

Two Kinds of Practical Knowledge

DISSERTATION ABSTRACT

Will Small

University of Chicago

The goal of this dissertation is to provide a unified account of two kinds of practical knowledge and the structure of intentional action. The topic of practical knowledge presents philosophical difficulties because of the modern tendency to picture cognition and the will as opposed: the will is concerned with changing the world, cognition with representing it as it already is. Thus the expression “practical knowledge” seems as if it can denote merely knowledge of a distinctive *topic* (compare: geographical knowledge). I defend the idea, however, that practical knowledge is a distinctive *form* of knowledge, or way of knowing (compare: perceptual knowledge). The aim of my dissertation, then, is to advance our understanding of what it is to know something practically. I do this by focusing on two kinds of practical knowledge and the ways in which they inform intentional action: *agent’s knowledge*—the knowledge an agent has of her intentional actions in progress and in prospect—and *knowledge how*—the general knowledge that enables an agent to perform particular actions. The investigation reveals that *intending to do something*, *doing something intentionally*, and *possessing a rational power* (that is, *bearing a skill*) are ways of knowing.

I argue, following G. E. M. Anscombe, that if an agent is doing something intentionally, she knows she is doing it. I extend this claim to intention for the future: if an agent intends to do something, she knows she is going to do it. An agent’s knowledge of her intentional actions in progress or prospect—of what she is doing intentionally when she’s doing it, and of what she’s going to do intentionally—differs, Anscombe claims, from knowledge she might have about her actions under non-intentional descriptions or about events that are befalling or will befall her. Famously, she thinks that agent’s knowledge is “non-observational”, by which she seems to have meant that it is based neither on observation nor on any sort of evidence; and this description has some intuitive plausibility:

for example, while I might discover that, in intentionally walking to school, I'm thereby enlarging the hole in my shoe, I don't discover that I'm walking to school. Likewise, with respect to future events, while someone who says "I am going to fail this exam" as a prediction bases his remark on evidence, someone uttering the same sentence as an expression of intention does not. Even though agent's knowledge does not seem to be beholden to the facts, Anscombe nevertheless thinks that it is a genuine form of knowledge.

Following Gilbert Ryle, I argue that having a certain kind of practical knowledge how to do something consists in possessing a distinctive kind of practical power. An agent's knowledge how concerning a certain task or activity differs, Ryle claims, from her theoretical or propositional knowledge concerning it. For instance, think of the difference between the knowledge one might acquire about how to drive a car with manual transmission, on the one hand, by reading a book or observing someone and, on the other, by taking lessons and practicing. Ryle argues that practical know-how cannot be reduced, and should not be assimilated, to knowledge of facts ("knowing that"); rather, it consists in practical capacities, abilities, and dispositions. Even though these kinds of powers do not consist (solely) in knowledge of facts, Ryle nevertheless thinks that they exemplify a genuine form of knowledge.

The very existence of each of these kinds of practical knowledge is often doubted. For instance, many philosophers have found incredible the idea that there could be knowledge of a worldly fact (e.g. that I am walking to school) without any basis in observation or other evidence; if there is any knowledge satisfying Anscombe's characterization, they think, it must be the knowledge one has of what one intends to do or be doing—of the content of one's intention, rather than the actuality of its execution. And, while Ryle thought knowledge how to be a species of knowledge distinct from what he called "knowing that", a number of contemporary philosophers have argued, in effect, that he drew a distinction where there is no difference. Even among those philosophers who accept the existence of one or both of these kinds of practical knowledge, there remain disagreements as to their nature, bases, and significance. I will argue that these two kinds of practical knowledge are of great significance: we need a proper understanding of both in order to make sense of intentional action.

In showing this, I will thereby show that agent's knowledge and knowledge how exhibit a certain unity. These two kinds of practical knowledge become intelligible only together, with intentional action, as it were, figuring as the middle term: moving in one direction, to understand agent's knowledge, we need to understand what it is knowledge of (intentional action in progress and in prospect), but, fundamentally, intentional action is an exercise of knowledge how; moving in the other direction, to understand knowledge how, we must appreciate the fact that it issues in intentional action, which we get in view only by recognizing that it is essentially cognized in agent's knowledge.

A unified account of this sort is somewhat unusual: though each kind of practical knowledge has attracted a fair amount of recent philosophical attention, it is notable that, for the most part, the attention has been individual. Each kind of practical knowledge is discussed by a different set of philosophers, who come at the topic with different agendas and commitments. The debate about the nature of knowledge how has taken place largely within epistemology and the philosophy of mind, while the debate about the nature of agent's knowledge has taken place largely within the philosophy of action. This kind of division of labour is of course understandable, but it has impoverished both debates: I argue that they both share a common structure, and that they have a common solution.

The common structure, very schematically, is this. One side of the debate, attempting to do justice first and foremost to the epistemic status of its target concept (be it agent's knowledge or knowledge how) as knowledge, conceives of its target as a special kind of theoretical knowledge. The other side complains that this undermines the practical status of the target, and re-conceives of the target accordingly—but in the process, they undermine or otherwise fail to do justice to its epistemic status. Neither side of the dialectic provides a satisfying account of practical knowledge, thus encouraging the suspicion that the concept is incoherent—that its two aspects are not co-satisfiable.

An examination of both instantiations of this schematic dialectical structure will reveal that, in each case, both sides of the debate share certain key assumptions, which must be rejected in order to get the relevant kind of practical knowledge in view. In order to escape this unsatisfying dialectic, we will find it necessary to discard certain deeply held assumptions concerning the nature of intentional action, knowledge, and powers, dispositions, and abilities. We will see that the contemporary debate

concerning agent's knowledge and intentional action assumes that intentional action is itself intrinsically unstructured and that practical reason informs it only from the outside; that the contemporary debate concerning knowledge how assumes that the powers, dispositions, and abilities that enable us to perform skilled actions are, similarly, intrinsically uninformed by practical reason; and that the tendency to view all forms of knowledge as kinds of theoretical knowledge is hard to dislodge. Only by rejecting these assumptions and freeing ourselves from this tendency will we be able to get the two kinds of practical knowledge, and thus intentional action properly understood, in view.

*

In Chapters 1–3, I articulate and defend the Anscombean idea that there is a *cognition condition* on intention and intentional action: an agent who is doing *A* intentionally has practical knowledge—agent's knowledge—that she is doing *A*, and an agent who intends to do *A* has practical knowledge—agent's knowledge—that she 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. Agent's 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 she is doing, but how and why she is doing it. Agent's knowledge is knowledge *in intention*, 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 she is going to do, but how and why she 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 done *A* intentionally. In the specific case of intentional action, the agent's knowledge in intention 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 she is doing, but that it will be no accident if she succeeds. Similarly with intentional action in prospect: an agent who intends to do something knows not just what she is going to do, but that it will be no accident if she gets it done.

In Chapter 1, then, I introduce the idea of a cognition condition on intention and intentional action via Anscombe's claim that if an agent is doing *A* intentionally, she knows "without observation" that she is doing *A*, and the purported counterexamples to it. I then introduce the contemporary debate in the philosophy of action that has inherited the issues bound up in Anscombe's claim. In inheriting these issues, the contemporary debate has transformed them: whereas Anscombe was arguing that intentional *action* is essentially *known*, the contemporary debate disputes whether or not *intention* involves *belief*. *Cognitivists about intention* (e.g. Velleman, Setiya) think that intention is a special kind of belief, and thus that when an agent intends to do *A*, she has a belief (of a special kind) that she is going to do *A*. Cognitivists conceive of themselves as capturing the insight contained in Anscombe's claim, while doing away with various excesses of her account. However, insofar as cognitivists conceive of practical reasoning as reasoning that concludes, or can conclude, in an intention, they thereby conceive of practical reasoning as reasoning that concludes, or can conclude, in a special type of belief. And insofar as reasoning that concludes in belief is theoretical reasoning, cognitivists think that practical reasoning is not a kind of reasoning distinct *from* theoretical reasoning, but rather

a distinctive kind of theoretical reasoning. *Non-cognitivists about intention* (e.g. Davidson, Bratman) deny that intention consists in a special kind of belief: they want to insist on the distinction between practical and theoretical reasoning, and they think that intention is a special kind of desire or “pro-attitude”. It seems, however, that they must thereby reject the idea that there is a cognition condition on intention and intentional action. The contemporary debate, then, presents us with a choice between two positions: either we accept the cognition condition, thereby placing the topic of agent’s knowledge at the centre of a philosophical account of agency, by denying the distinctively practical character of intention and practical reasoning (cognitivism), or we retain the distinctively practical character of intention and practical reasoning by rejecting the cognition condition, thereby relegating the topic of agent’s knowledge to the periphery of the philosophy of action (non-cognitivism). I end Chapter 1 by suggesting some reasons why we might both hope and believe that these positions do not exhaust the options.

In Chapter 2, I articulate my conception of the cognition condition. I show that its true ground and meaning come into focus through reflection on the calculative structure of intentional action in progress, and that 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. I conclude that the proper way to articulate the cognition condition is more demanding than we might have anticipated: if an agent is doing something intentionally, she knows *that* she is doing it, she knows *how* and *why* she *is* doing it, she knows *how to* do it, and she knows that *it will be no accident if she ends up succeeding* in having done it. Similarly, if an agent intends to do something, she knows that she is going to do it, she knows why (and potentially how) she is going to do it, she knows how to do it, and she knows that it will be no accident if she ends up having done it.

In Chapter 3, I use the rich account of the cognition condition developed in Chapter 2 to resolve the issues raised in Chapter 1. First, I show that the purported counterexamples do not get a grip on the cognition condition once it is properly understood, and why we need and ought not weaken it

in the ways that some have recently suggested. Secondly, I turn to the question of justification. The account of Chapter 2, it might seem, shows only that if an agent is doing something intentionally, she *thinks truly* that.... Thus it is necessary to evaluate the epistemological credentials of agent's *cognition*, in order to see whether it really amounts to agent's *knowledge*—and if so, in virtue of what. I do this by assessing several suggestions that have been made in the literature on agent's knowledge. It emerges that the correct account of the epistemology of agent's knowledge is latent in the account of the cognition condition and the structure of intentional action provided in Chapter 2: whereas the account in Chapter 2 emphasized the *metaphysical role* of knowledge how in constituting the temporal unity of intentional action, in Chapter 3 we come to appreciate that an agent's possession of knowledge how to do *A* gives her the *epistemic authority* to form intentions to do *A*, and thus to make knowledge claims of the form “I am doing *A*” and “I am going to do *A*”. The final section of Chapter 3 returns to the contemporary debate between cognitivism and non-cognitivism about intention, and shows why the appearance that these positions exhaust the options for thinking about the relation between intention and cognition depends on a conception of intentional action as intrinsically unstructured by practical reason.

In Chapters 4–5, I raise and resolve a problem for the account developed thus far. The account of the cognition condition developed in Chapter 2 relies heavily on the idea that intentional actions are calculatively structured: that one is doing *D* by doing *C* by doing *B*...and so on. But though many of the things we do, we do by means of doing other things, it would appear that there must be some things we can do *directly*, otherwise we could never do anything at all. The overwhelming majority of philosophers of action have concluded that, to ward off the threat of a vicious regress, there must be basic actions: deeds an agent's doing of which do not require her to take means to doing them, the basic things we *can just do*—“directly”, “immediately”, “just like that”. A parallel requirement arises when we reflect on the role knowledge how plays in justifying agent's knowledge. In order to be justified in thinking 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 that one knows how to do what one is doing (or intends to do). This demand is often satisfied by *procedural knowledge*: one knows that one can

do *C* by doing *B*. But if one is to exercise 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: there must be a way of knowing how to do something that doesn't consist in knowing by what means one can do it. Both the structure of intentional action and the structure of agent's knowledge appear to require foundations: *basic action*, and a non-procedural form of knowledge how—*skill*.

In Chapter 4, I argue that giving a coherent account of basic action is considerably more difficult than philosophers of action have realized. The tendency is to suppose that, whereas complex intentional actions issue from deliberation (which in turn makes essential use of procedural knowledge how), a basic action does not depend on deliberation—if it did, it wouldn't be basic—but instead consists in the exercise of skill. However, contemporary theorists of skill and basic action (e.g. Dreyfus, Enç, Setiya) have overstated the difference between acting from deliberation and acting from skill: the two have been represented not merely as distinct, but as profoundly disjoint. The result is an inadequate conception of basic action as alien to the conceptual order of which it is supposed to be the fundamental member. I argue that to resolve this dilemma—one horn of which is the inadequate conception of basic action, the other horn of which is the regress of doing things by means of doing other things—we must reject a seemingly plausible assumption: that if an agent can perform a basic action *A* without having to engage in deliberation about how to do *A*, then his doing of *A* will be an action that has no internal means-end structure. In order to reject this assumption, of course, we need to make sense of the possibility that there is another way, other than through deliberation, in which an intentional action might have an inner rational articulation. The suggestion, to be developed in Chapter 5, is that a basic action has an inner rational articulation because it is the exercise of skill. While it is not uncommon for philosophers of action to think that an account of intentional action must be grounded in skill, the way in which skill is commonly conceived of in fact prevents it from playing this kind of role; thus, the challenge in Chapter 5 is to articulate a conception of skill that makes it intelligible how its exercise could be a genuinely intentional action (and thus not alien to the conceptual order it is supposed to ground).

In Chapter 5, I pursue this goal first by evaluating the dominant contemporary conceptions of knowledge how: *anti-intellectualism* and *intellectualism*. Anti-intellectualists (e.g. Dreyfus, Noë) think that knowing how to do *A* consists in having a kind of practical ability to do *A*. It is typically assumed that Ryle is an anti-intellectualist, and there are obviously similarities between his view and contemporary formulations of anti-intellectualism (many of which are articulated, indeed, in order either to defend Ryle from intellectualist objections or to formulate his view more precisely so as to criticize it). But I show that the similarities are superficial. Contemporary anti-intellectualism offers, in effect, a *reductive* account of knowledge how: it seeks to reduce knowing how to do something to the possession of various abilities, dispositions, and powers where the relevant concepts *ability*, *disposition*, and *power* are presumed to be already understood. The argument of Chapter 4 already shows that such a conception of skill will not do, for it cannot but entail the inadequate conception of basic action. Ryle, however, sought to identify a *distinctive kind of power*: the power that is knowing how to do something. His account is resolutely *non-reductive*. Thus intellectualist arguments against anti-intellectualism—of which, as I show, some are convincing while some are not—do not touch a genuinely Rylean position.

Intellectualists think that knowing how to do *A* consists in knowledge of facts concerning how to do *A*. According to the well-known version of intellectualism proposed by Stanley and Williamson, one knows how to do *A* just in case (i) one knows, of a way in which one can or could do *A*, that one can or could do *A* in that way, and (ii) one can think of that way under a “practical mode of presentation”. I argue that this account fails for three reasons: first, the idea of practical mode of presentation is insufficiently explained; secondly, the conception of a way of performing an action cannot play the role that the account requires of it; and thirdly, the account of how knowledge how informs human action is inadequate—like anti-intellectualist accounts, it entails the inadequate conception of basic action. The reason that both intellectualism and anti-intellectualism cannot make proper sense of the idea that intentional action depends on the exercise of skill is that they both work within a common framework, which I call *practical dualism*. According to practical dualism, the various abilities, dispositions, and powers that we exercise in acting are not intrinsically informed by

practical rationality and cognition. Whereas intellectualists conceive of these non-cognitive powers as guided by genuinely rational and cognitive states, anti-intellectualists typically maintain that they amount to *sui generis* exemplifications of intelligence; nevertheless, anti-intellectualists still insist on a sharp break between the kinds of powers that make up a skill and fully-blown powers of cognition and practical reasoning (compare the conception of acting from deliberation and acting from skill as not merely distinct but disjoint, mentioned in connection with basic action above). In order to make progress, we must reject practical dualism by insisting that practical powers are intrinsically informed by practical reason.

By rejecting practical dualism, we are finally in a position to seriously entertain Ryle's thought, that knowing how to do something consists in a distinctive kind of practical power. Reflection on the failures of the inadequate account of basic action, and on the anti-intellectualist and intellectualist accounts of knowledge how that entail it, shows that this distinctive kind of practical power must be a *rational* power. Thus in the final section of Chapter 5, I articulate a conception of skill as a rational power that solves the problem of basic action. I find the materials for a proper understanding of the concept of a rational power in Aristotle (in *Metaphysics* Θ). Central to the account is the idea that a skill is a two-way power. I argue that in order to make sense of skill as a genuine form of knowledge, we must drop the blinkered focus on the expert's masterful exercise of his skill in skilled action that characterizes many philosophical treatments of knowledge how, and inquire into the broader condition of *bearing* a skill. It emerges that, in addition to exercising a skill by doing the activity in question, the bearer of a skill *practices* and *teaches* what she knows. Because it is only by paying attention to the broader context of bearing a skill that we can understand how a skill persists across its bearers and across time, the investigation elucidates what I call *the life-cycle of a skill*. When we recognize that the life-cycle of a skill consists in a distinctive kind of rational structure, we are finally able to make sense of the idea that a basic action has an inner rational articulation because it is the exercise of skill. We thus resolve the problem of basic action and thereby vindicate the conception of intentional action and agent's knowledge articulated in Chapters 1–3.