, by knowing how to do what one intends to do; such an account cannot mark the distinction between intention and aspiration.

has a genuine intention, it is because we doubt that he has the necessary skills.⁴¹

I originally quoted Baier with a view to developing this parallel: just as someone who is bringing off his end *D* without knowing how he is doing it is not doing *D* intentionally, so someone who purportedly intends to do *D* without knowing how he is going to do it does not really *intend* to do *D*; rather, he merely *aspires* to do *D*.⁴² Our discussion has primarily focused on drawing out the way in which intentions, like intentional actions, are themselves calculatively structured (though perhaps so only potentially).⁴³ The general knowledge how to do things implicated by the calculative structure of practical reasoning is knowledge of means, where this is understood on the model of deliberation: of figuring out what would be sufficient to achieve one's end, and selecting one such (set of) sufficient means. Such deliberation is often thought about in terms of planning⁴⁴—though as we saw when thinking about the rational re-specification of an intention in action in §2.3, the deliberative process need not be a “mental” process distinct, or even distinguishable, from the progress of the intentional action itself. But we have now uncovered another way in which an agent might not know how to do what he aspires to do, namely by failing to possess the skill or skills his plan requires of him. This suggests a distinction within general knowledge how, between two kinds of knowledge how to do something—*procedural knowledge how* and *skill*—about which I will have more to say later.

Assuming that what one intends to do is neither a basic action nor something one is doing for a non-instrumental reason, to intend to do something is to intend to do it for the sake of doing something else, and (at least by when the time to execute it comes) to intend to do it by means of doing something else. Intention, whether it is in action or for the future, is, essentially, calculatively articulated. This point is not refuted by the possibility of intentions to do things for which the means have not yet been decided upon, because insofar as such a thought is an intention, the efficacy that belongs to it as practical thought is realized (in part) by the figuring out and selection of such means—by deriving particular from general knowledge how—when or before the time to take them comes. As we have already seen, calculative articulation is self-conscious teleological articulation: the means–end structure of intention (whether in action or for the future) obtains in virtue of the agent's knowledge of it. Thus, if you intend to do something, you know what you intend to do, why you intend to do it, and how to do it, and you either know how you intend to do it or that you will know how you intend to do it at some point before doing it. The cognition condition applies to intention for the future as well as intentional action. This falls out of the fact that intention for the future is intentional action *in prospect*: what it is that one now intends to do is something whose subsequent realization will take the calculatively-articulated form we have already discussed, where this form will obtain precisely because the agent represents it as obtaining; one's present representation of one's action in prospect must also take this calculatively-articulated form if it is to be the representation that is the source or cause or principle of the

⁴¹ We might also doubt whether he has, or can get, the “objective” counterparts to the necessary skills, i.e. their instruments, such as, in this case, a guitar. Unfortunately, this fascinating topic lies beyond the scope of this essay.

⁴² This formulation may suggest another parallel: if someone is bringing off his “aspired” end *D* without knowing how he is doing it, then he is doing *D* not intentionally but “aspirationally”. But this is not right, as I show in §4 below.

⁴³ When I say that an intention to do something may be calculatively structured only potentially I mean more than that it is (merely) possible that it ends up becoming so articulated; I mean that it is in its nature to so articulate itself, and that it will be no accident if it does and some accident if it doesn't.

⁴⁴ The popularity of this helpful but ultimately one-sided way of thinking about practical thought is primarily owed to Michael Bratman (1987).

action, when its time comes.⁴⁵

3.2 Thinking that you'll try and thinking that you'll succeed

This is enough, I think, to see that Bratman's bookstore case from §1.2 can be dispensed with. According to Bratman (1987: 37), he intends to go to the bookstore, but he's agnostic with respect to the question whether he will even try to go there when the time comes, because he knows he's so forgetful; he concludes that an intention to do *A* doesn't "involve" even the belief that one will *try* to do *A*, let alone the belief or knowledge that one is actually going to *do A*. But I think that we can now see that Bratman's "intention" to go to the bookstore reveals itself to be simply a mere aspiration or wish: if he's considered his bad memory (and this, after all, is the ground of his agnosticism), then it's a pretty hopeless plan that doesn't involve writing a reminder on his hand or leaving a post-it on the dashboard or something. This is in part an *ad hominem* objection, because for Bratman intentions are just plans writ small. But it is not just *ad hominem*: if someone avows an intention to do something, we can investigate whether he really has an intention and not a mere aspiration or wish by asking him how he'll do what he "intends".⁴⁶

The kind of case that should concern us, then, is one where someone putatively intends to do something, despite being agnostic as to whether or not he will *succeed*, not whether or not he will *try*. Such a case would be directly in line with Davidson's original carbon-copier case. And indeed, someone who is impressed by carbon-copier counterexamples might feel aggrieved by the discussion so far, on the grounds that it seems tangential to the real issue: *belief in*—or even worse, *foreknowledge of*—*success*. It may be all very well to suppose that if you are doing something intentionally, you know why and how you're doing it, but this isn't the same as knowing that you will end up having done it. Likewise, it may be all very well to suppose that if you intend to do something, you know why and how you are going to do it, but the real question is whether or not you *will* do it. Knowledge of success is the issue; it is scepticism about eventual success that props up the agnosticism on which carbon-copier style worries trade. We must turn, then, to the question of the relationship between *doing* something (or intending to do something), *being* such that one will end up having done it, and *knowing* that one will end up having done it. This question does not concern in the first instance the calculative structure of intention and intentional action, but rather their temporal structure. Again, I begin with action, by examining the connection between two thoughts that, I will claim, belong to the knowledge in intention of an agent who is doing *B* intentionally: 'I am doing *B*' and 'I am going to do *B* (i.e. end up having done *B*)'.

4 The temporal structure of intentional action

4.1 Stopping and finishing

Here, a comparison with Davidson can take place rapidly, because—somewhat surprisingly—his account of the logical form of action sentences (1967) is silent about the truth-conditions of the kinds of sentence that

⁴⁵This section owes much to conversations with Sebastian Rödl and Candace Vogler.

⁴⁶Of course, most people can and do simply intend to go to the bookstore; the point is that the context Bratman describes is one in which the status of his "intention" is in question, legitimately subject to the question 'How?'.

comprise our topic, i.e. ‘S is doing B’ (or more precisely ‘I am doing B’). The object of his inquiry is *completed* actions, and thus sentences of the form ‘S did B’, which, as is well known, he analyses by quantifying over events along these lines: $(\exists x)(\text{did } B(S, x))$.⁴⁷ Anscombe provides the clue for thinking that this amounts to an oversight: “A man can *be doing* something which he nevertheless does not *do*, if it is some process or enterprise which it takes time to complete and of which therefore, if it is cut short at any time, we may say that he *was doing* it, but *did not do* it” (1963: §23). What would we quantify over in such a case?⁴⁸

Picking up on this clue, one strand of recent work in action theory has made much of the temporal structure of action. Rather than theorizing only completed actions, philosophers have increasingly been drawn to examine the temporal structure of actions in progress.⁴⁹ The crucial element that this brings into the discussion is marked in language by what linguists call *aspect*. In applying a state-predicate *F* to an object *a*, the form of predication admits contrasting tenses (*a was/is/will be F*). By contrast, in applying an event- or process-predicate *do B* to an agent (or patient) *S*, the form of predication admits contrasts of both *tense* (past, present, future) and *aspect* (imperfective, perfective); the aspectual contrast that interests us is between one type of imperfective, the progressive (*S was/is/will be doing B*), and the perfective (*S did/will do B*).⁵⁰

Davidson’s investigation into the logical form of action sentences characterized by perfective aspect conceals the fact that actions (and events more generally) are temporally structured. As Bernard Comrie puts it, perfective meaning “presents the totality of the situation referred to...without reference to its internal temporal constituency: the whole of the situation is presented as a single unanalysable whole, with beginning, middle, and end rolled into one” (1976: 3); but the whole of an event that takes time to happen cannot be located at a single present moment. Perfective meaning does not imply that the event or process is temporally structureless (cf. Comrie 1976: 21); it simply abstracts from its internal temporal constituency. By contrast, imperfective meaning, which is expressed in progressive judgments, “is crucially concerned with the internal structure of the situation, since it can both look backwards towards the start of the situation, and look forward to the end of the situation” (Comrie 1976: 4). By reflecting on the temporal structure of action, revealed by the aspectual contrast passed over by the standard Davidsonian approach, I will argue that there must be a cognition condition on intentional action (in progress and in prospect) that encompasses knowledge of eventual success, properly understood.

I already noted (in §?? above) one important feature of progressive judgments, their *broadness*: a progressive judgment ‘S is doing B’ may be true even if there is nothing that S is doing right now that amounts to doing B or taking some means to doing B. In this section, I focus on the other important feature of progressive judgments: their *openness*. The openness of the progressive is precisely what Anscombe points to when she

⁴⁷E.g. ‘Shem kicked Shaun’ becomes $(\exists x)(\text{Kicked}(\text{Shem}, \text{Shaun}, x))$; ‘I flew my spaceship to the Morning Star’ becomes $(\exists x)(\text{Flew}(\text{I}, \text{my spaceship}, x) \ \& \ \text{Motion-towards-and-terminating-at}(\text{the Morning Star}, x))$. See Davidson 1967.

⁴⁸Perhaps Davidson finds it so easy to countenance the otherwise problematic inclusion of events in his ontology (see e.g. his 1969) because he treats them (as he treats actions) as without internal structure.

⁴⁹The most influential work on this topic is Michael Thompson’s ‘Naive Action Theory’ (2008: Part 2). See also Falvey 2000; Rödl 2002, 2007: ch. 2; Moran and Stone 2009.

⁵⁰There is no present perfective, at least for events and processes that take time to occur: ‘John walks to school’ does not say of the present what ‘John walked to school’ and ‘John will walk to school’ (i.e. ‘It will be the case that John walked to school’) say of the past and the future; rather, ‘John walks to school’ says that John habitually walks to school. (By contrast, ‘John is walking to school’ says of the present what ‘John was walking to school’ and ‘John will be walking to school’ say of the past and the future.) In general, ‘S does B’ says that S does B habitually; there are other readings too (describing non-durative events; the historical present; certain performative uses), none of which concern us here.

notes that “a man can *be doing* something which he nevertheless does not *do*.” That is, while the truth (at t_1) of ‘S is doing B’ entails (at t_2) the truth of ‘S was doing B’, neither entails—at any time—the truth of ‘S did B’.⁵¹ Thus it could be true that S is (or was) doing B, even though S will never have completed this act of B-ing. The intelligibility of the thought that S was doing B, but didn’t do B (and isn’t still *en train*) depends on the event-predicate *do B* admitting a contrast between two ways in which something that is doing B can cease to do B: by *stopping* doing B and by *finishing* doing B. The predicate *cross-the-street* admits this contrast: one can have been crossing the street even if one didn’t and never will make it to the other side, if one’s street-crossing was cut short, perhaps by an oncoming bus. The predicate *walk* does not: there is no contrast between stopping walking and finishing walking. Following Antony Galton (1984: 67), I’ll call predicates that admit the stopping/finishing contrast *telic*, and those that do not *atelic*.⁵² To put the point slightly differently, telic predicates are those whose instances can be *interrupted*. An atelic predicate can be transformed into a telic predicate by specifying an end, either in terms of an objective to be achieved (e.g. from *walk* to *walk-to-school*) or in terms of a length of time (e.g. from *walk* to *walk-for-an-hour*). By specifying an end, one specifies the proper terminus of the event-form to which the predicate refers, whose instantiation on an occasion may then be subject to improper termination.

All of this applies regardless of whether the event in progress is an intentional action or not: it marks what belongs to the *representation of movement*, a topic preliminary to ours, the representation of intentional action (rational self-movement), which is a specific form of it.

4.2 The quick argument

Our present guiding question is: what is the relationship between doing A intentionally and ending up having done A intentionally? Clearly we are interested in the question as it pertains to telic concepts: our question pertains to the successful achievement of an end. We want to know whether knowing that one is doing A involves knowing that one will end up having done A—whether knowing what you’re doing involves knowing (or believing) you’ll succeed. There is a quick argument, based on the aspectual considerations already outlined, for the claim that it doesn’t. If John knows that he is walking home, then it is true that he is walking home. But we have seen that the proposition ‘John is walking home’ is inferentially unrelated, so to speak, to the subsequent truth of ‘John walked home’ (and its negation), and thus to the present truth of ‘John will have walked home’ (and its negation). So what John knows, when he knows he is walking home, does not involve knowing that he will succeed in walking home: his knowledge in intention does not amount to cognition of success. And given that I argued in §3 in favour of conceiving of intention for the future as intentional action in prospect, it would appear that if John intends to walk home, then his knowledge in intention that walking home is what he is going to do is not in fact cognition of success; that is, his “knowledge” that he is going to walk home turns out not to be knowledge that he will, in fact, walk home. The quick argument, combined

⁵¹ Thus it is a mistake to assume that an event- or process-predicate denotes a *type*. For an incomplete walking across the street falls under the concept *walk-across-the-street* even though there is no completed event—the token that would instantiate the type. See Thompson 2008: 134–8, especially n.21.

⁵² Galton is himself following Comrie 1976: 44ff. Any temptation one might feel to suppose that the distinction between telic and atelic predicates, as flat-footedly characterized here, maps straightforwardly onto Aristotle’s more complicated distinction between *kinesis* and *energeia*, of which it is undoubtedly reminiscent, should be resisted. For a recent and pioneering study of the latter distinction, see Burnyeat [[CITE]].

with the assimilation of intention for the future to intentional action in prospect, apparently shows that the cognition condition cannot extend to cognition of success.

But the quick argument is no good. It slides from the fact that neither ‘S is doing B’ nor ‘S was doing B’ entails ‘S did B’, ‘S did not do B’, ‘S will do B’, or ‘S will not do B’ to the claim that ‘S is doing B’ and ‘S was doing B’ are indifferent to the eventual doing of B. But this is not so; success and failure are not on a par: if they were, there would be no grounds for ever thinking that a sentence of the form ‘S is doing B’ is true.⁵³

We can see this by looking at a case in which something was doing something, even though it didn’t end up having done it: Unlucky John, who was walking home from work, even though he didn’t end up having walked home, because he was run down by a car. To walk home from work, John needed to walk down Avenue A, take a right at Boulevard B, and then turn left into C Close. Suppose John has walked down A to B, and is walking along B towards C: he is walking, but has not yet walked home; all the while he was walking down A, and during his progress so far along B, he is walking, but has not walked, home. Then, while crossing B in preparation to turn left into C, he is hit by the car; he doesn’t embark on walking down C, and hence doesn’t end up having walked home. Surely it is true, though, that he was run over *while walking home*. But if eventual success and failure were on a par with one another, then surely all we would be entitled to say is that what he was doing was walking down A and then along B. And in fact, it would be more appropriate to say, not that John *was walking* down A to B and then along B to C, but only that he *walked* down A and then along B *to the fateful spot*. That is, we are forced into scepticism about the very idea of an event’s being in progress. According to such a scepticism, events are things that have happened; there is no such thing as them happening. But such a position is incoherent, for the very idea of something’s having happened presupposes the idea that there was a time at which it was happening, and the idea of something’s happening is what is presently denied.⁵⁴

The argument from the premise that someone can be doing something that they nevertheless will not do to the conclusion that intentional action does not involve knowledge of success displays too little sensitivity to the aspectual considerations upon which it relies. Consequently, it proves too much: if the possibilities of eventual success and failure were on a par with respect to the truth of ‘S is doing B’, then the concept *event* would be annihilated.⁵⁵

If the fact that something is doing something neither entails successful completion nor is indifferent to the possibilities of eventual success and failure, then the general conclusion we should reach is this: if S is doing B, then it will be no accident if S succeeds in doing B, and it will be some accident if S fails.⁵⁶ This is no surprise: our understanding of the telic concept *do B* has two possibilities built in for ways in which something that is instantiating that concept can cease to instantiate it—either by finishing doing B or by stopping short of finishing. Stopping and finishing are not on a par; finishing doing B is proper to a process of doing B, stopping short improper.

⁵³Here I follow Rödl 2007: 29ff.

⁵⁴I am not, absurdly, supposing that the idea of something’s happening is intelligible independently of the idea of something’s having happened; the point is that the two ideas are interdependent, and equally implicated in the concept *event*.

⁵⁵Again, nothing turns on the fact that the event in my example is an intentional action (or rather, it would have been, had John made it home). I might just as well have illustrated the point with a tree falling over or a boulder rolling down a hill.

⁵⁶“Success” and “failure” here are to be understood in terms of the end specified by the telic concept; thus, one might succeed in dying or falling over.

4.3 The epistemological argument

We said that for *S* to be doing *B*, it must be no accident that *S* ends up doing *B*, and that to deny this, as the quick argument implicitly recommends, would be to annihilate the concept *event*. But now a second sceptical possibility rears its head (epistemological rather than metaphysical): how can we *know* that *S* is doing *B*? And how could *S* know that he is doing *B*? Our sceptic grants that if *S* did *B*, there was a time when *S* was doing *B*; and at that time, it would, as things turn out, have been *true* to say ‘*S* is doing *B*’. But what could have *justified* that assertion *then*? It looks as if all we could ever be justified in saying while an event is in progress is that it *seems* to be in progress. And the same looks like it goes for the agent of an intentional action: if asked what you’re doing, perhaps all you are justified in saying is something like, “As far as I know, I am doing *B*”.⁵⁷ While we can be justified, by the eventual occurrence of the event, in saying, later, that it *was* (then) in progress, and thus, though it turns out that it can be true that an event *is* in progress, such truths escape our knowledge until after the event’s resumé.

The objection is not that, if we are ever to claim that *S* is doing *B*, we must know that *S* will do *B*, for we have established that the progressive judgment does not entail the corresponding perfective judgment. Even if we could know, on an occasion, that *S* is doing *B* on the basis that we know that *S* will do *B* (and in any case it is not clear how we could rule out the myriad possibilities that would prevent *S* from doing *B*), this cannot be the only criterion for knowing that *S* is doing *B*, or else we could never know that *S* was doing *B* but didn’t end up having done *B*. Thus the sceptic does not claim that we retreat from ‘*S* is doing *B*’ to ‘*S* seems to be doing *B*’ on the grounds that we have not ruled out the possibility that *S* might stop doing *B* before finishing doing *B*. Rather, the worry is that, faced with a situation in which *S* is putatively doing *B*, we have no criterion by which to say whether *S* is doing *B* or *B'* or *B''* or...until *S* has done *B*, and if *S* doesn’t end up having done *B*, no criterion by which to say what, if anything, *S* was doing at all.⁵⁸

We can bring out the worry by comparing the case of Unlucky John with this case presented by Kevin Falvey:

A friend approaches me as I’m at the track one afternoon preparing to run. “What are you doing?” he asks, as I take off down the track. “I’m running a four-minute mile,” I call back, only to stop short coming off the first turn, exhausted, and plod slowly back to the starting line. “What happened?” my friend asks. “Well,” I say, “I didn’t run a four-minute mile—in fact I guess I *can’t* run a four-minute mile. But I was *running* a four-minute mile; I just didn’t complete it. You know, just as someone can be crossing the street but not cross the street, I was running a four-minute mile, even though I didn’t do it.” (2000: 23)

The point, I take it, is that Falvey didn’t end up having run a four-minute mile, he was never running a four-minute mile, and he was unjustified in claiming that he was. But if Falvey couldn’t transform a few strides down the track into a case of running a four-minute mile by making his claim, what was it that made John’s stretch of walking a case of walking home and entitled him to claim that that was what he was doing? If we are ever to be justified in applying the present progressive to a situation—if, that is, we are ever to be justified in claiming that something is happening—then we must see our way to legitimately distinguishing between Falvey and Unlucky John, even though neither of them ended up succeeding in their tasks. We need

⁵⁷Cf. Galton 1984: 86–7.

⁵⁸The point may be illustrated with a famous example: what makes it the case that someone is following the rule ‘plus 2’ when he is counting ‘2, 4, 6, 8...’ as opposed to following the rule ‘quus 2’? And how might this be known?

an account of why John *was* walking home and how he knew that he was, while Falvey *wasn't* running a four-minute mile, and thus that his claim that he was running one didn't amount to knowledge. But how can we explain this?

The first thing to notice is that our question—how to distinguish between cases where 'S is doing B' is true and cases where it is false—arises regardless of whether the event in progress we consider is an intentional action or not. As Anscombe notes, of the openness of the progressive: "The point...is in no way peculiar to intentional action; for we can say that something was falling over but did not fall (since something stopped it)," and she concludes from this that "we do not appeal to the presence of intention to justify the description 'He is Y-ing'" (1963: §23). Her remark, though, also serves to develop a result we arrived at earlier, that if S is doing B, then it will be no accident if S succeeds in doing B, and it will be some accident if S fails: in Anscombe's example, something was falling over but did not fall over *because something else stopped it*. Something external, accidental, to the process must enter into the explanation of how it could be that S was doing B, is no longer doing B, but didn't end up having done B; nothing external to the process needs to be mentioned to explain how it came to be that S, which was doing B, ended up having done B.

This allows us to distinguish between Unlucky John and Falvey, as we need to if we are to rebut the sceptical challenge provided by the epistemological argument. Neither of them achieved his end (John didn't walk home, and Falvey didn't run a four-minute mile). The reason John didn't make it home is because, though he was walking home, he was involved in an accident that interrupted his progress. (You can only be interrupted doing something if you're doing it.) But Falvey wasn't interrupted while running a four-minute mile; the reason he didn't run a four-minute mile is because he was never running one. What Falvey was trying to do (if indeed he was even trying to run a four-minute mile, and not simply mouthing off) was something difficult; something that lay beyond his power. John's failure to make it home is not aptly characterized as him failing to bring off a difficult task; it is best characterized as him suffering a horrific interruption in the course of carrying out an easy task.⁵⁹

Unfortunately, even though we can distinguish the cases in this way, we haven't answered the sceptic's challenge. It is true that John was interrupted and Falvey was not, but John's being interrupted *presupposes* that he was doing something; we need to entitle ourselves to the idea that he was doing something, that what he was doing was walking home, and that this fact could have been known—indeed, that it could have been known by John himself. The possibility of interruption presupposes a process's being underway; if a process *is* underway, then, if it is not interrupted, it will reach completion. As Sebastian Rödl puts it, "[t]he end of a movement in progress is present in the movement while it has not yet reached it" (2005: ch.6 msp.7). And this is what we need to understand, if we are to understand thoughts of the form 'S is doing B', and thus the form of thought in which we are primarily interested: the first-person progressive expression of intention in action, 'I am doing B'. In other words, to justify the application of the present progressive—to legitimately judge that S is doing B—we need to know that S will do B if it is not interrupted. This is weaker than needing to know that S will in fact end up having done B, but stronger than needing to know that it is possible that

⁵⁹On the deep differences between the superficially similar statements 'I would have run a four-minute mile if I could' and 'I would have made it home if I hadn't been run down', see Austin 1956.

S will end up having done *B*.⁶⁰ How can we know this? How can what is happening here and now contain a tendency towards some particular end, such that it will be no accident if it accomplishes that end? What justifies the application of the concept *do B* to an event in progress?

4.4 Justifying the application of the present progressive

It depends on the value of *B*. If a stone is rolling down a hill, then its being no accident that it will end up having rolled down the hill is explained by mechanics: once the stone (which has a certain shape and mass) has begun rolling (at a certain speed and acceleration) down the hill (which has a certain decline), it will be no accident that it reaches the bottom—indeed, it will be no accident that it reaches the bottom at a certain time, and having followed a certain trajectory. Only if something interferes will the stone not reach the bottom of the hill. If a dog is moulting, then its being no accident that it will end up having moulted is explained by appeal to the dog's life-form: moulting is a non-accidental feature of the life of dogs as such. (In other words, *dogs moult*.) We can know that our dog was moulting (and that he was about half-way through) before he met his untimely end because we know that dogs moult and what moulting is. On the other hand, it is possible to judge wrongly that a dog is moulting; somehow or other the cause of the hair loss we mistook for moulting or its onset comes to light. Similarly, it is possible to judge wrongly that a stone is rolling down a hill; if, early in its decline, the hill plateaus, then a stone's rolling down to the plateau might be mistaken for its rolling down to the bottom of the hill. "That stone is rolling down the hill," I say. "No," you say, "it was only pushed hard enough for it to roll down to the plateau."

Physical, chemical, and biological laws specify types of physical, chemical, and biological processes, respectively, and explain why such processes come to completion if nothing interferes.⁶¹ It is often objected that the phrase "if nothing interferes" is but shorthand for an indefinitely long disjunction "unless *p* or *q* or *r* or..." the net effect of which is to render vacuous the claim about what will happen: *S* will do *B*, unless it doesn't. However, I am in sympathy here with Rödl, who writes that "[t]he formula 'if nothing interferes' does not designate an unknown content but a *known form*.... That generic statements [i.e. statements that express laws] state what happens if nothing interferes characterizes their form, not their content" (2005: ch.6 msp.19).⁶² Michael Thompson helpfully suggests that,

⁶⁰To see this, suppose that *S* is putatively moving from α to σ , when we encounter *S* at π . It is possible that *S* will stop short of its putative end σ , and end up having moved from α to π , or from α to p ; and it is possible that *S* will overshoot σ and end up having moved (e.g.) from α to ω . More is required to know that *S* is doing *B* than that *S* is engaged in a process of which it is possible that the process ends with *S* having done *B*.

⁶¹I am abstracting from the important differences between such laws and thus such types of process—in particular, from what is distinctive about genuinely goal-directed processes—that a full treatment of the representation of movement would have to attend to.

⁶²Cf. Rödl 2005: "The progressive statement looks beyond what is here and now. It does not look in front, or back or to the side; it looks up to *what happens in general*" (ch. 6, msp. 7–8). It is clear that what happens in general—what a law says happens—does not always happen, because a process exemplifying what happens in general—a process exemplifying a law—may be interrupted, in which case it will stop without finishing. This shows that the generality of statements of what happens in general cannot be understood through the device of universal quantification. But it cannot be understood through some special kind of generic quantification, either: if the intelligibility of statements exhibiting progressive and perfective aspect depends on the generic statements they exemplify, then the generic statement cannot in turn be understood as quantifying over such statements. As Rödl puts it, "[t]he generality of generic statements does not characterize the quantity of the subject but the way in which subject and predicate are conjoined" (2005: ch. 6, msp. 15).

[t]hough actual uses of the sign “going to V” often escape the symbol to which an apt *Begriffsschrift* would restrict it, so that it can be said in some cases to express simple futurity, ... it is, in its primary and most interesting use, an instrument for the expression of imperfective aspect—it expresses, as we might put it, the “prospective imperfective”, and no tense at all. [CHECK](2008: 142–3)

I propose that the “prospective imperfective” judgment ‘S is going to do B’ is equivalent to the judgment ‘S will do B if nothing interferes’ as I just explicated it—that is, where the qualification signifies the form of the judgment, not part of its content.⁶³ It follows from this that neither ‘S is going to do B’ nor the equivalent ‘S will do B if nothing interferes’ is falsified by the subsequent failure of S to do B, as long as the reason S did not do B was that something interfered and prevented S from doing B. We can judge, after the fact, that S *was* going to do B, that S *would have* done B, but that S didn’t do B because something interfered. Of course, if someone judges that S is going to do B, but in fact S does not do B, then what he judged was going to happen did not happen; but whether or not his judgment was true or false depends not on this, but on whether S would have done B were it not for the interference. By contrast, even if S would have done B had the interference not occurred, the fact that S did not do B falsifies an earlier statement of “simple futurity”—that is, where subject and event-predicate are joined with future tense and perfective aspect—‘S will do B’.⁶⁴

Thus, though the thought that S will do B (simple future) is not contained in the thought that S is doing B, the latter nevertheless contains the thought that S is going to do B (prospective imperfective). What makes it the case that what S is doing is a case of doing B is that, in doing what it is doing, S is going to do B. All knowledge of what is happening is, for this reason, knowledge of what is going to happen; knowledge, that is, of what will happen if nothing interferes, not knowledge of what simply will happen. By contrast, to know what *will* happen is to have knowledge of what is going to happen *plus* the knowledge—hard to come by, this—that nothing will interfere.⁶⁵

4.5 The temporal structure of intentional action is explained by its calculative structure

On this picture, then, we distinguish between a case in which S is *actually* doing B from one in which S is only *apparently* doing B but is in fact doing B’ by apprehending what is happening as exemplifying a law that is part of a system of laws (e.g. chemical laws) that together determine a generic kind of process (e.g. *chemical process*); the law that what S is doing exemplifies and through which it is understood specifies what doing B is and explains why S is going to do B. (How this works in detail for various kinds of processes is an interesting

⁶³Just to be clear: I am not claiming that typical utterances in English of the form ‘S is going to do B’ necessarily express prospective imperfective judgments, and thus admit ‘S was going to do B but never did’; ordinary language is used for much more than doing philosophy.

⁶⁴Cf. Comrie 1976: “It is important to appreciate the difference between...expressions of prospective meaning and expressions of straight future time reference, e.g. between *Bill is going to throw himself off the cliff* and *Bill will throw himself off the cliff*. If we imagine a situation where someone says one of these two sentences, and then Bill is in fact prevented from throwing himself off the cliff, then if the speaker said *Bill will throw himself off the cliff*, he was wrong, his prediction was not borne out. If, however, he said *Bill is going to throw himself off the cliff*, then he was not necessarily wrong, since all he was alluding to was Bill’s intention to throw himself off the cliff, i.e. to the already present seeds of some future situation, which future situation might well be prevented from coming about by intervening factors. Indeed, *Bill is going to throw himself off the cliff* might well be shouted as a warning to some third party to prevent the future situation from coming about” (64–5). Our question, of course, is what it would be for the seeds of some future situation to be present. It is clear that it cannot always be intention, for *The house is going to burn down* and *The house will burn down* (both said while the house is burning) exhibit the same contrast; and even if Comrie is right that Bill’s intention is that wherein those seeds of the future presently lie, it remains to be seen what his intention must be if it is to function in that way.

⁶⁵Cf. Anscombe 1963: §52, Hampshire 1965: 65, and Galton 1984: 110.

question that deserves a fuller discussion; but it would be out of place here, where our concern is what kind of knowledge of the future is present in an agent's knowledge in intention.) However, when I start intentionally walking to school, even though I have a certain mass, and am moving at a certain speed, it is not these physical facts and the kind of mechanical explanation they are fitted to enter into that explains why it is no accident if I end up having walked to school intentionally. It is not as if once I start walking to school, I am carried there by something like gravity. Nor is there a chemical or biological explanation. In fact, there aren't *any* laws of the relevant kind about people walking to school, or even about students walking to school, as there are about dogs moulting.⁶⁶ When I am walking to school intentionally, I am not actualizing a tendency that is described by a natural law and which law thereby explains, and is exemplified by, what is happening. Rather, what I am actualizing, or rather executing, is my intention to walk to school.

Thus, what explains why it is no accident if I end up having walked to school intentionally is that (i) I have knowledge in intention that I am walking to school, that is, I *intend* to walk to school, which as we have seen involves (ii) *knowing why* I am making my way to school—to return a book to the library—and (iii) *knowing how* I'm getting there—I know I'm getting there on foot, I know how to get there on foot (I know the route), and I know how to get about on foot (I know how to walk). This knowledge in intention is what accounts, on the one hand, for the fact that my walking to school has a definite end-point (the returns desk in the library), and so I know what will count as me having finished walking to school, being half-way done, and so on; and on the other hand, for the fact that the parts and phases of my walking-to-school process amount to one process that luck and accident might interfere with, but which they do not constitute.⁶⁷

We can reach the same point by reflecting on the concept of a means, as it figures in the calculatively-articulated knowledge in intention that permeates the intentional action. The intention in action is articulated, and articulates itself, into means and ends, and it is internal to the idea of the means–end relation that the means are (jointly) sufficient for the production for the end. Thus, the knowledgable comprehension of the means as sufficient to produce the end, which is internal to the self-conscious teleological (that is, calculative) articulation of intention in action, is thereby a knowledgable comprehension of the fact that it will be no accident if one ends up intentionally producing the end, for if it would be some accident that the end would be produced, then this would show that the means are insufficient, and that any thought of them as being sufficient could not have been knowledge. And what I have said in this paragraph and the last, about what explains why it will be no accident if I end up having walked to school intentionally when I am walking to school intentionally, can be straightforwardly transposed to explain why it will be no accident if I end up having walked to school intentionally when I intend to walk to school (when, that is, I am going to walk to school intentionally): in this case too, the calculative articulation of the intention for the future (which may still be *in potentia*) shows why it will be no accident if the intended end is produced.

⁶⁶Of course, there might be *statistical* laws about, say, students walking to school. But a statistical law cannot explain the temporal unity of the events that exemplify it, because it generalizes over events that are antecedently comprehended (cf. n. 62 above). A law of the relevant kind would explain the temporal unity of the events that exemplify it. Perhaps Kant's "laws of freedom" do this; unfortunately I cannot pursue the question here.

⁶⁷Recall from n. 30 above the contrast between, on the one hand, my heart's pumping because it is circulating blood round my body, the teleology and temporal unity of which is explained by the generic statement 'the heart pumps blood in order to circulate blood round the body', and, on the other, the pumper's pumping because he is poisoning the inhabitants of this house, the teleology and temporal unity of which is explained not by the (false, or perhaps unintelligible) generic statement 'the pumper pumps in order to poison the inhabitants of the house', but rather by the pumper's intention.

Moreover, this account also explains why there is another way in which an intention (in action or for the future) might end up not being executed, other than if something external interferes: the agent may change his mind. If something external interferes, then there is a real process underway, or that is going to be underway when the time comes and towards which there is thus already a real tendency, that is prevented from reaching completion by something accidental to it. If the agent changes his mind, then the process (or tendency towards it) simply is no more; the real tendency towards the completion of the action has not been interfered with, but abolished. As with the “qualification” *if nothing external interferes*, this possibility is properly represented in the form of the agent’s knowledge rather than the content. However, unlike that “qualification”, which belongs to the form of knowledge of what is going to happen (knowledge of movement in progress or in prospect), the “qualification” *unless I change my mind* belongs specifically to the form of knowledge in intention (of intentional action in progress or in prospect), reflecting intention’s distinctively self-conscious causality.

We saw, in §2.5 above, that the distinction between doing *C* by doing *B* accidentally and doing it non-accidentally was explained by the source of the relevant teleology: in a machine’s design, in an organism’s life-form, in an agent’s calculatively-articulated intention. Here, we see that the non-accidentality of the accomplishment of the end specified by *do B*, or equivalently, the *temporal unity* of an event falling under the concept *do B*, has a similar kind of underlying explanation. Mechanical laws underwrite the temporal unity of mechanical events; chemical laws underwrite the temporal unity of chemical events; biological laws underwrite the temporal unity of biological events.⁶⁸ The design of a machine underwrites the temporal unity of the actions it carries out. And intentional actions, through their calculative articulation, underwrite their own temporal unity: *the calculative structure of an intentional action constitutes its temporal structure*.

So while Anscombe was right to insist that the presence of intention is not what explains the legitimacy of present progressive judgment *in general*, it in fact turns out that when the concept *do B* is an intentional action concept, the legitimacy of its present progressive predication *does* depend on the presence of intention in the agent. If an agent is doing *B* intentionally, then he is going to do *B*, where this consists in his intending to do *B*, of which intention his doing *B* intentionally is the execution (in progress).^{69, 70}

⁶⁸Cf. Rödl 2002: §3, Rödl 2005: ch. 6, msp. 11–12, and Galton 1984: 142.

⁶⁹It follows from these remarks that I am committed to what Michael Bratman calls “the Simple View”, namely, that if *S* is doing *B* intentionally, then *S* intends to do *B* [CHECK FORMULATION, CITE], the truth of which is much debated in the literature. I cannot engage that debate here; it must suffice for now to say that the account of the calculative and temporal structure of intentional action (in progress and in prospect) that I have been developing here constitutes what I take to be the deep ground for a defensible version of the Simple View.

⁷⁰It would be a mistake to read me as claiming here that we are to understand intentional action as the execution of a calculatively-articulated intention, where the latter notion is antecedently understood. Such an account would be received as unproblematic by many contemporary theorists, but our reflection on the temporal structure of action suggests that it should not be, as Rödl notes, “On any conception of the will, wanting to do something is an act of the mind that is such as to manifest itself in one’s doing it. Thus we have no idea what it is to want to do something if we do not comprehend what it is to do something. Statements about wanting have content only because progressive statements are possible, and therefore the latter are not possible through the former” (2005: ch. 6, msp. 3; cf. Boyle and Lavin 2010). What Rödl says of wanting can be said of intending, too: we have no idea what it is to intend to do something without knowing what it is to do something. Thus, it seems, we cannot appeal to intention to explain the possibility of intentional action (cf. Anscombe 1963: §4).

However, though I have argued that, where *do A* is an intentional action concept, we cannot understand ‘*S* is doing *A*’ independently of ‘*S* intends to do *A*’, this is not to say that I think we can understand ‘*S* intends to *A*’ independently of ‘*S* is doing *A*’. Rather, the view I have been advancing should be understood as revealing that the two forms are interdependent: on the conception I recommend,

We are now in a position to distinguish properly between Falvey and Unlucky John. What makes it the case that John was walking home, not walking three-quarters of the way home (supposing he made it half-way), is that he was walking home intentionally; he intended to walk home and he was executing that intention, and thus he knew that what he was doing was: *walking home*. By contrast, the explanation of why Falvey wasn't running a four-minute mile, though he may briefly have seemed to be, is that he wasn't in a position to have the relevant intention, let alone execute it. In our earlier discussion of intention and aspiration (§3.1), I said that an intention to do *B* requires knowing how one *can* do *B* (that is, knowing how to do *B*) though it need not require knowing how one *shall* do *B* (that is, knowing how one is going to do *B* on this occasion), though the intention understands itself as needing to become further calculatively-articulated so as to settle this question when (or before) the time to do *B* comes. Now, Falvey cannot run a four-minute mile, and so he cannot know how he can run a four-minute mile; in other words, he doesn't know how to run a four-minute mile, and so cannot intend to run one.⁷¹ Thus he is like Setiya with respect to dancing the tango; right now, Falvey can *aspire* to run a four-minute mile, but all he can *intend*, at the moment and in this respect, is to train.

One might wonder whether we need to say all this; after all, if Falvey cannot run a four-minute mile, doesn't this suffice to explain why he wasn't running a four-minute mile? (One cannot do what one cannot do, and thus one cannot be doing what one cannot do.)⁷² Why bother taking this detour through whether or not Falvey was in a position to intend, and thus whether he knows how, to run a four-minute mile? The answer is this. Falvey's case exemplifies a distinctive kind of situation that, though not uncommon, does not possess a feature fundamental to intentional action: namely, that it seems that there is no such thing as having run a four-minute mile, but only through luck. By contrast, most types of intentional action are such that an agent might "perform" a token of the type, but only through luck and thus, it is commonly accepted, not intentionally. It is easy to imagine scenarios in which agents have replenished a water supply or walked to school or made a three-point shot or baked a cake or become rich or pulled off a coup or a heist only through luck; their plans and skills—their knowledge how to do the things they did—were either inadequate

where *do A* is an intentional action concept, the fundamental form of judgment is 'S is doing *A* intentionally'—better still, the first-personal expression of knowledge in intention 'I am doing *A*'—where this involves understanding the temporal unity of the action as constituted by its calculative structure; that is, by the agent's practical knowledge in intention.

⁷¹Some (e.g. Snowdon 2003) will want to resist this suggestion, and claim that what distinguishes Falvey from someone who *can* run a four-minute mile is not the latter's having more knowledge, but the latter's having more power. I discuss the relation between practical powers and knowledge how in §§5ff.

⁷²Care must be taken with this point. The principles (i) that it is possible to have been doing *B* intentionally, but yet not end up having done *B* (and not through change of mind), and (ii) that one cannot be doing what one cannot do, are consistent only if we deny the superficially compelling principle (iii) that failure to do *B* if one tries suffices to show that one cannot do *B*. At the very least, (iii) must be modified as follows: one's failure to do *B* in particular conditions *C* even though one tried suffices to show that one could not do *B* in those conditions (i.e. conditions in which something interfered). But against even this, I sympathize with J. L. Austin, when he writes: "Consider the case where I miss a very short putt and kick myself because I could have holed it. It is not that I should have holed it if I tried: I did try, and missed. It is not that I should have holed it if conditions had been different: that might of course be so, but I am talking about conditions as they precisely were, and asserting that I could have holed it. There is the rub. Nor does 'I can hole it this time' mean that I shall hole it this time if I try or anything else: for I may try and miss, and yet not be convinced that I could not have done it; indeed, further experiments may confirm my belief that I could have done it that time although I did not. ...According to [the traditional beliefs enshrined in the word *can*], a human ability or power or capacity is inherently liable not to produce success, on occasion, and that for no reason (or are bad luck and bad form sometimes reasons?)" (1956: 166 n.1). See also the discussion of fallibility in §5 below.

to their goals as the situations played out or, though they were adequate, it was not in fact through them that those goals were achieved. In such cases, the agents did what they intended or aspired to do, but that their doing those things depended substantially on luck leads us to judge their actions were not intentional (under the descriptions that make reference to the goals). Now we know why this is so: the calculative articulation of practical thought is the rule that runs through an action; it is that which thereby constitutes the action's temporal unity and which accounts for its tendency towards completion. If there is nothing that accounts for the tendency towards completion, then nothing answering to the relevant description (the one that makes reference to the point of completion) is happening; if there is such a tendency, but something other than a sound, calculatively-articulated practical thought (i.e. an intention) accounts for it, then what is being done is not being done intentionally, and if it ends up having been "done", or having "come to pass", it will not be an intentional action.

Thus, if someone who aspires to do *B* brings it about that he has done *B*, we can ask two questions: was the present progressive judgment 'He is doing *B*' ever true? And, if it was, what explained it? If its truth is explained by the agent's sound practical thought, we have a case of intentional action; if not, not. It would be a confusion, however, to suppose that if an agent does what he aspires to do, we have a case of "aspirational action". Or at least, it would be a confusion to think that an action that satisfies an agent's aspiration could be "aspirational" in the same sort of way in which an action that is the execution of an agent's intention is intentional. This is because, whereas "intentional" specifies the principle of temporal unity of the event, "aspirational" specifies no such principle; rather, it marks the fact that though there is a "match" between what the agent wanted and what happened, it was some accident, from the point of view of practical thought, that the agent got what he wanted. Being "aspirational" or "lucky" is not a way of being an action, it is a way of failing to be an intentional action.⁷³ Or rather, it is a way of failing to be intentional under the relevant description: if there is action at all it is because the agent did *something* intentionally, but not what he aspired to do, and the coming to pass of that to which he aspired must be traced to an origin external to his practical thought (and powers).⁷⁴

4.6 The cognition condition, reformulated again

I have argued that someone who is doing *B* intentionally, or who intends to do *B* in the future, is going to do *B*; even though it is possible that he will not do *B*, it will be no accident if he ends up having done *B*. This is likely to be met with scepticism, for it follows by contraposition that if the agent is not going to do *B*, then neither is he doing *B* intentionally nor does he intend to do *B*. The distinction drawn between the prospective imperfective 'is going to do *B*' and the simple future 'will do *B*' shows that this is not absurd, but it will be suspected that it renders the point trivial. At the very least, it will surely seem as if the account I have offered must fail to do justice to what many think of as the central insight of Anscombe's claim about

⁷³ See n. 42 above.

⁷⁴ Consequently, if an agent is doing *A* intentionally, but his doing *A* is not a phase of a process that tends towards his doing *B*, then though he may be doing *A* intentionally because he wants or aspires to do *B*, we do not say that he is doing *B*. This is how we should interpret Anscombe's remark, mentioned in §2.3 above, that "the less normal it would be to take the achievement of the objective as a matter of course" (1963: §23), the less we are prepared to use the present progressive. (Anscombe in fact says that in such cases "the more the objective gets expressed *only* by 'in order to'", but if we are reading "in order to" as "with the intention to", then this cannot be right; rather, we should say "He's going to London because he wants/aspires/is trying to make his uncle change his will".)

agent's knowledge: the idea that an expression of intention says something not just about the speaker's state of mind, but about what is happening or going to happen in the world—just as might an expression of belief. That someone who is doing *B* intentionally is going to do *B* (going to end up having done *B*) is due to his intention to do *B*, I said, yet it is compatible with this that he ends up not having done *B*. Thus, the objection runs, the successful prosecution of the action is present merely in the agent's mind as an object of intention; the only kind of knowledge to be had here is knowledge of that state of mind, not knowledge of what will happen in the world. This is so regardless of any linguistic sleights of hand I might have made regarding the expressions 'is going to do *B*' and 'will do *B*'. If it turns out that knowledge of what is going to happen is not, as such, knowledge of what will happen, but rather knowledge merely of what the agent intends to do, then the spirit of Anscombe's claim has been abandoned, whether or not lip service has been paid to the letter.

The objection will not do, however. Consider a parallel case: a boulder has begun to roll down a hill. We know that it is going to roll down the hill to the bottom; that it will reach the bottom if nothing interferes (we do not know that nothing will interfere). There is a real process underway here; real tendencies, which are systematically captured by the laws of mechanics, are being actualized. This is not something "merely subjective", something not about the world but merely concerning the mind. Likewise, when someone has begun to walk to school intentionally, this too is a real process underway, a process which to be understood as the very process that it is must be understood as one that will terminate in that person's having walked to school unless something interferes. That we are not certain that he will end up having walked to school does not make it the case that our claim that he is going to walk to school expresses something only about his or our states of mind.⁷⁵

One might be prepared to grant this, however, and yet resist the seemingly stronger claim that knowledge that someone who intends to go shopping at the weekend is, by that very fact, going to go shopping at the weekend is knowledge of anything more than their state of mind. But such resistance amounts to a refusal to acknowledge the fact that intention is efficacious. It is true that if you know that someone wishes she could fly to the moon or believes that Hume died in 1776 then what you know is something about that person's state of mind and not (simply as such) something about what is going to happen in the world. This is because neither mere wishes nor beliefs are practical thoughts: they are not as such efficacious. But an intention, as distinct from other species of desire such as aspiration and more-distantly related acts of mind such as mere wish, is sufficient to produce the intended act, as being moved with a certain degree of force is sufficient (albeit in a specifically different way) for a boulder of a certain mass, size, and shape to roll down a hill of a certain decline. Assuming that what the agent intends is to do something "worldly"—like go shopping or walk home or write a novel—and not merely "mental"—like sing a song in her head or imagine what it would have been like to be Napoleon—then knowledge of her intention is knowledge of a real tendency towards something's happening in public reality. Our knowledge of what someone is going to do intentionally is no less objective and worldly than our knowledge, of an event of any kind, that it is going to happen.

But does *the agent* who is doing *B* intentionally need to *know* that it'll be no accident if she ends up having

⁷⁵Cf. Anscombe: "If I say I am going for a walk, someone else may know that this is not going to happen. It would be absurd to say that *what* he knew was not going to happen was not the very same thing that I was saying *was* going to happen" (1963: §52).

done *B*? Isn't it enough that it'll *be* no accident—why should she need to know it? One might worry that the account I have offered amounts to an implausibly intellectualized conception of agency, one that supposes that an agent somehow has some grasp of the complicated and quite possibly tendentious metaphysical views outlined in the present section; in other words, one might worry that I am supposing that one must become a philosopher of action before one can act intentionally.⁷⁶ But this would be a mistake. I have tried to situate my account of the temporal unity of intentional action within a broader framework, one which reveals that intentional action is a distinctive species of event in virtue of the form that the temporal unity of intentional actions takes. But the specific form that this unity takes is simply the calculative structure that belongs to intentional action insofar as it is, as it always is (at least *in potentia*), articulated into means and ends. The idea that someone who is doing *B* intentionally knows that she is going to do *B*, that it will be no accident if she ends up having done *B*, is just this: she knows that she is doing *B* by doing *A*, where this is the 'by' of calculative unity, and thus she knows that doing *A* is as such sufficient for doing *B*, and thus, in virtue of this sufficiency, that it will be no accident if, in doing *A*, she ends up having done *B*. Conversely, to suppose that it could be no accident that she would end up having done *B* despite not herself knowing that this was so would be to suppose that its being no accident derived from something other than her own practical thought, and thus to suppose that she is not doing *B* *intentionally*.

We can now conclude our investigation into the temporal structure of action. If an agent is doing *B* intentionally, then she is going to do *B*: it will be no accident if she ends up having done *B*, and she will end up not having done *B* only if she was prevented or changed her mind. That her success will be no accident is explained by her intention; if someone merely aspires to do *B*, and brings it about that she has done *B*, this is not an intentional action. For her to be doing *B* intentionally, she must know that she is doing *B* intentionally, and thus she knows that it will be no accident if she succeeds. That is, she knows she will succeed if she is not prevented. The same goes for intentional action in prospect as for intentional action in progress: if an agent intends to do *B* on Tuesday, then she will do *B* on Tuesday unless she is prevented (or unless she changes her mind, in which case the antecedent is no longer satisfied and thus her doing *B* on Tuesday is no longer on the horizon). It may be objected that she does not know with absolute certainty that she will succeed; but neither does she know with absolute certainty that the sun will rise tomorrow—for that event may be prevented, too.

Anscombe's intuitive idea with which we began in §1.1—that if I am doing *B* intentionally then I know (without observation) that I am doing *B*—has been revealed to be both more complex, and more deeply grounded in the structure of intentional action, than we might have thought. An intentional action is calculatively and temporally structured, and thus so must be its representation. This is so whether the action represented is in progress or in prospect. The calculative structure, which belongs essentially to the agent's practical thought, underwrites the temporal structure, which thereby too belongs to the agent's practical thought. If I am doing *B* intentionally, then I am doing it by doing *A* intentionally and because I am doing *C* intentionally, and for each of *A*, *B*, and *C* it will be no accident if I succeed in bringing them off.⁷⁷ Moreover,

⁷⁶Thanks to Anton Ford for pressing me to address this objection.

⁷⁷Two qualifications. First, I am assuming that my doing *B* is not a basic action, of which there will be more discussion below. Secondly, it is possible to do *B* intentionally because one takes oneself to be doing *C* intentionally, and yet one is not doing *C* intentionally either because one is not doing *C* or because one is doing *C* but only through luck (i.e. either of **conditions (i) and (ii)**)

this calculative and temporal structure is what I have knowledge in intention of, which knowledge constitutes that structure. This is what it is for ‘I am doing *B* intentionally’ to be true. And likewise for ‘I am going to do *B*’.

5 Fallibility and the carbon copier

We have derived a more adequate understanding of the cognition condition from its source in the calculative and temporal structure of intentional action (in progress and in prospect); now we are in a position to respond to those examples, like Davidson’s carbon-copier, that purport to show that one may be doing something intentionally, or intending to do something, while remaining agnostic about whether one is doing it, or going to do it. (We saw in the previous section how and why one’s knowledge in intention that one is doing *A* contains knowledge of eventual success—knowledge that one is going to end up having done *A*.) To defend the cognition condition I have outlined, then, I must show that either the carbon copier *does* know that he is making the copies while he is intentionally making them, or, if he doesn’t know that he’s making them, he *isn’t* making them intentionally. Now, as it stands, Davidson’s carbon-copier case is underdescribed (this is unsurprising—it originally appeared as a parenthetical remark). Happily, though, when we flesh it out, either it ends up being a case of intentional action that satisfies the cognition condition, or it ends up not being a case of intentional action. The relevant kind of fleshing out that is required consists, unsurprisingly, in providing the calculative structure within which we are to situate, and thus grasp, the copier’s action and its temporal unity.

First, consider the following mundane expansion of the story. The copier needs to produce ten copies of a document to give to his boss. He reasons that the fastest way to do this would be to interleave carbon paper between the ten pages, and make the copies all in one go. So he does this, and he presses very hard on the pages as he writes. Suppose that when he inspects the documents, he finds that the impression carried through to the seventh copy, but not beyond. What happens next? Remember, the copier needs to give ten copies to his boss; thus, having made only seven copies, his inspection reveals not that he has failed but that he *hasn’t finished yet*. (It is not as if, not having made all ten copies in one go, he won’t be able to give ten copies to his boss.) So he places the seven copies he has made to the side, and writes again on the eighth, again pressing firmly, so that the impression carries through to the ninth and tenth pages. He has now made ten copies, intentionally; and all the while, from beginning to write on the top sheet until checking the tenth impression, he was making ten copies. Moreover, he knew that this is what he was doing; his practical knowledge that he was making *ten* copies, not seven, is what *sustained* his action through to the next phase: the calculative unity of his action constitutes its temporal unity. It’s true that he was initially trying to make all ten copies at once, but the calculative and temporal context of his action—the facts that he needed to make ten copies,

might not obtain). (A parallel situation is one in which one is doing, or intends to do *B*, because one aspires to do *C*.) If one is not doing *C* intentionally, where this is not just because one is going to do *C* but is not doing it yet, then one does not have knowledge in intention that one is doing *C*, or that one is going to do *C* (in the relevant sense in which that knowledge is, or rather would be, contained in the knowledge that one is doing *C*). Does this impinge upon whether or not one is doing *B* intentionally, or whether or not one knows that one is? It depends—on whether or not the error in one’s practical thought or performance in virtue of which one, though doing *B* intentionally, is not doing *C* intentionally is grievous enough that it impugns the status of *C* as the measure by which one assesses and regulates one’s doing *B*.

it didn't really matter whether he did them all in one go or not, and that he continued to do things with the intention of making more copies until there were ten of them—reveals that this was just one possible means he could have taken. As it happened, the means he chose only amounted to getting seven tenths of the job done; so what he thought was the whole of his action turned out to be a proper part. But this is no objection to maintaining that while he was doing the parts, he was doing the whole, and thus no objection to maintaining that while he was doing the parts, he knew he was doing the whole.⁷⁸

What happens if we flesh out the case in a way where it matters that the copier produces all ten copies in a single act of carbon-copying: if his job or his life hangs on it? In such a case, if the first inspection reveals that only seven copies have been made, then this *does* amount to the revelation that the copier has failed, and not merely that he hasn't finished yet, because what determines whether or not he has finished is not his having made ten copies, but his having finished writing on the top sheet (once). But it's not clear that a copier under such a strange demand is in a position to intend to make all ten copies at once. He can aspire to, of course; he can give it a shot and hope for the best. But if you aspire to make ten copies and you bring it about that ten copies are made, you have not necessarily acted intentionally. If you bring it about that ten copies are made through intention, or—what comes to the same thing—through knowledge in intention, then this is intentional action; if you bring it about through luck, it is not. So if the copier doesn't have knowledge in intention that he's making ten copies, because he's not in a position to genuinely intend to make ten copies, he is not making ten copies intentionally, even if he ends up having brought it about that there are ten copies of the document.⁷⁹

One way to take Davidson's carbon-copier example would be as material for a sceptical argument, an action-theoretical parallel to the argument from illusion. According to this argument—call it *the argument from failure*—an agent who is doing *A* intentionally does not know simply in virtue of his practical thought and agency that he is doing *A*; it “practically seems” to him as if he is doing *A*—that is, he is trying to do *A* and he knows this—but it would be unjustified of him to infer from the way things seem to the way they are, because when an agent tries and fails to do something, things are for him, from the point of view of his practical thought and agency, just as they would be if he succeeded. Knowledge that he is *actually* doing what he intends to be doing or is trying to do would require supplementing his practical thought with something extra-practical: observational knowledge, say, of what he is doing. If the carbon-copier ends up having succeeded in making his copies, then it will have been the case while he was making the copies, Davidson tells us, that

⁷⁸One might worry that this vindication of the claim that the carbon-copier, in making ten copies intentionally, thereby knows that he is making ten copies, fails to show that his knowledge is (all) in intention, and that it is not derived, at least in part, from observation. Perhaps agent's knowledge consists of two factors, non-observational knowledge of what one intends to be doing, and observational knowledge of what one is actually doing (Keith Donnellan (1963: 407) and J. L. Mackie ([CITE] 104–5) advance such views; see Falvey 2000 for criticism). But this would be a mistake: when the copier observed that he had only made seven copies, he was not observing what he *was doing*, but rather what he *had done*. As Falvey puts it, “Observation enables the agent to notice and correct mistakes; but what counts as a mistake here is determined by what the agent is doing” (2000: 29); observing a mistake presupposes knowing what you're doing, which knowledge therefore cannot be observational. Cf. in this context Anscombe's remarks on the role of the eyes as “an aid” in writing (1963: 53) and her discussion of the “two knowledges,” one practical, the other speculative, involved in any intentional action (1963: 88–9).

⁷⁹Cf. Thompson (unpublished ms. [CITE]); but compare his remarks, in ‘Naive Action Theory’, on whether intending to do something involves thinking that one will end up doing it (2008: 102–3, 145 n.29), or whether asserting ‘*S* is doing *B*’ in any way commits the speaker to the eventual truth of ‘*S* did *B*’ (2008: 144).

he was making them and making them intentionally. But while he is making them, supposing that he is in fact making them, things are for him as they would be if he were merely trying to make them; thus, all he knows is that he is trying.

But the argument is too quick. To see this, return to our second expansion of the case, where the carbon-copier tries but fails to make his ten copies all in one go. Consider what was true when he was doing whatever he was doing in order to make the copies—pressing very hard with his pen, say. Plausibly, it is true that he was then *trying* to make the copies, but it is not true that he was then *making* the copies. Contrast Unlucky John, considered before his accident: he *was* then actually walking home. (Perhaps it is also true that he was trying to walk home, but this should not detain us.) Our reflections in §3.1 and §4.4 above have put us in a position to distinguish between two different notions of “failing to do something”: (i) being either *prevented* from doing something one had a genuine intention to do or *interrupted* in the course of intentionally doing it, and (ii) failing to bring about a state of affairs or effect that one *merely aspired* to bring about, which, with respect to the relevant description, one was never doing intentionally. The failed copier’s failure is roughly of a piece with Falvey’s failure⁽ⁱⁱⁱ⁾ to run a four-minute mile, not of a piece with Unlucky John’s failure⁽ⁱ⁾ to make it home; and, as we have seen, bringing about a state of affairs or effect that is the content of a mere aspiration is not acting intentionally, under the relevant description. The argument from failure simply equivocates between these two notions of failing to do something. Trying to base an argument from failure against the possibility of practical knowledge of intentional action on the fraughtness of difficult tasks like making ten carbon-copies at once or running a four-minute mile is like trying to base an argument from illusion about perceptual knowledge on the difficulty in distinguishing far-away square towers from far-away round towers: no one should be convinced by either.⁸⁰

However, my remarks about the second way of fleshing out the carbon-copier case might make it sound as though actions that involve doing something in one shot, so to speak, can never be intentional. For what made it the case (in the first variation) that the copier was making ten copies the whole time was that it didn’t matter whether he made them all in one go and he knew this and knew how to cope with fact that, as it turned out, his first pass resulted in the production of seven copies. But surely it is possible to intend to sink a birdie putt, and indeed to sink one intentionally, even though if you miss it the best you can then do is make par; there are a whole host of kinds of actions that depend on being done in one shot, or might depend on it in particular circumstances. And it might seem that I have committed myself, implausibly, to saying that, when someone misses a birdie putt because of an error of performance—not because someone yelled while she was striking the ball, or because a gull swooped and made off with the ball while it was rolling towards the hole, but just because she mis-hits it—she thereby reveals herself to have had merely the aspiration of making a birdie, and the intention (of underwhelming ambition) to just finish the hole, however many strokes it ends up taking.

⁸⁰If you “try” to do something and fail, it may or may not be clear to you or anyone else whether you failed to “do” it because you were prevented or because you weren’t in a position to do it intentionally. (The cases of temporary bodily paralysis that are so popular in the literature can I think be construed, with enough back story, either way; I doubt that whatever intuitions we might have about such esoteric cases amount to data that an account of intention, action, or trying is under an obligation to respect.) The same thing goes for trying and succeeding, as I have already emphasized (§3.1 above): there is a difference between bringing about one’s goal through luck and bringing it about through knowledge (that is, between lucky action and intentional action); sometimes it is hard to make out in practice, and individual cases can be argued about.

Saying this would be a mistake, but doing so is not compelled by the conception of the cognition condition on intention and intentional action I have been advancing. The birdie-putt case brings out a third way in which someone's action might fail, in addition to (i) and (ii) above: (iii) failing to do what one intended to do because one's action is the exercise of a *fallible capacity*. That the skills of putting, shooting a basketball, and so on are fallible is obvious to anyone who possesses or admires them; but we can and sometimes do make mistakes in every area of intentional action. Our ability, in many cases, to "pick up" an intention that has "fallen to the ground" means that we can, in such cases, overcome our mistakes and do, and know that we are doing, what we intend to do. But it would be wrong to infer that, when the possibility of picking up the intention and carrying on isn't possible because of the one-shot nature of what one intends to do, one's exercising a fallible capacity *successfully* counts as a piece of luck that undermines the status of one's action as intentional and one's practical thought as an intention. That a skilled golfer *might* have missed her putt, placing a birdie irrevocably beyond reach, does not impugn the status of the putt that she holes through the exercise of her skill: that her capacity is fallible does not mean that every exercise of it is flawed.

This is not to say that a skilled golfer can *never* be said to have been lucky in making a putt—she might hit a putt slightly off-line, due to a flawed exercise of either putting judgment or execution, that nevertheless "lips in" when it might as well have "lipped out"—but only that she should not *always* be said to be lucky in this sense. (Recall, in any case, that the distinctions between intention and aspiration, between intentional action and making things happen merely by luck, are not binary.) We must not conflate the first and third senses in which one might fail to do something: the world may, from time to time, "do her a favour" by, so to speak, permitting what ought to have been insufficient to get the job done, but the only sense in which it does her a favour every time she succeeds is by not permitting e.g. a gull to fly off with the ball. Indeed, it is part of what a skill is that it provides (fallible, yet resilient) security against particular points of vulnerability that beset the kinds of action and activity for which the skill is useful; not only is it no accident that when an agent who possesses the skill of doing *A* exercises that skill, she ends up having done *A*, it is no accident if she ends up having done *A* despite the presence of certain typical hampering factors, that might interrupt or prevent an agent who, say, possessed the skill to a lesser degree.

This might seem like a lot of fuss to make. Why shouldn't we just accommodate apparent counterexamples along the lines of Davidson's carbon-copier, by weakening either the content or the scope of the cognition condition? Although space will unfortunately not permit me to go into the details of all of the various proposals made on this score, we are already in a position to see that the basic answer is this: if we suppose that the carbon copier could be making ten copies intentionally without knowing that he is doing so—either because he knows only that he is trying to make ten copies (Davidson), or that he is doing something "more basic", like writing firmly on the top sheet, with the aim of making ten copies (Setiya), or because it's just an atypical case (Gibbons)—then we are simply taking for granted the calculative and temporal unity of the action while at the same time denying the presence of that which could account for it, the agent's knowledge in intention. Anscombe, commenting on her deployment of Aquinas's slogan, "practical knowledge is the cause of what it understands", writes that it "means more than that practical knowledge is observed to be a necessary condition of the production of various results; or that the idea of doing such-and-such in such-and-such ways is

such a condition. It means that without it what happens does not come under the description—execution of intentions—whose characteristics we have been investigating” (1963: §48). The cognition condition is woven into the very fabric of the intentional action that is cognized in agent’s knowledge. The pull of the carbon-copier case is felt only if we imagine that the structure of action will take care of itself, independently of the agent’s knowledge of it; but because her knowledge is constitutive of the action, calculatively and temporally, we can’t simply take the action for granted and then ask whether or not she has knowledge of it, and if so how. Thus, reflection on the calculative and temporal structure of intentional action reveals both why we must and how we can reject the idea that carbon-copier style cases pose a threat to the cognition condition, and that they must be accommodated by weakening the condition’s content or scope.

6 Why there must be basic action

6.1 Practical *knowledge*?

However, even if we should not weaken the content or scope of the agent’s cognition, perhaps we should weaken its *kind*. Recall that contemporary cognitivism’s appropriation of Anscombe’s thought proceeds in accordance with a programme of epistemological “modesty,” as Setiya puts it, of “setting aside the claim to knowledge” (2007: 25) in favour of a claim about what the agent *believes*. The point of weakening the condition from knowledge to belief can’t be to respond to the worry that the belief might be *false*, because if I am doing *A* intentionally, then my belief that I am doing *A* is true: the cognition condition is a condition on something that, if it obtains, guarantees that the cognition in question is true.⁸¹ Rather, the point of weakening the claim in this way must be to respond to a worry about *justification*. Consequently, one might reasonably wonder how my account of an agent’s *cognition* in intention is supposed to amount to *knowledge* in intention.

A number of proposals about the epistemology of agent’s knowledge have been made in the recent literature. Unfortunately, I cannot consider them all here. I want to focus on what strikes me as the most plausible proposal, in light of the account developed thus far in this essay. The proposal, due to Setiya (2008; 2009b), is this: “One is epistemically justified in forming the belief that one is doing *A* involved in doing *A* intentionally, only if, and because, one knows how to do *A* (2009b: 128).⁸² The knowledge how that is supposed to justify, in part, one’s knowledge in intention must be general knowledge how, not particular knowledge how, for the latter is already contained in one’s knowledge in intention.⁸³ This account fits well with the story I have told: whereas I focused on the metaphysical role of general knowledge how in constituting the calculative and temporal unity of an intentional action (in progress and in prospect), Setiya points to its epistemological role. And given the account of the structure of intentional action provided here, according to which the metaphysics and epistemology of intentional action are really the same thing, looked at from different angles,

⁸¹We have also seen how and why if I intend to do *A*, it is true that I am going to do *A*.

⁸²Setiya thinks that knowledge how to do *A* is a necessary condition on one knowing in intention that one is doing *A*; he remains neutral on whether one also needs “knowledge of ability in the simple conditional sense” (2008: 407): if I intend to be doing *A*, I will be doing *A* in fact, or whether “it is sufficient that one knows how to do *A* and has no reason to doubt one’s own ability” (2009b: 134, variables modified).

⁸³See §3.1 above. If one’s knowledge in intention is knowledge of what one is going to do in the future, it may as yet contain particular knowledge how only potentially.

as it were, it should be no surprise if general knowledge how is the metaphysical basis of intentional action and the epistemological basis of knowledge in intention.

In §3.1 above, I distinguished between two kinds of general knowledge how: procedural knowledge and skill. Setiya draws, in effect, the same distinction, from the angle of the kind of action it is knowledge how to perform: knowledge how to perform a non-basic action—one performed by performing other actions as means—“consists in knowing basic means and knowing how to take those means,” whereas knowledge how to perform a basic action *A*—performed not by taking further means—“consists in the disposition to do *A* in execution of my intention” (2009b: 135, variables modified). It is relatively straightforward to apply Setiya’s suggestion about the justification of an agent’s cognition of her intentional action to cases of non-basic action, where the general knowledge how deployed is procedural knowledge. Procedural knowledge that one can do *B* by doing *A* is not evidence that one is in fact doing *B* by doing *A*; rather, one applies, or puts into practice, one’s procedural knowledge in here-and-now doing *B* by doing *A*, deriving one’s intentional action, and thus one’s calculatively-articulated knowledge in intention, from one’s intention to do *B* and the procedural knowledge.

However, it is less clear how non-procedural general knowledge how—skill—justifies an agent’s cognition of her basic intentional action:

We can say at least this. If I know how to take those basic means, this knowledge consists in the disposition to execute the corresponding intentions. Since I have this disposition, it is no accident that, when I intend and thus believe that I am [taking the basic means], I am doing so in fact. Exercising basic knowledge how ensures non-accidentally true belief. To say this is not to endorse an epistemology on which its being no accident that a belief is true suffices for it to count as knowledge. But it does preempt a residual source of skepticism, that the truth of beliefs formed without sufficient prior evidence could only be a matter of luck. When they are constituted by intentions and one knows how to perform the relevant actions, that is not the case. (Setiya 2009b: 136)

How can a disposition amount to knowledge in more than name, without adopting a reliabilist framework? Setiya suggests that the view that (non-procedural) knowledge how “does epistemic work helps explain why it should count as knowledge” (2009b: 134–5), but it remains quite unclear how it is supposed to perform any epistemic work unless it already counts as knowledge.⁸⁴ In the final sections of this essay, I turn to the parallel questions of whether there can be calculatively basic actions, and whether skill amounts to a genuine form of knowledge. There are compelling reasons to think that in the absence of positive answers to both questions, we will have failed to understand practical knowledge and the structure of intentional action. Unfortunately, there are also compelling reasons to think that positive answers cannot be given.

6.2 Calculatively basic action and skill

As we have already seen, the A—D order is many things at once: it is the order of action explanation and the order of practical reasoning, and it displays the structure of intentional action and the structure of the agent’s

⁸⁴Setiya’s claim that his account shows why knowledge how cannot be propositional doesn’t quite work: he claims that if it were, it would justify knowledge in intention inferentially, but knowledge in intention is non-inferential. But procedural knowledge how is, or at least can be, propositional; knowledge in intention of non-basic action is not by that fact inferential knowledge. What this shows is that one piece of propositional knowledge may non-inferentially justify another. Therefore the argument that non-procedural knowledge how—skill—is non-propositional must be made on other grounds.

knowledge in intention. An intentional action's calculative structure can be traversed, as we have seen, via the *reason-seeking* question 'Why?' and the *means-seeking* question 'How?'. Anscombe thinks that the chain of 'Why?'s, which articulates and extends the A—D order considered as the order of action explanation, comes to an end. It ends when the question 'Why are you doing *D*?' receives in response a "desirability characterisation" (e.g. 'Because it's *fun*' or 'Because it's *pleasant*', 'Because it's good for my *health*', and 'Because it *befits* an *X* to do *D*'); issues concerning the "decency as an answer" of such desirability characterisations to the question 'Why?' transcend action theory—they "belong[] properly to ethics" (1963: §40).

It seems that the chain of 'How?'s must also come to end as well, on pain of regress. Given what we said about the A—D order, we could express the regress in different ways. With an eye towards the way in which the A—D order captures the structure of intentional action, we might think that, while many of the things we do by means of doing other things, there must be some things we can just do, without taking any means to them. Reflecting now on how the A—D order illuminates the structure of agent's knowledge, we might think with Jennifer Hornsby that "some things—at the end of [the] 'by'-chains, as it were—must be done without knowledge of procedures. These are things the agent does 'directly'. They are basic things, ...which we are inclined to say the agent is *able to simply do*" (2005: 114). And recalling that the A—D order is also the order of practical reasoning, we might think that if, as Anscombe puts it, "The mark of practical reasoning is *at a distance* from the immediate action, and the immediate action is calculated as the way of getting or doing or securing the thing wanted" (1963: §41), then, unless practical reasoning is to go on for ever in vain, without coming to a conclusion, there must be immediate actions, the doing of which need not be mediated by calculation resulting at some supposedly 'more' immediate action for which, of course, the same problem would arise.⁸⁵

From the fact that the A—D order is the order of the structure of intentional action, the order of agent's knowledge, and the order of practical reasoning, it is plausible to conclude that these regresses, of the dependence of intentional action on intentional action, of the agent's knowledge how to do things, and of practical reasoning, are in fact *the same regress*. We need an account, then, of something that stands to the question 'How?' as the notion of a desirability characterisation stands to the question 'Why?'. And whatever we make of Anscombe's move in proposing that the question 'What is intentional action for?' belongs not to action theory but to ethics, a parallel move with respect to our question 'How is intentional action possible?'—to defer it to some other branch of philosophy—seems implausible. The suggestion made by Hornsby (1980a; 1980b; 2005; 2007), among others,⁸⁶ is that the regress is stopped by what we might call *calculatively basic action*.⁸⁷

The idea is that the calculatively basic things done are things an agent (i) can simply do—that is, do without doing them by means of doing something else—and (ii) knows how to do otherwise than on the basis of knowing that he can do them by means of doing something else. An agent's knowledge how to do things cannot be exhausted by procedural general knowledge how. Thus Hornsby's proposal: an agent is able

⁸⁵Cf. Enç 2003: 51; Ruben 2003: 61; O'Brien 2007: 164.

⁸⁶Enç 2003 and O'Brien 2007 make similar proposals, and I will occasionally draw on them as well.

⁸⁷What I call "calculatively basic", Hornsby refers to as "teleologically basic" (1980b: 88). I distinguished between calculative and other forms of teleology in §2.5. The term *basic action* has philosophical currency beyond the topic, to which the present discussion is restricted, of calculatively basic action.

to simply do things when he has, not procedural knowledge how to do them, but “practical knowledge” how to do them:

Someone whose knowledge how to do *A* is practical is able to simply do *A* (at least so long as it is actually possible for her to do *A*). Our practical knowledge provides for our doing all of the things and engaging in all of the activities which we do or engage in as agents. We can do, and may have reason to do, a very great deal more than we are able to simply do; but our doing these other things requires procedural knowledge that we can put into practice. Practical knowledge, when it has been learned and not forgotten, is what enables an agent to get started, as it were. (Hornsby 2005: 115)

A calculatively basic action, then, is the exercise of this kind of practical knowledge, which in the context of this essay we are better off giving a more precise name, non-procedural knowledge how, or skill. As we have seen, Setiya suggests that skill, in this sense, consists in a disposition to perform a basic action in execution of an intention.

On this conception of basic action, different things may be basic for different agents. This is because different agents have different skills, in virtue of acquiring them through learning and practice, as an example from John Searle illustrates: “For a good skier, making a left turn can be a basic action. He just intends to do it and he does it. For a beginner to make a left turn, he must put his weight on the downhill ski while edging it into the slope, stem the uphill ski, then shift the weight from left to right ski, etc., all of which are reports of the content of his intention in action” (1983: 100). Searle’s example brings out a further important point, that although bodily movements may be among the calculatively basic things done (as shifting one’s weight is, for the novice skier), so may actions “further out” (as making a left turn is, for the good skier). Now, the good skier’s left turn may be composed of a sequence of movements of just the same types as the beginner, but the point is that they do not figure for the former, as they do for the latter, as objects of his practical thought and attention, means taken to his end of making a left turn. As Anscombe puts it, “[i]n general, as Aristotle says, one does not deliberate about an acquired skill; the description of what one is doing, which one completely understands, is at a distance from the details of one’s movements, which one does not consider at all” (1963: §30). Cases in which the agent does not conceive of a bodily movement as his means to his end—where the thing “further out” is done “directly”—are not hard to find. A favourite example of Hornsby’s is that of performing speech acts, where, in one’s mother tongue at least, one does not make detailed adjustments to one’s mouth, tongue, etc. as a means towards correct pronunciation.⁸⁸ Similarly with the complex movements of the fingers made in typing or playing the piano—or even tying one’s shoelaces: if the agent made it his end to make those movements of the fingers, he would be best off typing or playing the piano or tying his shoelaces, for *those* are the things he does directly.

However, even with these remarks about the scope of basic action in place, Hornsby’s proposal, so far as I have described it, remains highly schematic. For, we have no idea yet what it is to be “able to simply do” something, or to do something “readily”, “directly”, “forthwith”, or “just like that”—to use some other phrases common in the literature;⁸⁹ indeed, our problem is precisely that our grip on what it is to do something

⁸⁸See, e.g., Hornsby 1980a: 389; Hornsby 2005: 120ff.; Hornsby 2007: 173.

⁸⁹Hornsby, Enç, and O’Brien try to capture the practical “directness” of basic action via locutions like these: one is “able to simply *A*” (Hornsby 2005: 115); one is “simply able to [*A*]” (Hornsby 2007: 168); one can do *A* “just like that” (Hornsby 1980b: 88); one “has immediate access” to doing *A* (Enç 2003: 52); one knows how to do *A* “without needing to know how to do it” (O’Brien 2007: 164); one knows how “to execute [*A*] just like that” (O’Brien 2007: 164); etc. Such locutions are rife in contemporary action theory,

has so far centrally involved the notion, captured by the A—D order, of doing it by doing something else. Similarly, we have no idea yet what it is to have non-procedural knowledge how to do something as opposed to procedural knowledge how to do something, or whether the former will be recognisable as knowledge. We have names for our regress-stopping concepts, but we do not yet have accounts of them. Hornsby's proposal will require some unpacking before it is quite clear what it amounts to, and whether it is satisfactory. But we will be in a better position to understand and evaluate her proposal if we familiarise ourselves with some of the difficulties in thinking about basic actions and practical knowledge that lie ahead.

7 Why there cannot be basic action

7.1 The initial segment argument

In 'Naive Action Theory', Michael Thompson inverts the focus Anscombe takes on the A—D order. Whereas Anscombe thought of intentional action as *subject* to the question 'Why?' (1963: §5), Thompson understands it, in the first instance, as an *answer* to the question 'Why?'. His "provisional[] attempt to isolate [the] category" of intentional action is this: "S's doing *A* is an intentional action (proper) under that description just in case the agent can be said, truly, to have done something else *because he or she was doing A*" (2008: 112, variables modified). And thus we could gloss his definition like this: intentional actions are those actions to which a certain sense of the question 'How?' is given application; the sense is of course that in which the answer gives a means. On Thompson's view, then, what it is to do something intentionally is to do it by doing something else intentionally. Thus, it seems, his definition can at best characterize only nonbasic intentional action; what illumination it offers the category of intentional action cannot be total. But this is no oversight; Thompson has an argument for the claim that his definition holds with full generality, which therefore amounts to an argument against the possibility of basic actions.

Thompson argues that intentional actions are shot through with calculative complexity of the part-whole variety. The idea of calculative part-whole complexity has obvious application to those intentional actions that are "intuitively resolvable into a heterogeneous collection of sub-actions that are themselves clearly intentional—organs, as it were, of the whole. Such is the relation of egg-breaking and egg-mixing to omelet-making, of brick-laying and door-framing to house-building, and of writing the letters 'a' and 'c' to writing the word 'action'" (2008: 106–7). But it seems equally obvious that this intuitive resolution "will come to a limit" (2008: 107): there are, intuitively, no intentional sub-actions that stand as parts to writing the letter 'a'. At such a point, then, we seem to find our final resolution of a nonbasic action—writing the word 'action'—into basic actions, which themselves have no intentional sub-actions as parts.

However, Thompson argues, "such resolution is not necessary": "*wherever* a completed individual action is under a description [of it as having taken time] it will be possible to find a true naive rationalization in which that description appears *in the explanans*. Even actions that, like arm-raising, do not divide in this way need not, after all, be viewed as pointlike" (2008: 107). The argument for this claim—which would underwrite Thompson's definition of intentional action as an answer to the question 'Why?' or, equivalently, as subject to the question 'How?'—is given in terms of considerations pertaining to those actions that involve moving

independently of focused discussions of basic action

things, including moving oneself. The idea is that if, for example, Smith intentionally raised his arm from his side (α) to his ear (ω), in order to scratch it, then he raised his arm from his side to, say, his shoulder (π), and intentionally: he raised it to his shoulder because he was raising it to his ear, and he knew while he was doing it that he was doing it and why, as we could have discovered if we had stopped him and asked him. If this is right, then it seems that every segment of the trajectory α – ω was an intentional action. Thus, Thompson claims, “[a]cts of moving something somewhere intentionally always have an initial segment that is also an act of moving something somewhere intentionally” (2008: 111). Even the movement α – β isn’t safe: it must have an “initial segment” α – α' , which itself is an intentional action that has as its own parts intentional actions, and so on. If this argument (call it *the initial segment argument*) is sound, then all intentional actions that involve moving something somewhere, even if it is only one’s limbs that are thus moved, are irreducibly calculatively complex, where the complexity is of the part–whole variety.

One might suspect that the initial segment argument turns on the peculiarities of local motion. But, as Sebastian Rödl points out, the the notions of part and whole are applicable in virtue of the in/complete contrast that belongs to the *temporal structure* of action: “‘S did A’ can be true only if there was a time when S was doing A, but had not yet done A, a time when the action was underway, but not yet complete. The idea of someone’s having performed an action has application only where the idea of his having been performing it has application” (2002: 329–30, variables modified). But what S had done, while S was doing but had not yet done A, is an action by means of which S was doing A, to which the same argument applies.⁹⁰ According to Thompson and Rödl, then, every part, in the relevant sense, of an intentional action is itself an intentional action, and every intentional action has proper parts in the relevant sense; thus there cannot be basic actions. The calculative structure of intentional action and its temporal structure apparently issue conflicting demands: there must and yet cannot be basic actions.⁹¹

7.2 What does the initial segment argument show?

What is the upshot of the initial segment argument for Hornsby’s proposal? It is not immediately obvious. By virtue of the thought that a basic action may be “at a distance from the details of one’s movements,” in Anscombe’s phrase, the proposal evades the objection that basic action is a kind of action-theoretical prime mover. Hornsby does say that skill, and thus basic action as she conceives of it, is what enables an agent

⁹⁰Thus Rödl says that though Thompson’s argument shows “that it is *not impossible* that every part of an intentional action is, again, an intentional action,” he (Rödl) “claim[s] that it is *necessary*” (2002: 332, n.12). One might say that Thompson’s argument gives spatial application to what is essentially a temporal point: the lines in space and time along which Smith moves are isomorphic.

⁹¹The initial segment argument engenders natural suspicion. After all, we have come some distance from the intuitive idea that egg-breaking is a part of omelette-making, and that both are intentional actions. Anton Ford (ms.: §3) argues that the temporal structure of actions, though it ineliminably brings in the in/complete contrast, does not thereby bring in the notions of *part* and whole, but rather *phase* and whole; he urges us not to confuse the parts of an action with its phases (for example, the parts of an action can overlap, but an action’s phases cannot). Most interestingly for us, Ford claims that while the parts of an action are, as such, actions, the phases of an action are not, as such, actions. One might try to exploit Ford’s distinction, then, to block the possibility of reiterating the initial segment argument: if the initial segment of an action is a phase, not a part, and thus not itself an action, then the argument cannot get going. But this proposal will not work, because for any phase of an action there is a *corresponding* part. (This is not objection to Ford; his goal is not to rebut the initial segment argument.) The parts of *walking from α to γ* may well include *hacking away the underbrush* and *looking out for snakes*, but among their number are also *walking from α to β* and *walking from β to γ* . The fact to which Ford points—that parts of an action do not *as such* line up serially, and that it is quite unusual if they do—does not dislodge the thought exploited by the initial segment argument: that any action can be resolved into parts that *do* line up serially, and that such resolution can never come to an end.

to “get started”; but we need not take her to be expressing the thought that the agent does the basic thing *first*, with whatever is done nonbasically consisting of something else done, or accruing to the basic thing, afterwards. Nor is she committed to thinking that a basic action must be *atomic*—that it has no parts which are themselves actions. Recall Searle’s example of the good skier. The basic thing done is making a left turn. There is a sense in which he does this by putting his weight on the downhill ski while edging it into the slope, etc. But what interests Hornsby is the way in which the skier putting his weight on the downhill ski while edging it into the slope, etc. by making a left turn—this is the conception of doing one thing by doing another that her conception of basic action exploits. And the same thing goes for each of the component parts of the left turn, and their combination in a single left turn: each is done by making a left turn, even though some of the components at least will be over and done with well before the turn has been made.

But it is not so easy to evade the challenge posed by the initial segment argument. We said that on Hornsby’s view, a basic action has parts—it is not atomic. But are these parts intentional actions or not? If each part is an intentional action, then it depends on practical thought. Now, Hornsby does not claim, absurdly, that every non-basic action is preceded by an occurrent thinking of a procedural fact.⁹² So it is not the presence or absence of such a thought *process* that makes an action non-basic or basic. Furthermore, it is not the case, Hornsby thinks, that mere *possession* of a piece of procedural knowledge of the form ‘One can do *B* by doing *A*’ makes it the case that one’s doing *B* necessarily employs this knowledge, or that one is doing *B* by doing *A* (in the calculative sense of ‘by’). Rather, it is whether the action is an *exercise* of procedural or non-procedural knowledge how. But if procedural knowledge is the medium of calculative thought, as Hornsby thinks, and if calculative articulation is essential to intentional action, as Thompson and Rödl think, then it seems that behaviour that issues from non-procedural knowledge cannot be intentional action.

What, then, of the possibility that a basic action has parts, but which are not themselves intentional actions? The problem is this. The calculative unity of these parts in the action is not accounted for, and thus neither is the temporal unity of the action as a whole; the only possible source of the action’s structure lies beyond the agent’s practical thought, in which case the action is not intentional.

Return to the A—D order. The agent is doing *D* by doing *C* by doing *B* by doing *A*. All the things he is doing are unified: *A* as means to *B*, *B* (and thus *A*) as means to *C*, *C* (and thus *B* (and thus *A*)) as means to *D*. And the agent knows all this. Moreover, that the things he is doing are so unified and that he thinks them to be so unified are not two independent realities that just so happen, happily, to obtain. His thought, that they are so unified, is *knowledge*; but it could not be knowledge if they were not so unified. And his knowledge is *practical*, the cause of what it understands: the things he is doing are so unified because he thinks that they are. So the unity of the things he is doing and his knowledge of that unity come together, necessarily.

Suppose I have set the man to work doing *C* by doing *B*. He knows all right that he is doing *B*, and he knows that he is doing *B* by doing *A*. But suppose he does not know that he is doing *B* for the sake of doing *C*—I have concealed the purpose of his work from him. As far as he is concerned, he is doing *B*, not because he is doing *C*, but because he is doing *D*—carrying out orders, earning his pay. In this case, though he is doing *C* (by doing *B* by doing *A*), and though there is a purposive nexus between his doing *B* and his doing *C*, he

⁹²See e.g. Hornsby 1980b: 81.

is not doing *C* intentionally; I have kept the purposive nexus from him. The source of the “means-end” unity of doing *B* and doing *C* is external to the agent’s practical thought, alien to him. For instance, if my plan fails, and he is doing *B*, but not in such a way as to be getting *C* done, then this is my problem, not his: that he is, as far as *my* calculation is concerned, meant to be doing *C* in doing *B* can set no standard for *his* doing *B* if this is not how he conceives of the point of doing *B*.

Return now to a purportedly basic action, *A*, with parts, 1, 2, 3, that are not themselves intentional actions. The agent is not doing *A* by doing 1, 2, and 3 where the sense of ‘by’ is to introduce means. There is still a sense in which the agent is doing *A* by doing 1, 2, and 3; 1, 2, and 3 are, after all, parts of his doing *A*. But that they are parts, and that they contribute to his doing *A* in whatever ways they do, are, by hypothesis, beyond the agent’s practical thought, though he may know of them by observation, hearsay, or the memory of his training. From the agent’s practical point of view, doing *A* appears “pointlike”; the details of his movement are alien to him. He might know that he is doing *A*, but he cannot understand his doing *A*. And so his knowledge that he is doing *A* cannot be practical knowledge, because practical knowledge is the cause of what it understands, the source of the calculative, and thereby temporal, unity of the action.

Thompson argues that,

as Aristotle (for example) teaches, skill or craft or *technē* often drives out deliberation. What is done in accordance with skill in doing *B*, or in the exercise of a practical capacity to do *B*, is not, as such, determined by deliberation or reflection—unless by a peculiarity of the skill itself (which might involve measurement and calculation, say, as laying carpeting does). But the absence of reflection does not make the action thus skillfully performed, making a pot of coffee, as it might be, or raising a hand, into a sort of unanalyzable whole; egg-breaking certainly does not lose its character as an intentional action after the agent’s thirty-fourth omelet. Why should we suppose that acquisition of the type of skill that interests us, skill in moving a limb or object along this or that type of path, must deprive movement along sub-paths of their status as intentional? (2008: 108)

A skilled omelette-making is no less calculatively articulated for being skilled as one done in accordance with the steps of a recipe—one is still seasoning the eggs in order to make it taste nice. But the idea of knowing how to make an omelette that figures in the thought that making an omelette can be a basic action for somebody seems to treat omelette-making as an “unanalyzable whole”. Lucy O’Brien, for example, puts her version of Hornsby’s proposal like this: “for any agent, in a given situation, there will be a class of actions, members of which, the agent is entitled to take as actions that she knows how to execute successfully, and any process of determining what she is to do will result in a command to execute one of those actions that she is entitled to take as knowing how to execute just like that” (2007: 164). It appears that the agent’s cognitive relationship to the action—here, a command—leaves it such that any complexity in the action is external to the agent’s practical thought. It is as if the command to make an omelette is issued by the mind, whose work here is now done, and automatically “carried out” by the body, or sub-personal systems within it.

In a similar vein, Enç conjectures that “[o]nce the intention to tie one’s shoe lace is formed, no subsequent practical reasoning is required to figure out how to satisfy the intention. If the content of the intention matches an item in one’s repertoire of basic acts, and if the conditions are right, the formation of the intention becomes sufficient for its execution by the lower subsystems” (2003: 71). And he summarizes his view in the claim that “knowing how to do something, *B*, without needing to use one’s knowledge of how to do something, *A*, in order to do *B* boils down to a capacity rational agents have of getting *B* done without needing to cognitively

control how it is done” (2003: 71). This is a good example of the conception of basic action and practical knowledge that underpins this many contemporary approaches. Implicit is the idea that cognitive control could only be involved via the exercise of procedural knowledge. The control over the parts of a basic action and their unity is assigned not to practical thought, but to lower, sub-personal systems. But this precisely undermines the status of the action as intentional: the explanation of why it will be no accident if an agent in the midst of performing a basic action *A* will end up having done *A* is to be located beyond her practical thought. This puts her doing *A* on a par with her digesting a piece of meat, the temporal unity of which event is similarly explained by the characteristic function of certain sub-personal systems; one cannot save this kind of view by suggesting that “the lower behavioural systems only ‘know how’ to make good on certain commands” (Enç 2003: 71) or that one’s fingers “just get on with the job by themselves, as it were, now that I have learned to type” (Steward 2009: 301–2)—the scare quotes and ‘as it were’s give the game away. The language of *delegation* (and the like), which is so pervasive in contemporary literature on basic action, has its home in describing the interactions of multiple rational agents, not one rational agent and various sub-agential, non-rational systems.

Hornsby is sensitive to these difficulties. When she tries to explain what it is to be *able to simply do A*, to do *A directly*—without mediation by calculation, practical reasoning, procedural knowledge—she warns us against a variety of misunderstandings of the notion of *simplicity* or *directness* in play. The exercise of skill is not “something that we *simply do*” in the sense in which digesting is something that we simply do—something that lies beyond our control (2005: 126). Nor are skills themselves simple, or simple as such; many of them are very complex. And the actions and activities in which the exercise of skills consists need not be simple, as reflection on the good skier and the initial segment argument have already shown (cf. Hornsby 2007: 173). These three clarifications, however, provide little in the way of a positive account of what it is to be able to simply do *A*, such that this notion might explicate that of practical knowledge and show how an agent can “get started” acting, in the relevant sense. We have been told that the simplicity of skill and basic action consists neither in the skill, nor in its exercise, nor in the activity in which the exercise consists. But in what sense, then, are basic actions basic, and how are they supposed to stop the regress with which we began while retaining their status as intentional actions? Hornsby’s account of the phrase “able to simply do *A*” provides no positive insight, and the positive suggestions made by O’Brien and Enç seem only to undermine the proposal by encouraging the thought that basic action ends up outside the sphere of practical thought altogether, by making the parts, and thus the whole, of a basic action non-intentional. It seems we are left without a response to the initial segment argument.

We have seen that thinking that basic actions have no parts, or that their parts are not intentional, is untenable. But our only grip on what it would be for the parts to be intentional is for them to owe their articulation to calculation. Yet an action that is done by doing its parts, which are calculated to amount to doing the whole, is precisely not a basic action. What we need to understand, then, is this: *how can there be actions that are calculatively articulated*—remember, this means that the agent understands the articulation, and his understanding is its source—*which calculative articulation is not owed to calculation?*

8 The structure of practical knowledge

8.1 A parallel problem

It helps, I think, to consider what I take to be a parallel problem in thinking about perceptual experience. I'll quickly rehearse a line of thought familiar from the work of John McDowell. If the only way in which a belief could be justified were by another belief, then we would not be able to think of the world bearing rationally on our thought; for, only thought could rationally bear on thought. Thus there is a great temptation to think that it must be possible for a belief to be justified by something other than a belief; and the natural candidate—for it is the rational bearing of the world itself that we are interested in—is experience. But it is surprisingly difficult to give an account of experience on which it is fit for this purpose, and as long as this is so the very ideas of beliefs with empirical content, and believing one thing on the basis of another, remain in jeopardy. It is easy to lapse into a version of the Myth of the Given: the thought that experience provides external constraint on thought from outside the sphere of rational relations and yet somehow this constraint is supposed to be rational. The difficulty consists in conceiving of experience as somehow simultaneously *within* thought—so that its bearing can be *rational*—and yet *beyond* thought—so that it can provide *external* constraint. As is well known, McDowell's elegant solution to this puzzle is to conceive of experience as, like thought, having conceptual content; however, while conceptual capacities are actively employed in thought, they are passively drawn into operation in experience. That conceptual capacities are operative in experience itself, and not merely in the response in thought to experience, secures the rationality of the transaction between mind and world; that their operation is passive secures the external constraint on thought.

McDowell's offers a diagnosis of why such a view seems difficult; we have difficulty seeing how a natural phenomenon like sensory perception could be “anything but externally related” to conceptual activity (1996: 89). What we need, he thinks, is to relax our conception of the natural, to allow that a transaction in nature—the receipt of sense-impressions—could at the same time be a transaction in the space of reasons. And he notes, perhaps surprisingly, that his focus on problems concerning perception “was not essential” (1996: 89):

Kant says “Thoughts without content are empty, intuitions without concepts are blind”. Similarly, intentions without overt activity are idle, and movements of limbs without concepts are mere happenings, not expressions of agency. I have urged that we can accommodate the point of Kant's remark if we accept this claim: experiences are actualizations of our sentient nature in which conceptual capacities are inextricably implicated. The parallel is this: intentional bodily actions are actualizations of our active nature in which conceptual capacities are inextricably implicated. (1996: 90)

What could it be for conceptual capacities to be inextricably implicated in action—in particular, in basic action, whose basicness, McDowell agrees with Hornsby, means that one does not do “the thing in question by doing something else in a sense that brings one's means-end rationality into operation” (2007: 327)? His response is less illuminating than one might hope:

So what difference does it make, according to me, for activity to be permeated with rationality...? ... When a rational agent catches a frisbee, she is realizing a concept of a thing to do. In the case of a skilled agent [i.e. one who is performing a basic action], she does not do that by realizing other concepts of things to do. She does not realize concepts of contributory things to do, in play for her as concepts of what she is to do

by virtue of her means-end rationality in a context in which her overarching project is to catch the frisbee. But she does realize a concept of, say, catching *this*. ... When a dog catches a frisbee, he is not realizing any practical concept; in the relevant sense, he has none. The point of saying that the rational agent, unlike the dog, is realizing a concept in doing what she does is that her doing, under a specification that captures the content of the practical concept that she is realizing, comes within the scope of her practical rationality—even if only in that, if asked why she caught the frisbee, she would answer “No particular reason; I just felt like it”. (2007: 327)

The idea seems to be that basic action is “permeated with rationality”—it “inextricably implicates conceptual capacities”—insofar as it comes within the scope of the agent’s practical rationality by being subject to the reason-seeking question ‘Why?’. But though the basic action is apt to be part of a larger whole whose articulation and unity is the deed of the agent’s means-end rationality, whatever structure the basic action has is not aptly characterised in terms of the operation of that means-end rationality.

It is hard to see this picture as more than a recapitulation of our problematic. We knew from the beginning that a basic action must be capable of being subject to the question ‘Why?’, for it must be capable of answering the question ‘How?’—indeed, capable of answering it in a special, regress-stopping way. Our problem has been that we cannot see how a basic action, insofar as any structure it has is not the deed of the agent’s means-end rationality, could genuinely be an answer to the question ‘How?’, genuinely subject to the question ‘Why?’. An agent without a practical understanding of the articulation of his movement cannot be expressing practical knowledge when he says ‘I am doing A’.

The shape that McDowell’s picture takes is this. A basic action, if it is to stand a chance of doing the work that it needs to do, must be subject to the question ‘Why?’; it must fall within the scope of the agent’s means-end rationality. Yet that means-end rationality cannot enter into the constitution of the basic action itself, or else it would no longer be basic. It can merely relate a basic action to a larger action or plan. However, this starts to seem a bit as if a basic action is mythical in the way in which the Given was: it is supposed to stand in rational relations while being intrinsically nonrational. And thus to alleviate this suspicion, McDowell insists that even basic action is permeated with rationality, that it inextricably implicates concepts. But this makes it clear that the inextricable implication of concepts cannot consist in the fact that a basic action comes within the scope of the agent’s practical rationality by being subject to question ‘Why?’, for its being so subject *presupposes* the inextricable implication of concepts (this is another way of putting what is wrong with the ‘delegation-to-subpersonal-systems’ approach to basic action).

The trouble is that we only have one picture of what the implication of concepts in the practical domain looks like: their deployment in means-end rationality. The problem about perception was resolved by distinguishing two different ways in which the same (theoretical) conceptual capacities could be operative. In *Mind and World*, the idea was that we could distinguish between the active putting together of concepts into propositional form in judgment and the passively given, propositionally articulated, content of experience. In more recent work, McDowell’s thought is rather that the passively given content of experience is not propositional, but intuitional; this intuitional unity of the sensory manifold in experience is still conceptual, even though it is not discursive, because “every aspect of the content of an intuition is present in a form in which it is already suitable to be the content associated with a discursive capacity, if it is not—at least not yet—actually so associated. That is part of the force of saying, with Kant, that what gives unity to intuitions is the same function that gives unity to judgments” (2008: 264). The parallel between action and perception suggests, then, that we

need to be able to distinguish between two ways in which practical conceptual capacities could be operative in action. And the idea of a second way in which practical conceptual capacities could be operative in action might be able to do duty as the substantive account, for which we have been looking, of Hornsby's proposal that basic actions are the exercise of skill.

8.2 Skill

What we need is for basic actions to have a rational structure of the same form as actions structured by the A—D order, but where this structure is not the result of active mediation in practical thought, that is, of calculation from general to particular. The suggestion is this: the capacity for calculation, whose exercise in practical reasoning accounts for the articulation and unity of the A—D order, also accounts for a corresponding articulation and unity in basic action.⁹³ The same function—the capacity for practical thought—gives unity to both. Now, the idea that conceptual capacities are *passively* operative in action seems to be an obvious non-starter. The contrast we need to exploit is not that between activity and passivity; rather, it is that between mediated and immediate activity. The idea we need is one of the immediate realization of rational structure: the rational structure of a basic action is due not to its being worked out, but to its being *immediately instantiated*.

In calculation, the capacity for practical thought is actualized in the synthesis of different representations of action: in deriving, from an intention to do *B*, and procedural knowledge that one can do *B* by doing *A*, my intentional action (doing *B* by doing *A*, here and now), I actively bring together a representation of doing *B* as end and a representation of doing *A* as means, and my doing *A* in order to do *B* is the same reality as this synthetic representation. This is what it is for calculation to occur, for means and end to come together in practical thought. The calculative articulation of a non-basic intentional action by practical reasoning is the result of actively synthesizing practical representations.

A skilled action too must be calculatively articulated, on pain of not being an intentional action, not subject to Anscombe's question 'Why?', and thus not fit to be the means by which the agent performs a non-basic action. But though this articulation must ultimately have the same source (the same function, the same capacity for practical thought), it must be actualized in a different way. It cannot arise through active synthesis, through being derived by practical reasoning; it must lie already articulated in the skill itself, waiting to be immediately instantiated in the skilled action. A skill, then, must be a conceptually-, and thus rationally-, articulated capacity, whose exercises are conceptually—that is, calculatively—articulated intentional actions, the parts of which are themselves intentional actions, unified by practical thought in its non-procedural form.

In seeing that skill cannot be a mere capacity or disposition, but must be a *rational* and *conceptual* capacity or disposition, we are able to assuage part of the worry about the foundational role that Setiya's proposal for the epistemology of knowledge in intention engenders (§6). A full treatment of the epistemology of skill would entitle us to the claim that it constitutes a genuine form of knowledge, and explain why this knowledge consists in a rational capacity, rather than a piece of propositional knowledge that may or may not be associated with a non-rational capacity. It is progress enough, however, to have shown that the metaphysical role that basic

⁹³Here I echo McDowell: "The capacity whose exercise in judging accounts for the unity of the content of judgments—propositional unity—also accounts for a corresponding unity in the content of intuitions" (2008: 260).

action plays in the structure of intentional action can be fulfilled only if it is the exercise of a kind of general knowledge how that does not merely secure the epistemological point that it will be no accident if an agent's cognition in intention is true, but that the capacity that secures that point is a rational, conceptual capacity, thus preempting the worry that the epistemology of intentional action must bottom out in reliabilism.

Conclusion

I have argued that there is a cognition condition on intention and intentional action. If an agent is doing *B* intentionally, he has knowledge in intention that he is doing *B*, that he is doing *B* because he is doing *C* and, if it is a non-basic action, that he is doing *B* by doing *A*; his knowledge in intention that he is doing *B* is knowledge that he is going to do *B*, knowledge of success, in this sense: he knows that it will be no accident if he ends up having done *B*. A parallel condition holds, perhaps surprisingly, for an intention for the future to do *B* and knowledge in intention that one is going to do *B*. In both cases, agent's knowledge is of what is happening, or what is going to happen, in the world; it is not knowledge merely of the agent's state of mind. This knowledge is practical: it the cause of what it understands; without it what is happening, or what is going to happen, is not an intentional action, and any practical thought of the agent's about what is happening, or what is going to happen, does not hit the heights of intention. The cognition condition is grounded not in reflection about the directions of fit and guidance born by different kinds of mental states, nor in the indicative content of linguistic expressions of intention, but in the calculative and temporal structure of intentional action itself. Examples that purport to show that the content or scope of the condition must be weakened, or that the condition should be completely rejected, have been revealed to be ineffective, because they trade on confusions, remedied via reflection on the structure of intentional action, about the different kinds of failure to which intentional action may be susceptible. Cognition in intention amounts to knowledge in intention because intention and intentional action presuppose general knowledge how to do things, which comes in two kinds, procedural knowledge how and non-procedural knowledge how (skill). The calculative structure of intentional action—of doing *B* by doing *A*—and the structure of procedural knowledge—knowledge that one can do *B* by doing *A*—presuppose the possibility of basic action, and that basic action is the exercise of skill. Though reflection on the temporal structure of intentional action initially suggested that this demand could not be satisfied, we saw our way to resolving the puzzle by recognizing that skill cannot be a non-rational capacity to act, but must rather be a rational capacity that is characterized by the same general structure as practical thought more generally, albeit in a specifically different way.

References

- Anscombe, G. E. M. 1963: *Intention*. Second edition. Oxford: Blackwell. Reprinted by Harvard University Press, 2000.
- 2005: *Practical Inference*. St Andrews Studies in Philosophy and Public Affairs. Exeter: Imprint Academic.
- Austin, J. L. 1956: 'Ifs and Cans.' *Proceedings of the British Academy*, 42, 109–132. Cited as reprinted in Austin 1961, pp. ??
- 1961: *Philosophical Papers*. Oxford: Oxford University Press.
- Baier, A. 1970: 'Act and Intent'. *Journal of Philosophy*, 67(19), 648–658.

- 1971: 'The Search for Basic Actions.' *American Philosophical Quarterly*, 8(2), 161–170.
- Boyle, M. and Lavin, D. 2010: 'Goodness and Desire.' In S. Tenenbaum (ed.) *Desire, Practical Reason, and the Good*, pp. 161–201. Oxford: Oxford University Press.
- Bratman, M. 1987: *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.
- Comrie, B. 1976: *Aspect: An Introduction to the Study of Verbal Aspect and Related Problems*. Cambridge textbooks in linguistics. Cambridge: Cambridge University Press.
- Davidson, D. 1967: 'The Logical Form of Action Sentences.' In N. Rescher (ed.) *The Logic of Decision and Action*. Pittsburgh: University of Pittsburgh Press. Cited as reprinted, with Criticism, Comment, and Defence, in Davidson 1980, pp. 105–148.
- 1969: 'The Individuation of Events.' In N. Rescher (ed.) *Essays in Honor of Carl G. Hempel*, pp. 216–234. Dordrecht, Holland: D. Reidel Publishing Company. Cited as reprinted in Davidson 1980, pp. 163–180.
- 1971: 'Agency.' In R. Binkley, R. Bronaugh, and A. Marras (eds.) *Agent, Action, and Reason*. Toronto: University of Toronto Press. Cited as reprinted in Davidson 1980, pp. 43–61.
- 1973: 'Freedom to Act.' In T. Honderich (ed.) *Essays on Freedom of Action*. London: Routledge and Kegan Paul. Cited as reprinted in Davidson 1980, pp. 63–81.
- 1978: 'Intending.' In Y. Yovel (ed.) *Philosophy of History and Action*. D. Reidel and The Magnes Press, The Hebrew University. Cited as reprinted in Davidson 1980, pp. 83–102.
- 1980: *Essays on Actions and Events*. Oxford: Oxford University Press.
- Donnellan, K. 1963: 'Knowing What I am Doing.' *Journal of Philosophy*, 60(14), 401–409.
- Enç, B. 2003: *How We Act: Causes, Reasons, and Intentions*. Oxford: Oxford University Press.
- Engstrom, S. 2009: *The Form of Practical Knowledge: A Study of the Categorical Imperative*. Cambridge, MA: Harvard University Press.
- Falvey, K. 2000: 'Knowledge in Intention.' *Philosophical Studies*, 99, 21–44.
- Finkelstein, D. H. 2003: *Expression and the Inner*. Cambridge, MA: Harvard University Press.
- Ford, A. ms.: 'Varieties of Naive Action Explanation.' Unpublished manuscript.
- Galton, A. 1984: *The Logic of Aspect: An Axiomatic Approach*. Oxford: Clarendon Press.
- Gibbons, J. 2010: *Seeing What You're Doing*, vol. 3, pp. 63–85. Oxford: Oxford University Press.
- Goldman, A. I. 1970: *A Theory of Human Action*. Englewood Cliffs, NJ: Prentice Hall.
- Hampshire, S. 1965: *Freedom of the Individual*. New York: Harper & Row.
- Hornsby, J. 1980a: 'Action and Ability.' In R. Haller and W. Grassl (eds.) *Language, Logic, and Philosophy*, Proceedings of the Fourth International Wittgenstein Symposium, pp. 387–391. Vienna: Hölder-Pichler-Tempsky.
- 1980b: *Actions*. London: Routledge and Kegan Paul.
- 2005: 'Semantic Knowledge and Practical Knowledge.' *Proceedings of the Aristotelian Society, Supplementary Volume*, 79(1), 107–130.
- 2007: 'Knowledge and Abilities in Action.' In C. Kanzian and E. Runggaldier (eds.) *Cultures: Conflict-Analysis-Dialogue*, Proceedings of the 29th International Wittgenstein Symposium, pp. 165–180. Frankfurt: Ontos Verlag.
- McDowell, J. 1996: *Mind and World*. Paperback edn., with a new introduction by the author. Cambridge, MA: Harvard University Press.
- 2007: 'Response to Dreyfus.' *Inquiry*, 50(4), 366–370. Cited as reprinted in McDowell 2009, pp. 324–328.

- 2008: 'Avoiding the Myth of the Given.' In J. Lindgaard (ed.) *John McDowell: Experience, Norm, and Nature*. Oxford: Blackwell.
- 2009: *The Engaged Intellect: Philosophical Essays*. Cambridge, MA: Harvard University Press.
- McGinn, C. 1996: *The Character of Mind: An Introduction to the Philosophy of Mind*. Oxford: Oxford University Press, second edn.
- Moran, R. 2004: 'Anscombe on 'Practical Knowledge''. In J. Hyman and H. Steward (eds.) *Agency and Action*, Royal Institute of Philosophy Supplement: 55, pp. 43–68. Cambridge: Cambridge University Press.
- Moran, R. and Stone, M. J. 2009: 'Anscombe on the Expression of Intention: An Exegesis.' In C. Sandis (ed.) *New Essays on the Explanation of Action*, pp. 132–168. Houndmills: Palgrave Macmillan.
- Müller, A. W. 1979: 'How Theoretical is Practical Reason?' In C. Diamond and J. Teichman (eds.) *Intention and Intentionality: Essays in Honour of G. E. M. Anscombe*, pp. 91–108. Ithaca, NY: Cornell University Press.
- O'Brien, L. 2007: *Self-Knowing Agents*. Oxford: Oxford University Press.
- Rödl, S. 2002: 'Practice and the Unity of Action.' In G. Meggle (ed.) *Social Facts and Collective Intentionality*. Frankfurt: Hänsel-Hohenhausen.
- 2005: *Kategorien des Zeitlichen: eine Untersuchung der Formen des endlichen Verstandes*, vol. 1748 of *Suhrkamp Taschenbuch Wissenschaft*. Frankfurt am Main: Suhrkamp, 1. Aufl. edn.
- 2007: *Self-Consciousness*. Cambridge, MA: Harvard University Press.
- Ruben, D.-H. 2003: *Action and its Explanation*. Oxford: Clarendon Press.
- Ryle, G. 1949: *The Concept of Mind*. London: Hutchinson. Page numbers refer to the 1973 Penguin University Books Edition.
- Searle, J. R. 1983: *Intentionality*. Cambridge: Cambridge University Press.
- Setiya, K. 2007: *Reasons Without Rationalism*. Princeton: Princeton University Press.
- 2008: 'Practical Knowledge'. *Ethics*, 118(3).
- 2009a: 'Intention'. In E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy*. Fall 2009 edn.
- 2009b: 'Practical Knowledge Revisited'. *Ethics*, 120, 128–137.
- 2010: 'Sympathy for the Devil'. In S. Tenenbaum (ed.) *Desire, Practical Reason, and the Good*, pp. 82–110. Oxford: Oxford University Press.
- Snowdon, P. 2003: 'Knowing How and Knowing That: A Distinction Reconsidered'. *Proceedings of the Aristotelian Society*, 104(1), 1–29.
- Steward, H. 2009: 'Sub-intentional Actions and the Over-mentalization of Agency'. In C. Sandis (ed.) *New Essays on the Explanation of Action*, pp. 295–312. Houndmills: Palgrave Macmillan.
- Thompson, M. 2008: *Life and Action: Elementary Structures of Practice and Practical Thought*. Cambridge, MA: Harvard University Press.
- Velleman, J. D. 1989: *Practical Reflection*. Princeton: Princeton University Press.
- 2000: *The Possibility of Practical Reason*. Oxford: Oxford University Press.
- 2007: *Practical Reflection*. The David Hume Series on Philosophy and Cognitive Science Reissues. Stanford, CA: CSLI Publications.
- Vogler, C. 2002: *Reasonably Vicious*. Cambridge, MA: Harvard University Press.