

Part One

Rylean Responsibilism

Die Grammatik des Wortes „wissen“ ist offenbar eng verwandt mit der Grammatik des Wortes „können“, „imstande sein“. Aber auch eng verwandt mit der des Wortes „verstehen“. (Eine Technik ,beherrschen‘.)

The grammar of the word “know” is evidently closely related to the grammar of the words “can”, “is able to”. But also closely related to that of the word “understand”. (To have ‘mastered’ a technique.)

(Wittgenstein 1953, § 150)

Chapter 1

Ryle on Know-how and Intelligence

This chapter provides the groundwork of Rylean responsibilism in the form of a normativist account of Ryle's notion of intelligence and of his characterization of know-how as an intelligent ability.

In § 1.1, I introduce the concept of know-how in terms of the phenomenon it is supposed to explain. I argue that that this phenomenon, what Ryle calls "intelligent practice", can be understood as *normative* practice. This notion of normative practice is then further developed and clarified in § 1.2.

§ 1.3 is devoted to the idea that the concept of know-how is a concept of achieving success in activities, or of meeting the normative demands of those activities. This, in turn, requires the concept of an ability and in particular or a reliable ability within a certain range of normal situations. § 1.4 will clarify these issues.

In § 1.5, I present a sketch of the full conception of know-how as a capacity to achieve success in virtue of an understanding of an activity's normative requirements. As I will show in § 1.6, the crucial distinctions and explanatory demands involved here are entirely parallel to those in the debate about rule-following. Finally, § 1.7 discusses the loose boundaries of the concept of know-how – its vagueness and its context-dependence.

§ 1.1 Intelligence and Normative Practice

Any philosophical inquiry into the nature of a concept must start with an intuitive characterization of the explanandum. Otherwise, it is difficult to express the very question to be addressed. I shall therefore start with what I take to be a pre-theoretic and uncontroversial survey of the intuitive homeland of the concept of know-how. In fact, Gilbert Ryle has already provided such a survey. He offered a rich range of examples of know-how:

Ryle's Range of Cases¹

| | |
|--|--|
| playing chess | (Ryle 1945a, 5–6, 14; Ryle 1949, 29–30, 40–41) |
| arguing, reasoning, inferring | (Ryle 1945a, 6–7, 12–13; Ryle 1949, 29, 46) |
| speaking a language | (Ryle 1945a, 14; Ryle 1949, 29, 41, 50, 297) |
| behaving ethically | (Ryle 1945a, 12–13; Ryle 1949, 297) |
| practicing philosophy or science | (Ryle 1945a, 15) |
| writing with literary style | (Ryle 1945a, 14–15) |
| making jokes | (Ryle 1945a, 8, 10; Ryle 1949, 29–30) |
| playing cards | (Ryle 1949, 29) |
| writing poetry, composing limericks | (Ryle 1949, 48, 49) |
| feinting to do something | (Ryle 1949, 50) |
| calculating, practicing mathematics | (Ryle 1945a, 15; Ryle 1949, 34–35, 125) |
| solving anagrams | (Ryle 1949, 49, 125) |
| spelling a word | (Ryle 1949, 125) |
| mountaineering | (Ryle 1949, 42) |
| cooking | (Ryle 1945a, 8, 12–14) |
| playing cricket | (Ryle 1949, 125) |
| reading a map | (Ryle 1945a, 15) |
| driving a car | (Ryle 1945a, 10) |
| designing dresses | (Ryle 1945a, 8) |
| fishing | (Ryle 1945a, 12; Ryle 1949, 29) |
| practicing as a medical doctor, performing surgery | (Ryle 1949, 48, 125) |
| swimming | (Ryle 1949, 48, 293–294) |
| conducting battles, working as a general | (Ryle 1945a, 8, 14–15) |
| acting as a clown | (Ryle 1949, 33) |
| examining something | (Ryle 1945a, 15) |
| ‘seeing’ or appreciating jokes | (Ryle 1945a, 10; Ryle 1949, 29) |
| playing golf | (Ryle 1949, 78) |
| behaving appropriately at a funeral | (Ryle 1945a, 8) |
| sculpting | (Ryle 1949, 50) |
| persuading a jury | (Ryle 1945a, 8) |
| shooting, hitting the bull’s eye | (Ryle 1945a, 15; Ryle 1949, 45) |
| pruning trees | (Ryle 1949, 29) |
| boxing | (Ryle 1949, 48) |
| tieing knots | (Ryle 1949, 50, 54–55, 125) |
| trading professionally | (Ryle 1949, 48) |

I take it that this list of activities may serve as a criterion of adequacy for any philosophical explication of the concept of know-how. Whatever know-how may be, it must be possible to understand that these activities are typical examples for things one may know how to do.

¹ I have done my best to be able to claim that this list is, to my knowledge, complete – if not in the references, then at least in the examples listed.

For the sake of convenience, I will sometimes use the simple abbreviation “S knows how to A”, where “S” stands for a subject and “A” for an activity.

However, one may wonder why I have chosen to speak of ‘activities’ rather than, say, of ‘actions’. And with good reason. It is an obvious fact that all or nearly all of the examples can and typically are exercised as intentional actions – and as I shall discuss in chapter 3, this is not an accident. The reason for my choice of words is, however, rather simple. Activities, as I shall use this notion, include intentional actions, but also further elements. For example, some stretches of behavior, particularly very long processes, are not always naturally described in terms of individual actions, whereas the notion of an activity seems more natural.² Most importantly however, one sometimes exercises one’s know-how not as an intentional action, but entirely automatically. Still, such performances are nevertheless genuine exercises of the know-how in question. At this point, I cannot do more than bluntly state this claim, but I will elaborate, discuss and defend this in chapter 3. Until then, however, nothing I say depends on this point. Thus, in the earlier parts of this book, when I say that an ‘act’ or ‘performance’ is an exercise of know-how, this can simply be read as a notational variant of ‘action’. From chapter 3 onwards, however, it will become crucial to have the broader notion of acts, performances and activities.

One may also worry that the notion of an activity is unclear – either because it is not clear how we should understand it in general or because it is not clear where we should locate the boundary between variants of one activity and distinct activities. For now, I will assume that the concept of an activity is sufficiently clear in order to be put to work in my proposal. I shall come back to these questions later in this section and in § 1.7.

Thus, the concept of know-how is at home in explaining and understanding the performances of certain activities. Ryle calls this “intelligent practice” (1949, 27). He points out that people’s performances can exhibit a variety of “qualities of mind” (1949, 26), some of which in turn are picked out by what he calls “intelligence-concepts” (1945a, 3). These concepts form a proper subset of the “mental-conduct epithets” or “mental-conduct concepts” (1949, 9, 26) with which the whole of *The Concept of Mind* is concerned (cf. 1949, 9–11). For example, one may do something wisely, thoughtfully, correctly, successfully, as well as inefficiently, stupidly, and so forth. According to Ryle, these concepts “belong to the family of concepts ordinarily surnamed ‘intelligence.’” (1949, 26)

² One may couch this point in terms of the general ontological view that actions are a subclass of *events*, whereas activities are a subclass of *processes*. For an illuminating discussion of these issues, see Hornsby (2012). But I shall leave this open.

Thus, the touchstone for a philosophical explication of know-how is its role in explaining and understanding ‘intelligent practice’ – and this, Ryle proposes, through the examination of the family of ‘intelligence concepts’. But how are we to understand this family of concepts? Paul Snowdon, a dedicated critic of Ryle’s view, has complained:

Ryle does not say or give any hint as to the boundaries of this range of concepts. [...] The scope of intellectualism is, then, more or less impossible to determine. (Snowdon 2011, 62)

But this criticism does not go far enough. Ryle employs the family of intelligence concept in order to demarcate the explanandum for which the concept of know-how is then introduced as the explanans. Thus, not only do *intellectualist* accounts of know-how depend on a precise understanding of intelligence concepts, *every* account of know-how does.

As a first step in order to answer to the question what Ryle’s “intelligence concepts” have in common, it is important to note that he does not only include concepts which refer to what one would ordinarily conceive of as the *presence* of intelligence under the term “intelligence concept”. Cases of *lack or deficiency* in intelligence – such as stupidity – are also included. Ryle makes this perfectly clear in both of his central texts on know-how:

Consider, first, our use of the various intelligence-predicates, namely, “wise,” “logical,” “sensible,” “prudent,” “cunning,” “skilful,” “scrupulous,” “tasteful,” “witty,” etc., with their converses “unwise,” “illogical,” “silly,” “stupid,” “dull,” “unscrupulous,” “without taste,” “humourless,” etc. (Ryle 1945a, 5)

The mental-conduct concepts that I choose to examine first belong to the family of concepts ordinarily surnamed ‘intelligence’. Here are a few more determinative adjectives of this family: ‘clever’, ‘sensible’, ‘careful’, ‘methodical’, ‘inventive’, ‘prudent’, ‘acute’, ‘logical’, ‘witty’, ‘observant’, ‘critical’, ‘experimental’, ‘quick-witted’, ‘cunning’, ‘wise’, ‘judicious’, and ‘scrupulous’. When a person is deficient in intelligence he is described as ‘stupid’ or else by more determinate epithets such as ‘dull’, ‘silly’, ‘careless’, ‘unmethodical’, ‘uninventive’, ‘rash’, ‘dense’, ‘illogical’, ‘humourless’, ‘unobservant’, ‘uncritical’, ‘unexperimental’, ‘slow’, ‘simple’, ‘unwise’, and ‘injudicious’. (Ryle 1949, 26)

In a survey of the debate on know-how, John Bengson and Marc Moffet (2011b) have thus been inspired to an interesting orthographical strategy:

Hereafter, we reserve ‘intelligence’ (lowercase ‘i’) for intelligence in the narrow sense, namely, that which is intelligent but *not* stupid, idiotic, and so forth; we will use ‘Intelligence’ (capital ‘I’) as an umbrella term covering all states of intellect and character, including intelligence (in the narrow sense), stupidity, and idiocy. (Bengson & Moffet 2011b, 5–6)

It is certainly true that a clarification of Ryle's terminology is important. It seems to me, however, that simply distinguishing "Intelligence" from "intelligence" does not do much to clarify any of these terms and that it misses an important opportunity: There already *is* an umbrella term other than capital-'I'-Intelligence which is well-established and which helps to bring out the *point* of the various concepts subsumed under the label of capital-'I'-Intelligence. What I have in mind is the concept of *normativity*. However, it will take some time to spell out how this can be made to work.

To begin with, I take it to be obvious that all the terms cited by Ryle contain *evaluations*.³ These concepts do not describe performances in a detached or uninvolving manner but contain judgments of these performances as to whether or not – and to what extent – they live up to certain standards. In particular, they assess the quality of these performances against the background of the normative requirements of the activity in question.⁴

The norms or standards which govern our activities are manifold, of course. Efficiency, success, and originality are paradigmatic examples of the norms Ryle subsumes under the label "intelligence". Many of them are highly general and their application differs a lot from case to case. One may also conceive of these norms differently and deny that, say, efficiency in playing chess has much to do with efficiency in catching fish. According to such a view, there would be no such thing as *the norm of efficiency*, but rather several norms for the efficient practice of a certain activity (or family of activities). But for now, these differences do not make a difference. However many standards are in play at a certain point, and however narrowly or broadly they are to be conceived, it remains true that all of them are recognizably norms or standards against the background of which an activity is evaluated.

On this basis, I therefore propose to understand Ryle's notion of "intelligent practice" (1949, 27) – in the wider, capital-'I'-sense pointed out by Bengson & Moffett (2011b) – as *normative practice*. What we are interested

³ Next to the passages quoted above, Ryle mentions further examples in four other passages, and all of these examples clearly contain evaluations. He says that "we characterise either practical or theoretical activities as clever, wise, prudent, skilful, etc." (1945a, 1) and that "thinking and doing do share lots of predicates, such as 'clever,' 'stupid,' 'careful,' 'strenuous,' 'attentive,' etc." (1945a, 2) Ryle examines "[a]dverbs expressing intelligence-concepts (such as 'shrewdly,' 'wittily,' 'methodically,' 'scrupulously,' etc.)" (1945a, 3) And he writes: "When a person knows how to do things of a certain sort, we call him 'acute,' 'shrewd,' 'scrupulous,' 'ingenious,' 'discerning,' 'inventive,' 'an expert cook,' 'a good general,' or 'a good examiner,' etc." (1945a, 14)

⁴ The example "strenuous" (Ryle 1945a, 2) comes closest to a counterexample to this view. But the contrast between acting strenuously and acting with ease can be understood as the contrast between failing and succeeding at meeting the norm of efficiency.

in when we use the concept of know-how is behavior engaging in activities which in turn are governed by norms. In fact, these norms are also able to explain the otherwise puzzling co-existence of positive and negative ‘intelligence-predicates’ in Bengson & Moffett’s general category. The hallmark of the positive evaluations which refer to the phenomenon of genuine intelligence is that the norms of the activity in question are in fact met. In contrast, negative evaluations refer to the fact that the performances under discussion fail to meet these norms.

I take it that this proposal is something Ryle should have applauded because he makes very clear how important the normative dimension of know-how is, even if he does not use this term himself. He writes:

When a person knows how to do things of a certain sort (e.g., cook omelettes, design dresses or persuade juries), his performance is in some way governed by principles, rules, canons, standards or criteria. (For most purposes it does not matter which we say.) (Ryle 1945a, 8)

What is involved in our descriptions of people as knowing how to make and appreciate jokes, to talk grammatically, to play chess, to fish, or to argue? Part of what is meant is that, when they perform these operations, they tend to perform them well, i.e. correctly or efficiently or successfully. Their performances come up to certain standards, or satisfy criteria. (Ryle 1949, 29)

Ryle is certainly correct in assuming that terminological differences between ‘standards’, ‘criteria’, and related terms do not make a difference for the present questions. However, all of these expressions clearly refer to normative concepts. Thus, it should be obvious that Ryle explicitly endorses the idea that know-how is closely connected to the normativity of activities.

§ 1.2 Clarifying the Normativity of Activities

I have argued that the normativity of practice is at the heart of the concept of know-how since it is at the heart of the specific phenomenon for which this concept is supposed to account – what Ryle called ‘intelligence’. However, Ryle does not pick up these phenomena systematically. He does not explicate in which sense activities are governed by norms and in which sense one ought to conform to these norms. Therefore, at least some amount of clarification is in order.

In this section, I offer three such clarifications. First, I argue that the normativity of practice suggests a way to individuate activities. Second, I discuss the bindingness of the norms of an activity. And third, I comment on the question of their nature and their ontological status.

As for the first point, I hold that we can individuate each activity uniquely by giving the unique set of norms which governs it as well as the way these norms are weighed with respect to each other. Certainly, the individual norms which are members of this set can be of different importance for the activity in question. Some norms may form the core of the set, while others are merely at the periphery, and there may even be several gravitational centres more or less balancing one another.

Consider a swimmer. What is the activity she is engaged in? At a first level of description, we may say that she is swimming, period. At a second level of description, she is swimming the breast stroke. The norms and standards in play are clearly different in these different perspectives. One may be swimming very well, but not thereby swimming the breast stroke at all. The converse, however, does not hold. Swimming the breast stroke well is one of several ways of swimming well. A third level of description would be one where we describe her as swimming a race. Again, the norms and standards in play change. One may be swimming very well, but not thereby doing well at swimming a race. Again, the converse is false since everybody who does well at swimming a race also does well at swimming. Further, some performances will count both as swimming the breast stroke and as swimming a race, others will count as only one of them and some will count as neither while still being exercises of swimming.

These phenomena can be accounted for in terms of the hypothesis that activities are to be individuated in terms of weighed sets of norms – those which govern the activity. The set of norms which identifies the activity of swimming is a proper subset of the set of norms which constitutes the activity of swimming the breast stroke and again a proper subset of the set of norms which constitutes the activity of swimming a race. And these latter two sets intersect, but do not coincide.

On these grounds, I shall go on employ the view that activities are identified by the weighed sets of norms which govern them – particularly in § 2.2 and § 4.5. Thus, part of the full justification for this claim will hopefully stem from the way in which it proves useful at these later stages.⁵

This brings me to my second point, the question how binding the norms of a certain activity are. Crucially, to say that the activities one may know how to engage in are governed by norms is not to say that one ought to engage in these activities in a moral or an otherwise non-instrumental sense

⁵ One may object that the identity criteria of the individual norms are unclear and that therefore the identity criteria of an activity in terms of a set of such norms are equally unclear. This, however, is not an argument against my view but something I entirely endorse. As I shall argue in § 1.7, know-how is a vague concept anyway.

of ‘ought’. The fact that some activity is governed by certain norms simply means that *if* one’s performances are evaluated as exercises of *that* activity, it is *these* norms against the background of which the assessment is to be conducted. In other words, *if* the question is how good a performance is *as* an instance of A-ing, *then* what one *ought* to do is such-and-such. But it is an entirely open question whether or not one ought to act in that way all things considered.

For example, one may assess a policewoman’s shot at another person in many different ways. According to the norm not to harm others, it will count as morally bad. According to the norm to prevent a serial killer from continuing to commit murders, it may count as excusable, acceptable or even heroic, depending on one’s moral principles. But regardless of these assessments, it is clear that one can *also* assess the policewoman’s shot as an exercise of her marksmanship. Then, one can judge, for example, the speed, efficiency and dexterity of her performance. And clearly, a person’s marksmanship and her moral judgment or character can be assessed quite independently.

However, this does not mean that such things cannot interfere at all. There may very well be cases where excellence at a certain activity is not morally neutral because the acquisition of the know-how in question goes hand in hand with a process of stupefying of one’s capacity for moral emotion and ethical judgment. Torturers come to mind, and some may argue that marksmanship should be regarded in the same way. Further, maybe possessing or exercising certain know-how already *is* a moral achievement. Think of activities like rescuing people from drowning in the sea or from burning to death in a car or building which is on fire. But such interconnections are clearly only a local phenomenon. The vast majority of know-how concerns activities where the mere possession or non-possession of the relevant know-how is morally neutral.⁶

My third and final point of clarification concerns the ontological status or the norms of an activity. While the general philosophical question of norms and normativity is indeed very important, I venture to be very brief in the context of this book. My proposal is intended as entirely neutral with respect to this question, and I contend that an account of know-how *should* be neutral in these respects.

⁶ I cannot comment here in more detail on the the precise demarcation of those cases where having know-how itself is morally good or bad. This issue clearly seems to be connected to the question to which extent the relevant practices are automatized or habitualized, to which extent they are second nature. I will discuss this general issue at length in chapter 3, particularly in § 3.1.

To begin with, there are myriads of different kinds of activities which one may know how to engage in (cf. *Ryle's Range of Cases* on page 14) – just as there are myriads of different facts of which one may have propositional knowledge. And just like the debate about propositional knowledge tries to abstract away from the particularities of the facts one may know, the debate about know-how also tries to abstract away from the particularities of the activities one may know how to engage in. And part of what we abstract away from here is precisely the nature and the metaphysical status of these activities and of the norms which govern them.

For example, with some of these activities, there is a sense in which what it takes to succeed is ‘objectively out there’. Intuitively clear cases along these lines involve hunting or in foraging for food. In other cases, the relevant norms seem entirely socially constructed and in this sense ‘subjective’, for example when it comes to dressing well for a certain social occasion or behaving politely with respect to a specific group of people. Probably, there is a continuum of cases along the axis of objectivity and mind-independence versus subjectivity, social constructedness and mind-dependence.

Other questions concern the question how to account for the norms themselves and for our relation to them. Some explanations begin with an independent ontology of norms and spell out how they have a grip on us. Other explanations start with our normative attitudes and spell out how these come to be calibrated and institutionalized in such a way that we can finally speak of there *being* norms for us to follow.⁷ Probably, different norms and activities will have different statuses.

An account of the concept of know-how should be able to encompass all of these intuitive assessments and all philosophical accounts of how normativity in general and the normativity of activities in particular is to be classified, explained or otherwise elucidated. This is why the account offered here will be neutral in all of these respects.

§ 1.3 Living Up to What it Takes

So far, I have argued that the concept of know-how explains the engagement of people in normative activities. For this general phenomenon of the normativity of practices, Bengson & Moffett “use ‘Intelligence’ (capital ‘I’)

⁷ Ryle seems to suggest briefly that this is his preferred view (cf. Ryle 1945a, 4). Norms as independently existing abstract objects would probably appear to him just as mysterious and problematic as propositions (cf. Ryle 1930). However, his more positive suggestions on the nature of the mind in general and the intellect in particular make these worries much less problematic. I shall discuss this in § 2.5.

as an umbrella term" (2011b, 5). Now, I shall go on to explicate intelligence proper, what they call "intelligence (lowercase 'i')". In this section, I will discuss a necessary condition of know-how, achieving success at the activity in question. Then, in § 1.5, I complete the first step of my Rylean account of know-how by explicating know-how as a capacity to succeed in virtue of an understanding of the norms which govern the activity in question.

The two passages just quoted on page 18 contain an important further insight: Ryle points out that the concept of know-how does not only signify that the activity in question *has* a normative dimension, but also that the person who possesses know-how *lives up to* these normative requirements. To repeat:

What is involved in our descriptions of people as knowing how to make and appreciate jokes, to talk grammatically, to play chess, to fish, or to argue? Part of what is meant is that, when they perform these operations, they tend to perform them well, i.e. correctly or efficiently or successfully. Their performances come up to certain standards, or satisfy criteria. (Ryle 1949, 29)

Thus, to know how to do something is to know how – to have an ability or capacity – to do it *well*. That is, to know how to engage in an activity involves the characterization of one's performances as *good* performances of that activity – as cases of achievement or success. Ryle expresses this idea as follows:

We often say of a person, or of a performing animal, that he can do something in the sense that he can do it correctly or well. [...] Many of the performance-verbs with which we describe people and, sometimes with qualms, animals, signify the occurrence not just of actions but of suitable or correct actions. They signify achievements. Verbs like 'spell', 'catch', 'solve', 'find', 'win', 'cure', 'score', 'deceive', 'persuade', 'arrive', and countless others, signify not merely that a performance has been gone through, but also that something has been brought off by the agent going through it. They are verbs of success. (Ryle 1949, 125)

'Know' is a capacity verb, and a capacity verb of that special sort that the person described can bring things off, or get things right. (Ryle 1949, 128)

What does such a success come down to? I think that this question can be answered only in very general terms. To achieve success at an activity, is to meet the norms which govern that activity. *Knowing how to A is knowing how to A well in the sense that one's performances meet the norms of A-ing.*

This can be seen most clearly on the basis of Ryle's own explication. He is concerned with performing "well, i.e. correctly or efficiently or successfully" (Ryle 1949, 29) – that is, with the norms of correctness, efficiency and success. However, we should also note some further cases in point,

the norms which underlie what Ryle calls “a few more determinative adjectives of th[e] family [of intelligence-concepts]: ‘clever’, ‘sensible’, ‘careful’, ‘methodical’, ‘inventive’, ‘prudent’, ‘acute’, ‘logical’, ‘witty’, ‘observant’, ‘critical’, ‘experimental’, ‘quick-witted’, ‘cunning’, ‘wise’, ‘judicious’, and ‘scrupulous’.” (Ryle 1949, 26) We may thus place wisdom, inventiveness, originality, and wits alongside correctness, efficiency, and success.

For the present point, it is not important how many norms are involved, how they are related to each other and whether or not they may even be united in a single scale of proficiency. What matters is that each activity one may know how to perform is governed by norms which determine which candidate performance counts as a good performance of that activity, which does not, and to which degree this is the case. This, I contend, is what Ryle has in mind when he speaks of performing “well, i.e. correctly or efficiently or successfully” (Ryle 1949, 29). By using the abbreviation “i.e.”, Ryle makes clear that performing well *encompasses* the “more determinative adjectives” which follow (Ryle 1949, 26) and that performing well is *exemplified* by things like correctness, efficiency, and success.

I have argued that knowing how to engage in an activity is knowing how to live up to the norms of that activity – however exactly these are characterized. Given that know-how is thus connected to achieving success in the activity in question, it is important to see how strong this connection is. Given that I know how to read pretty well, it is not merely true that I *can* read. Rather, my know-how is an *ability* or a *capacity* to do so:

When a person is described by one or other of the intelligence-epithets such as “shrewd” or “silly”, “prudent” or “imprudent”, the description imputes to him not the knowledge, or ignorance, of this or that truth, but the ability, or inability, to do certain things. (Ryle 1949, 27)

More precisely, Ryle holds that know-how is a special kind of ability, a *skill* or a *competence*. While the terms ‘skill’, ‘competence’ and ‘know-how’ may evoke different associations, one being ‘a form of knowledge’ and the others being something else, I share Ryle’s contention that they denote the very same concept. Thus, somebody knows how to A just in case she has the skill to A and just in case she has the competence to A. *Know-how, competence and skill are one.*

A detailed defense of this conceptual policy will be offered in chapter 7, where I also give an explanation of the undeniable fact that it is sometimes entirely acceptable to use the expression ‘knows how to’ in cases without ability and therefore without skill or competence. Further, I will discuss the alleged counterexamples against this claim at length in chapter 5.

While Ryle never explicitly asserts that skill is know-how and know-how is competence, it is easy to find him switching back and forth between characterizing something as know-how and as a skill. For example, on the question “[w]hen a person knows how to do things of a certain sort”, he considers what happens when “[a] good experimentalist exercises his skill” (Ryle 1945a, 8).⁸ I contend that whenever Ryle uses the expression ‘skill’ at all, it is possible to replace this expression with ‘know-how’ and vice versa, without altering the content and coherence of what he says. And the same holds for ‘competence’.⁹

This has also been stressed by others who have wondered about Ryle’s point in choosing the expressions he chooses:

[R]eferences to know-how are not singled out as especially significant from the large array of other action-related terms that Ryle uses to characterize his position and to accompany the notion of intelligent behavior: ‘competence’, ‘capacity’, ‘skill’, ‘habit’, ‘bent’, ‘disposition’, ‘performance’, ‘practice’, ‘procedure’, ‘operation’, ‘task’, and ‘exercise’ all appear. ‘knowledge-how’ is just one convenient label out of many that Ryle took to point in the same general direction. It seems that he saw in nature a certain distinctive sort of mental prowess, and he took the various words he selected to be good ways of identifying that for his readers. He was not particularly invested in whether any particular locution, e.g. ‘know-how’, did the job perfectly. (Glick 2011, 428–429)

Ryle’s *occasional* use of ‘knowledge-how’ (and related forms) is philosophically idiomatic; it is merely an abbreviation for the sort of prowess or expertise, whatever it ultimately is, in virtue of which one performs intelligently. This usage is natural, for it is common in English to speak of the things that one does intelligently as what one knows how to do. (Sax 2010, 512)

Thus, the conceptual policy of treating know-how, skill, and competence as one is well-founded. This is true for the interpretation of Ryle’s texts, and it will be developed and defended systematically over the course of this book. Like know-how, skill and competence also refer to intelligent abilities – to abilities to live up to what it takes to do well in a certain activity.

⁸ This passage is quoted fully on page 52, where I discuss the general point Ryle tries to make on that occasion.

⁹ Obviously, I cannot discuss this for *all* of Ryle’s uses of these terms. To allow an evaluation of the textual basis of my claim, I shall at least list the pages where Ryle uses ‘skill’, ‘skilled’, etc. (Ryle 1945a, 1, 4–5, 7–8, 10, 13, 15; Ryle 1949, 28–29, 32–34, 40, 42, 45, 48–49, 58, 60, 78, 80–81, 134, 139, 145, 191, 260, 279, 294–295, 312–313) and where he uses ‘competence’, ‘competent’, etc. (Ryle 1945a, 1; Ryle 1949, 17–18, 31, 48–49, 54–56, 70–71, 77, 111–112, 129, 131, 194, 308, 312, 316).

§ 1.4 Reliable Ability and Normal Situations

If know-how is an ability, as just shown in § 1.3, the question arises what *kind* of ability it is. I will come back to this question § 1.5. In the present section, however, I would first like to discuss the concept of ability in general.

First, it is important to stress that ability is opposed to mere possibility.¹⁰ To take an example from the contemporary literature on know-how:

Lottery (Stanley & Williamson 2001, 414–415)

[I]f Hannah wins a fair lottery, she still does not know how to win the lottery, since it was by sheer chance that she did so.

Clearly, I, like Hannah, *can* win a fair lottery because it is *possible* for me to win it, but I do *not know how* to win the lottery because I do not have the *ability* to do so. But why? Is it really problematic to say that I *do* have the ability to win the lottery simply because I am *able* to win the lottery, and that I am able to win the lottery simply because I *can* win the lottery? Of course, it is perfectly fine to talk like this. But if ‘able to’ is used as equivalent with ‘can’, then we should draw a distinction between being able to do something and having the ability to do so.

In order to see this, let me come back to Ryle again. He writes:

Now successes are sometimes due to luck; a cricketer may score a boundary by making a careless stroke. But when we say of a person that he can bring off things of a certain sort, such as solve anagrams or cure sciatica, we mean that he can be relied on to succeed reasonably often even without the aid of luck. He knows how to bring it off in normal situations. (Ryle 1949, 125)

Thus, knowing how to A is knowing how to meet the norms of A-ing, in Ryle’s words, “reasonably often” – that is, *reliably*. A genuine ability which suffices for know-how always requires *sufficient* reliability, unlike my being able to win the lottery where I am not reliable at all. A physician who regularly fails to cure an illness will, other things equal, not count as knowing how to cure that illness, even if he was successful in a small number of cases. Normally, an overwhelming number of failures shows that a small number of successes were just lucky coincidences. Knowing how to A requires *sufficient reliability* at living up to the norms of A-ing.

¹⁰ I deny that every possibility is an ability, but I accept that ability a *form* of possibility. There has been prominent criticism of the idea that abilities can be understood in terms of possibilities, for example by Anthony Kenny (1975). But this criticism only brings out that possibility cannot be sufficient for ability – something I endorse – and it does not falsify my view that possibility is necessary for ability (cf. e.g. Maier 2010, sect. 4.2; Carr 1979).

This point is nicely illustrated by the following pair of examples which Katherine Hawley has brought into the contemporary debate:¹¹

Avalanche (Hawley 2003, 27)

[C]onsider Sally, out for an ill-advised winter walk in the hills. She has no idea what to do in the event of an avalanche. When an avalanche occurs, she mistakes the snow for water, makes swimming motions, and—luckily—escapes the avalanche, since in fact the way to escape an avalanche is to make swimming motions. [...] [H]er two failings (not knowing how to escape from an avalanche and mistaking snow for water) cancel each other out. Despite her success, she does not know how to escape avalanches.

Lucky Cake (Hawley 2003, 27)

[C]onsider Shelley, who attempts to make a cake by throwing together the ingredients she discovers on opening the kitchen cabinet. Luckily for her, the cupboard contains flour, sugar, butter, and eggs, and she makes a passable cake. Of course, she doesn't know how to make a cake under normal circumstances, which include selecting the ingredients.

Ryle notes that whenever somebody in a particular instance does something well and meets the norms of the activity in question, it is still an open question whether or not she exercised a reliable ability to do so:

In judging that someone's performance is or is not intelligent, we have [...] in a certain manner to look beyond the performance itself. For there is no particular overt or inner performance which could not have been accidentally or 'mechanically' executed [...]. (Ryle 1949, 44)

Ryle exemplifies this point as follows: "We observe, for example, a soldier scoring a bull's eye. Was it luck or was it skill?" (Ryle 1949, 45) The fact that the soldier met the norms of shooting on target extremely well does not settle the question whether he is a skilled marksman who actually *knows* how to score a bull's eye or whether he merely had beginner's luck and actually does not know how to score a bull's eye after all. Making a good shot was a *possibility* for him, but maybe this was *merely* a possibility. It is an open question whether or not he exercised a reliable ability to do so.

A second example is this: "A drunkard at the chessboard makes a move which upsets his opponent's plan of campaign." (Ryle 1949, 45) Again, the fact that the drunkard's move meets the norm of disturbing his opponent's plans very well does not settle the question whether he is a skilled chess-player who actually *knows* how to upset his opponents' plans or whether he

¹¹ Hawley also mentions a third case, which I shall quote as *Annoying Smoker* on page 163. But I will discuss this example later because it differs in important ways from *Avalanche* and from Shelley's lucky baking.

merely luckily moved a piece of wood, has no skill at playing chess at all and does not know how to disturb his opponents' plans at all. Making a good move was a *possibility*, but maybe this was a *mere* possibility. It is an open question whether he exercised a reliable ability.

These distinctions and examples suggest that mere possibility does not suffice for know-how. Instead, reliable ability is a necessary condition for know-how. In this sense, the concept of ability connects possibility, which could otherwise very well be *mere* possibility, with full-blown know-how. But it is easy to misunderstand abilities: The sense of ability which helps build the bridge to know-how is not the sense of merely being able to do something which comes too close to mere possibility. Instead, it is the sense of having the ability to do something in the sense in which ability already involves some degree of reliability.

These degrees of reliability at living up to norms are part of the reason why know-how is *itself* a gradual concept. Know-how is gradable:

[I]t is proper and normal to speak of a person knowing in part how to do something, i.e. of his having a particular capacity in a limited degree. An ordinary chess-player knows the game pretty well but a champion knows it better [...]. (Ryle 1949, 58)

Thus, one person may know how to play chess better than a second person, while both can be credited with the know-how to play chess in general. And a natural reason for this would be that both are *sufficiently* reliable in satisfying the norms which govern chess, while the first person meets them more reliably than the second one. Sufficient reliability marks at least part of the threshold between knowing how to A, but not being very good at it, on the one hand, and not knowing how to A at all. But it is not always clear where exactly this threshold lies and whether it applies to all cases. In § 1.7, I shall come back to this issue.

However, the insight that reliable ability is necessary for know-how must be further refined. As Ryle is fully aware, somebody might regularly fail at exercising an activity well, but still possess the knowledge how to do so. Suppose that the physician considered earlier regularly fails to cure the illness in question, but that this can be traced back to extreme stress, a dizzying illness the physician has herself, insufficient supply of the medicine needed, constant distraction by ongoing fighting around her or other factors along these lines. It may well be that such a physician *does* know how to cure the illness, despite the fact that she regularly fails to do so. This is because nobody who knows how to A is reliable in *all* situations. Know-how only requires reliability when things are *normal*. As Ryle puts it, a skilled person "knows how to bring it off in normal situations." (1949, 125)

Ryle does not give us any clue as to which situations are normal situations. But I do not think he or anyone *can* say much about this problem. I contend that the only way to characterize normal as opposed to abnormal situations is purely negative and entirely derivative of the activity under consideration. We can say: *A situation is normal with regard to A-ing just to the extent in which nothing impedes A-ing substantially.*¹²

Of course, something can be impeded in weaker and stronger ways. But this does not show that this characterization is inadequate. These degrees of the severity of the impediments of A-ing mirror exactly the degrees of the normality of the situations for which A-ing is considered. We may thus rephrase the above statement by saying that the more A-ing is substantially impeded, the less normal the situation is with regard to this activity. To illustrate, consider an archer trying to shoot at a target. One way in which the exercise of her competence to hit the target may be impeded is by strong wind. Equivalently, one way in which the situation may be abnormal with regard to the activity of shooting at her target is in there being strong wind. The strength of the impediment mirrors the distance from the normal case. The stronger the wind, the less normal the situation for archery.

I conclude that know-how involves ability which involves sufficient reliability at meeting the norms of an activity, while this reliability is measured against the background of normal situations in which success is required.

To end this section, I would like to point out that what I have expressed here in terms of reliable abilities and normal situations can often also be found in terms such as ‘affordance’, ‘option’ and ‘specific’ vs. ‘general ability’ or, equivalently, ‘narrow’ vs. ‘wide ability’.

For example, one can translate back and forth between the latter distinction od ‘specific’ vs. ‘general ability’ (cf. Maier 2010; Clarke 2015) and the vocabulary I have been using so far. On the one hand, to have the *general* ability to A is to have the reliable ability to A in situations which are sufficiently normal with regard to A-ing, *independently* of the question whether the current situation is sufficiently normal or not. On the other hand, to have the *specific* ability to A is to be able to A in the *current* situation – paradigmatically, because one has the general ability to A and because the current situation is sufficiently normal with regard to A-ing. Another way to say that somebody has the specific ability to A is to say

¹² This view has already been defended more generally. For example, Justin Fisher has argued that this is true of all dispositions (cf. Fisher 2013). And Sebastian Rödl has given an elaborate argument for this idea with respect to all powers and capacities, tracing it back to Aristotle and Thomas Aquinas (cf. Rödl 2010, 144). However, I bracket the question if know-how is itself a disposition until § 5.6.

that A-ing is an *option* for her here and now (cf. Maier 2013). In this sense, as I am writing these sentences, I have the general ability to play volleyball fairly well, but lacking a ball, I do not have the specific ability to do so – playing volleyball is not one of my current options.

There is also a third way of saying that somebody has the specific ability to A or, equivalently, that she has A-ing as an *option* here and now. Namely, one may say that A-ing is one of the *affordances* which presents itself to her here and now. The term ‘affordance’ has been made famous by James Gibson (1977; 1979)¹³ and has since been employed and discussed in psychology and philosophy (cf. e.g. Chemero 2003; Michaels 2003). On a standard view, “[a]ffordances are [...] understood as an organism’s possibilities for action in some situation.” (Rietveld 2008, 976) Thus, suitable features of the environment – those which Ryle labeled ‘normal situations’ – and suitable general abilities come together to *create* options or affordances to exercise the respective ability:

The *affordances* of the environment are what it *offers* the animal, what it *provides* or *furnishes*, either for good or ill. The verb *to afford* is found in the dictionary, but the noun *affordance* is not. I have made it up. I mean by it something that refers both to the environment and the animal in a way that no existing term does. (Gibson 1979, 127)

Affordances, then, are relations between the abilities of organisms and features of the environment. Affordances, that is, have this structure: Affords- ϕ (feature, ability). (Chemero 2003, 189)

After this quick survey of the interconnections between related debates and concepts, I shall now return to the development of my Rylean conception of know-how as an intelligent ability. When I continue to use the term ‘ability’, this should be understood as ‘general ability’ in the sense distinguished above. And when I occasionally speak of ‘options’, ‘affordances’, ‘normal situations’ or ‘specific abilities’, I refer to the fact that it is possible for a (general) ability to be exercised in the given situation.

§ 1.5 Intelligence as Normative Guidance

So far, I have argued that Ryle makes two important points about the relation between know-how and normative practice. First, § 1.1 and § 1.2 have shown that the activities one may know how to perform always involve a

¹³ This concept has some important predecessors, in Gestalt psychology, however, namely what Kurt Lewin (1935) has called ‘Aufforderungscharakter’ and what Kurt Koffka (1935) has called ‘demand-character’, often alternatively translated as ‘valence’.

dimension of normative adequacy. They are governed by norms which determine the quality of performances as exercises of that activity. Second, § 1.3 and § 1.4 have argued that knowing how to do something means knowing how to perform well in the sense that one's performances *satisfy* those norms with sufficient reliability.

On this basis, however, Ryle makes an important *third* point¹⁴ – and this is the insight missed by those who take Ryle to be an anti-intellectualist in the sense discussed in the Introduction.¹⁵ In a crucial passage the beginning of which has already been quoted on pages 18 and 22, he writes:

What is involved in our descriptions of people as knowing how to make and appreciate jokes, to talk grammatically, to play chess, to fish, or to argue? Part of what is meant is that, when they perform these operations, they tend to perform them well, i.e. correctly or efficiently or successfully. Their performances come up to certain standards, or satisfy criteria. But this is not enough. [...] To be intelligent is not merely to satisfy criteria, but to apply them; to regulate one's actions and not merely to be well-regulated. A person's performance is described as careful or skilful, if [...] [h]e applies criteria in performing critically, that is, in trying to get things right. (Ryle 1949, 29)

Ryle's talk of 'trying to get things right' resonates with my earlier explication of intelligence in terms of normativity. In this vein, the third point in this passage can be understood as follows. To know how to do something not only means to be capable to engage in an activity which is governed by norms, and to reliably meet those norms in one's performances. It requires further that one's performances are *guided* by the norms of the activity in question. This is Ryle's point when he writes that to exercise a competence is "not merely to satisfy criteria, but to apply them; to regulate one's actions and not merely to be well-regulated" and that it means "trying to get things right" (Ryle 1949, 29). Elsewhere, and as I already quoted on page 18, Ryle makes the same point, stating that exercises of a competence are "in some way governed by principles, rules, canons, standards or criteria." (Ryle 1945a, 8) And again, "to perform intelligently is to apply criteria in the conduct of the performance itself" (Ryle 1949, 40).

¹⁴ The same point has recently been discussed by Ernest Sosa in his 2015 book *Judgment and Agency*. A crucial passage from this book is this: "Competences are a special case of dispositions, that in which the host is disposed to succeed when she tries, or that in which the host seats a relevant skill, and is in the proper shape and situation, such that she tries in close enough worlds, and in the close enough worlds where she tries, she reliably enough succeeds. But this must be so in the right way." (Sosa 2015, 23). However, I will have to leave a detailed discussion of the similarities and differences between Sosa's and my views for another occasion.

¹⁵ Notable exceptions who see a distinction along these lines include Fantl (2008, 455) and Hornsby (2011, 82).

All of these formulations – ‘applying criteria’, ‘trying to get things right’, ‘being governed by standards’ – are ways to make the same point. They refer to the meeting of normative demands *by* striving for them (cf. Kremer 2016, 8) – in a word, to *normative guidance*.

This notion of normative guidance is crucial. Obviously, being guided by the norms governing an activity requires a specific relation to those norms. Ryle, however, is notoriously vague as to what this relation exactly is. Still, there is one passage in which he explicitly mentions something which can account for this relation:¹⁶

Understanding is a part of knowing *how*. (Ryle 1949, 53)

In other words, those who know how to A must have a *grip* on the norms of A-ing. They must have an *understanding* of the norms governing the activity in question, an understanding *in virtue of which* they eventually meet those norms. Thus, exercising one’s knowledge how to A means, first, *having* an understanding of what it takes to A well and, second, *acting on* such an understanding – in Ryle’s terms: ‘applying criteria’, ‘trying to get things right’, ‘being governed by standards’. Again equivalently, know-how requires a *conception* of the relevant activity and the *guidance by* this conception in one’s practical conduct.

These considerations provide the final step in explaining Ryle’s concept of “intelligence”. This *is* indeed an umbrella term. But it does not concern the mere *presence* of certain norms which govern the activity in question – as Bengson & Moffett (2011b) would have it for capital-‘I’-Intelligence. And it does not refer to the mere *satisfaction* of those norms – as they would have it for lowercase-‘i’-intelligence. What makes a performance intelligent in the full sense of the term is that it meets the norms of the activity in question *in virtue of its being guided by an understanding of those norms*.¹⁷

¹⁶ It should be noted that Ryle in fact *merely* mentions this notion of understanding and that this sentence is taken from a different, and more specific, context of understanding what it is that one, or somebody else, is doing. Still, this suggestion will turn out to be the key to spelling out more fully what Ryle otherwise merely sketches – the very *intelligence* of know-how. Also, I am not alone in giving pride of place to the concept of understanding and to this remark by Ryle. Katherine Hawley discusses this point with respect to ‘internalist’ conceptions of know-how (Hawley 2003, 28). And most importantly, Bengson & Moffett (2007; 2011c) even identify know-how with a suitable state of understanding. I will comment on this view, objectivist intellectualism, at various later points in this book, especially in chapters 8 and 9.

¹⁷ I have recently seen that Benjamin Elzinga has just made a very similar proposal. He writes: “For Ryle, S knows-how to Φ iff (1) S is able to reliably live up to the normative standards governing Φ -ing, and (2) S Φ ’s responsibly.” (Elzinga 2016, 2) While I cannot spell out all the parallels and the few remaining differences here in detail, I take the Rylean *responsibilist* account of know-how presented in Part One of this book to be very much in Elzinga’s spirit.

This view fits the description recently suggested by Adam Carter and Duncan Pritchard. They stress the fact that know-how is a genuine cognitive achievement and propose that, in know-how,

the ability and success components can't merely be 'accidentally satisfied' conjunctively; instead, we stipulate that the success element be *because of* the ability element. [...] What results, then, is a cognitive-achievement account of knowledge-how that is a viable alternative to both intellectualist views (given that knowledge-how and cognitive achievement come apart in both directions), *and* to neo-Rylean anti-intellectualists accounts according to which knowledge-how is just a matter of possessing certain abilities. (Carter & Pritchard 2015, 194–195)

I have argued that know-how is an intelligent ability – an ability to meet the norms governing an activity in virtue of being guided by an understanding of those norms. Bringing out the full force of the idea of normative guidance, and filling in Carter & Pritchard's sketch with a full account of know-how,¹⁸ will occupy the lion's share of Part One of this book. In the remainder of this section, I will strengthen the intuitive support for this line of thought by distinguishing what I want to call a *mere ability* from genuine know-how.

What is a *mere ability*?¹⁹ While genuine know-how is an ability to do something well *because* one does it intelligently, *mere ability* is an ability to do something well, but *not* because one does it intelligently. In other words, genuine know-how involves conforming to the norms of an activity in virtue of being guided by an understanding of them. But mere ability *only* involves conforming to those norms and nothing further. Let me develop and spell out this difference with the aid of some of Ryle's ideas and examples.

In the passage quoted on page 30, I have omitted two sentences which provide an example of mere ability as opposed to genuine know-how:

Clocks & Seals (Ryle 1949, 29)

The well-regulated clock keeps good time and the well-drilled circus seal performs its tricks flawlessly, yet we do not call them 'intelligent'. We reserve this title for the persons responsible for their performances.

Thus, Ryle holds that a clock keeping the time correctly or a circus seal performing artistic tricks do not have an understanding of the norms of time-keeping and of animal artistics, respectively. Therefore, their performances cannot be understood as intelligent, and therefore they cannot be counted as knowing how to keep the time or as knowing how to perform artistic

¹⁸ A further recent proposal to account for know-how as a cognitive achievement stems from Ben Kotzee (2016), who explicitly draws on virtue.

¹⁹ This way of expressing the distinction has already been proposed in very early discussions of Ryle's view (cf. e.g. Brown 1970, 218; Brett 1974, 295).

tricks, respectively. They do these things *well*, but they do not do them *intelligently*. This seems to be uncontroversially correct as far as clocks are concerned. But non-human animals are a more complicated topic on which I shall comment briefly in § 6.7. In any case, Ryle's insight can also be motivated and applied within the realm of human persons:

The ability to give by rote the correct solutions of multiplication problems differs in certain important respects from the ability to solve them by calculating. (Ryle 1949, 42)

These two abilities are introduced in order to illustrate the distinction between what Ryle calls “pure or blind habits” on the one hand and “intelligent capacities” on the other hand (cf. Ryle 1945a, 15; Ryle 1949, 42). This is exactly the distinction between *mere* reliable conformity to norms on the one hand and reliable conformity in virtue of normative guidance on the other hand – between mere reliable ability and genuine know-how.

Being able to give correct solutions of multiplication problems does not by itself qualify as possessing the know-how to multiply. For if one gives these correct solutions by sheer rote, then one does not possess an understanding of multiplication and one does not give the correct answers in virtue of being guided by this understanding. Instead, one merely parrots what one has learned by heart. This is also an achievement, but a very different one – one of remembering things correctly and of speaking a language.

Another case in point is swimming. It is sometimes said that newborn human babies ‘know how to swim’, but that they lose this ability over time and that, later on, they have to learn how to swim anew. From what I understand, it is true that human newborns possess a set of reflexes which involve holding one’s breath when their face is underwater and making rudimentary swimming motions when a certain portion of their body is immersed in water. Let us suppose that triggering these reflexes actually constitutes a way for newborns to swim. However, I also take it to be clear that this is not an instance of know-how or competence. It is not only that grown-ups who have learned how to swim and do possess this competence are *better* at swimming than newborn babies. This difference does not suffice to make the distinction between competence and mere ability. Rather, the babies’ ability is a *mere* ability because it consists of nothing but a pattern of reflexes. It does not involve what competence requires and what grown-up swimmers have, an understanding of swimming in virtue of which they swim.

The notion of intelligence in the sense of normative guidance is indeed crucial for demarcating cases of know-how from other capacities. It is instructive to consider some further intuitive verdicts here. For example:

Digestion (Stanley & Williamson 2001, 414)

Digesting food is not the sort of action that one knows how to do.

This is certainly true. Despite the fact that I *can* digest food, and that I do so *reliably*, it would be wrong to think that I *know how* to do so. But why?

There are clearly norms which govern the digestion of food and which an effective digestive system meets reliably. One does not have to be a reductive naturalist about normativity to appreciate the fact that digestion can work well, efficiently or badly or anything in between, and that this fact stems from the biological function of keeping the organism nourished. Of course, the norm that an organism is to be kept well-nourished and the way in which digestion plays a role in meeting this norm are very different from such norms as the ones governing language use and other things. But this is beside the point. The reason why digesting food is not something one may know how to do must lie beyond the normativity of this phenomenon – however basic and undemanding that normativity may be.

And it does. In order to digest food efficiently, one does not rely on an understanding of the norms of sustenance and nourishment. It is not that one digests food efficiently *in virtue of* one's understanding of these norms. Rather, this understanding is entirely independent from the digesting of food. Some people may come to understand the workings and the biological functions of one's digestion very well. But even they do not digest food *in virtue of* this understanding.

This example may appear problematic since digesting food is certainly something which we do not 'do' in the sense of intentional action. But as I shall argue in chapter 3, the distinction between intentional actions and other performances is not coextensive with the distinction between exercises of know-how and other acts. To foreshadow, digestion is something which we never do *intentionally*, and therefore it is something which we cannot *know* how to do. By contrast, things which we sometimes do intentionally and sometimes entirely automatically can perfectly well happen in virtue of guidance by an understanding, even in the completely automatic cases.

But even bracketing these questions for the moment, it is easy to see that even some intentional actions are not exercises of know-how. I take it that the clearest cases of this are intentional *basic* actions.²⁰

²⁰ This has also been noted by several other philosophers. Alva Noë mentions the example of breathing while sleeping (cf. Noë 2005, 279 fn. 2, crediting David Chalmers). Paul Snowdon considers blinking with one's left eyelid (cf. Snowdon 2011, 76). And as David Carr notes quite generally, a description of such very simple and basic abilities as skills or as cases of know-how is "odd", it "transgresses good sense" (Carr 1981a, 53). Further proponents of this view include Katzoff (1984, 61–69), Snowdon (2003, 12) and Stanley & Williamson (2001, 440).

Many actions are performed by doing something else. Basic actions, however, are just those actions where this is false. I take it that there are very clear cases of basic actions. For example, opening or closing my eyes or my mouth are things which one may do and have the ability to do. But it would be mistaken to think that I know how to perform these acts, that I have the competence to perform them. I have the *ability* to perform them, but this is a *mere ability*.²¹

These phenomena are perfectly in line with the account I have given so far. In fact, there is even a sense in which my proposal *predicts* them. Given that know-how is a capacity to engage in a normative activity, I have said that such norms determine which acts amount to *good* performances of that activity, and to which degree this is the case. But the relation of an act's *counting* as a performance of an activity is *just the inverse* of the relation of that activity's being performed *by* performing that act. In other words, that my moving a piece of wood on a board *counts* as making a smart move in our current game of chess is just another way of saying that I *make* a smart move in our current game of chess *by* moving a piece of wood on a board. Given that know-how is only concerned with normative practices, it follows that it cannot be concerned with basic actions.

Of course, one may argue over the question whether an allegedly basic action is indeed basic.²² But I contend that, *however* one is inclined to draw the line between basic and non-basic actions, one will *thereby* also draw a line between what one can actually know how to do and what one cannot know how but merely be able to do.²³ I return to this point in § 1.7.

§ 1.6 Normative Guidance and Rule-Following

I have argued that the concept of know-how as an intelligent ability can be elucidated in terms of the difference between such cases of know-how on the one hand and cases of mere ability on the other hand. To conclude this section, I would now like to follow up on Ryle's explicitly declared terminological liberalism. As I have already quoted on page 18, he states:

²¹ Kieran Setiya (2012) holds that I *do* know how to perform basic actions because I am able to perform them *intentionally*. But he basically agrees with the view I have offered when he says that "knowing how to perform a basic action is being disposed to act on the relevant intention when one has it" (Setiya 2012, 296). After all, such a disposition is clearly compatible with mere ability.

²² I think that this is what lies behind David Carr's remark that it may be acceptable to describe basic actions as exercises of competences "in exceptional circumstances largely irrelevant to present concerns" (Carr 1981a, 53).

²³ One might object that the notion of a basic action is too unclear, or even that there are no truly basic actions (cf. Lavin 2013). But I will have to bracket this here.

When a person knows how to do things of a certain sort (e.g., cook omelettes, design dresses or persuade juries), his performance is in some way governed by principles, rules, canons, standards or criteria. (For most purposes it does not matter which we say.) (Ryle 1945a, 8)

In § 1.1, I have suggested that we should use the general notion of norms and normativity in order to describe the crucial phenomenon of intelligence. Now, I would like to point out that it is equally possible to use 'rule' as the central term and tell exactly the same story. If we do so, we can see how two important philosophical debates turn out to deal with essentially the same problem – first, the debate about know-how and, second, the debate about rule-following and the normativity of meaning which stems from Ludwig Wittgenstein's *Philosophical Investigations* (1953) and from the influential interpretation of part of this work in Saul Kripke's *Wittgenstein on Rules and Private Language* (1982).²⁴

This connection can be brought into view by the relevant paradigm examples which include mastery of a natural language, of the use of a certain linguistic term, of a mathematical method or of a certain term in a formal language such as the expression '+'. It is not an accident that these examples are also discussed as paradigm cases of activities one may know how to engage in. Ryle mentions them explicitly – cf. *Ryle's Range of Cases* on page 14 – and many other philosophers do so, as well.

A further obvious parallel between the debates under consideration is that philosophers discussing rule-following typically make a distinction which coincides with the distinction between having a mere ability and having a competence – the distinction between merely conforming to a rule and actually following it in the sense of being *guided* by it. It is a common cause of confusion when talk of 'following' a rule is not sufficiently clarified with regard to these two senses. Further, when it comes to the concept of a rule, what is typically meant is exactly what I have pointed out above – the fact that there is a distinction between doing something correctly or incorrectly, better or worse, and so forth. Instead of 'rules' which determine these assessments, we can also speak of norms, criteria, and standards.

Thus, Ryle is correct to point out that terminology does not matter. To illustrate, consider a quotation from a classic discussion of rule-following:

I intend no particular theoretical implications by talking of rules here. The topic is that there is such a thing as the correct and incorrect application of a term [...]. I shall talk indifferently of there being correctness and incorrectness, of words

²⁴ For a brief overview of these complex debates, see Miller & Wright (2002) and Glüer & Wikforss (2009).

being rule-governed, and of their obeying principles of application. Whatever this is, it is the fact that distinguishes the production of a term from mere noise, and turns utterance into assertion—into the making of judgment. (Blackburn 1984, 281–282)

Furthermore, the debate about the so-called normativity of meaning which has spawned from the debate about rule-following can be understood along the same lines. For the normativity of the meanings of linguistic expressions stems from the normativity of the performances of employing them. Again, this can be illustrated with a quotation from a seminal texts on this problem:

The normativity of meaning turns out to be, in other words, simply a new name for the familiar fact that, regardless of whether one thinks of meaning in truth-theoretic or assertion-theoretic terms, meaningful expressions possess conditions of *correct use*. (On the one construal, correctness consists in *true* use, on the other, in *warranted* use.) (Boghossian 1989, 513)

Finally, Ryle's distinction between know-how and propositional knowledge and his argument that the former cannot be reduced to the latter – his famous regress argument against propositionalist intellectualism which I will discuss in chapter 9 – also have analogues in the debate about rule-following. Arguably, the very same considerations underlie Wittgenstein's equally famous argument for the conclusion:

[T]here is a way of grasping a rule which is *not* an interpretation, but which, from case to case of application, is exhibited in what we call ‘following the rule’ and ‘going against it’. (Wittgenstein 1953, § 201)

I contend that the Rylean point that *know-how* cannot be defined as the *propositional knowledge* of the norms governing the activity in question, but must rather be understood as the competence to follow these norms, is analogous to the Wittgensteinian insistence that grasping the rules of a language in the sense exhibited in *following* them cannot be defined as grasping these rules in the sense of *interpreting* them *as such-and-such*, but must rather be understood as having “mastered a technique”, as a “practice” (Wittgenstein 1953, §§ 199, 202). In fact, one of Wittgenstein's favorite ways of expressing this kind of grasp of rules, “Jetzt weiß ich weiter”, is standardly translated as “Now I know how to go on” (Wittgenstein 1953, § 180).²⁵ And Ryle explicitly declares:

Knowing a rule is knowing how. It is realized in performances which conform to the rule [...]. (Ryle 1945a, 7)

²⁵ I will comment on the question of the German translations of ‘knows how to’ in § 7.3.

In fact, the problem of know-how on the one hand and the problem of rule-following and the normativity of meaning on the other hand are not only closely connected. It is even possible to support the stronger claim that the latter debate is a special case of the former. To see this, it is important to realize that the debate about rule-following and the normativity of meaning concerns the status of linguistic and conceptual competences and norms. But speaking a language or using a concept is only one of many examples of activities one may know how to do. In this sense, the debate about know-how can be seen as a generalization of the debate about rule-following and the normativity of meaning. Given how often philosophers discussing rule-following appeal to analogies between speaking a language and activities like playing chess, this conclusion should not be surprising.

I have followed Ryle in disregarding subtle terminological differences between rules, norms, standards and the like, as long as the point of these terms is sufficiently clear. Especially when it comes to the term 'rule', some philosophers seem to agree with much of the content of my view, but to disagree with using this term. One example among many is Erik Rietveld. On the one hand, we share an undemanding notion of normativity:

The notion of normativity implied here is a very basic one; it is revealed when we distinguish better from worse, correct from incorrect, optimal from suboptimal, or adequate from inadequate in the context of a specific situation. (Rietveld 2008, 974)

On the other hand, Rietveld takes the term 'rule' to be inappropriate in characterizing this notion. Discussing the example of the normative assessment of a work of craft, he writes:

Appreciation of the object (in context) by the craftsman is normative without being (explicitly or implicitly) guided by rules. (Rietveld 2008, 979)

Part of the explanation of this disagreement may lie in a disagreement about the phenomenon of entirely unreflective and automatic exercises of skills:

With respect to such an episode of unreflective action no additional story about operations at the mental level is needed (for example in terms of having a propositional representation of a goal, following explicit or implicit rules, or reflecting consciously or even unconsciously). (Rietveld 2008, 993)

This shows beautifully how Rietveld – I should repeat, alongside many other philosophers – takes the notion of a rule to be entangled with the idea of propositional representations and of the conscious application of these propositions. This view belongs to what Ryle forcefully rejects under the

label ‘the intellectualist legend’. But unlike Rietveld, Ryle has no qualms about using the term ‘rule’ himself. As just quoted on page 37, Ryle even goes so far as to equate knowledge of rules and know-how (Ryle 1945a, 7).

Accordingly, I maintain that it remains to be unproblematic to use the term ‘rules’ for the norms governing an activity the understanding of and guidance by is the defining feature of know-how. Still, the systematic problem underlying Rietveld’s terminological choices certainly remains to be addressed. How should we conceive of entirely automatic and unreflective performances and how can we understand them to be guided by norms or rules? I shall provide an account of these problems in chapter 3.

In sum, I contend that Ryle’s and Wittgenstein’s concerns are indeed analogous.²⁶ While I cannot discuss the intricacies of this analogy between know-how and rule-following in detail in this book, I find it nevertheless important to point out that I take my considerations to also have consequences for the problem of rule-following and vice versa, for better or worse. In any case, I hope that this analogy helps to bring out Ryle’s crucial notion of intelligence as normative guidance which will continue to be at the core of the explanatory project of this book, particularly in chapter 4.

§ 1.7 Clarifying Loose Boundaries

In this section, I would like to elaborate on the Rylean view laid out so far and comment on the loose boundaries of the concept of know-how. I shall argue that ordinary ascriptions of know-how are both vague and context-dependent. While this view ties in very well with Ryle’s stated position, he does not discuss it explicitly. I shall therefore develop these considerations independently and only occasionally come back to Ryle’s texts.

To begin with, it is important to realize that a performance will be evaluated very differently depending on the activity as an exercise of which it is interpreted. To see this, let me start with an example.

Park Chess

Consider myself walking through the park and stumbling upon a boy and a girl who move chess pieces on a chessboard. I stop and watch them for a while, noticing the unusual pattern in which the pieces are set. After observing the girl’s next move, I criticize her for moving her rook diagonally and point out that this is forbidden. They laugh at me and tell me that they are not engaged in a game of chess, at all. They are playing a game of draughts on the chessboard,

²⁶ Given these obvious interconnections, it is surprising that the current debate about know-how is conducted largely in disregard of the debate about rule-following – some notable exceptions notwithstanding (cf. e.g. Tanney 2009).

using chess pieces instead of the regular flat draughts pieces. Obviously, the girl's performance is impermissible as a performance of playing chess, but perfectly fine as a performance of playing draughts. And as it turns out, it is not only a permissible move, but a very smart one, too. Accordingly, I cannot appeal to the norms of chess in order to assess their performances, but I must switch to the norms of draughts instead. Of course, I might consider whether or not the children should, as a matter of good taste, continue to use chess pieces when playing a game other than chess. But I sensibly decide not to and walk away.

In such a case, performances are wrongly assessed as exercises of a given activity. While I falsely take the children to play chess, what they actually play is draughts. Such a misconception can not only occur from the perspective of a spectator, but also from that of a practitioner. In the case just imagined, the children might also have reacted to my criticism quite differently. Maybe they were indeed playing chess and had a perfect grasp of all of its rules except the one rule which states which moves are permitted with a rook. After my suggestion, they correct their own conception of the norms which govern their activity and continue to play correctly.

These examples raise some further questions which are very important for the idea that such activities are governed by norms.

First, it is plausible to distinguish what I would like to call minimal and evaluative norms or rules.²⁷ In the case of chess, this is the distinction between the rules of the game which determine the permissible moves on the one hand and the strategic norms which state which moves in which situations are good ways to win the game against a given opponent on the other hand. Along these lines, one may wish to distinguish between playing chess *at all* – being successfully guided by the *minimal* norms of chess – and playing chess *well* – being successfully guided by the minimal *and* at least some of the *evaluative* norms of chess.

But it is not entirely clear how strictly this distinction can be maintained. In the variation of *Park Chess* I gave above, the children were clearly playing chess, even if there was a minimal norm – the rule for the permitted moves of the rook – which they did not understand and apply. Accordingly, even if there is such a distinction, it is at least not clearly true that meeting all of the minimal norms is necessary for engaging in an activity.

An alternative application of the distinction between minimal and evaluative norms would be the idea that, never mind the preconditions of merely engaging in an activity, it is a necessary condition for having a *competence* to

²⁷ This is obviously related to John Searle's famous distinction between constitutive and regulative rules (cf. e.g. Searle 1969; Searle 1995). But I will leave these interconnections to Searle's distinction and the surrounding questions implicit here.

understand and be guided by *all* of the minimal norms involved in the activity in question. Along these lines, the children in my example would count as engaging in playing chess, but not as knowing how to play chess. This view is more plausible, but nevertheless entirely optional. Alternatively, it is equally plausible to describe the children as having the competence to play chess but as not being very good at it.

This issue has already come up in § 1.4. There, I said that there is a threshold between knowing how to do something, but being rather bad at it, on the one hand, and not knowing how to do something at all. It is easy to imagine a continuum of cases of people playing chess with a complete ignoramus at the one end and an aspiring grandmaster at the other end. Where exactly should we draw the boundary we are looking for?

Obviously, this problem brings out the fact that the concept of know-how is vague. The considerable sophistication of the debate about vagueness notwithstanding, I think it is safe to follow one of the classic characterizations of vagueness, which I take to be the common sense view, namely:

To say that an expression is vague [is] to say that there are cases (actual and possible) in which one just does not know whether to apply the expression or to withhold it, and one's not knowing is not due to ignorance of the facts. (Grice 1989, 177)

The vagueness of know-how can be understood in terms of the question what exactly the norms of a given activity are, and how they are weighed against each other. As I argued in § 1.2, this is how activities can be individuated.

What I shall try to show now is that this idea can make sense not only of the vagueness of the concept of know-how, but thereby also of its context-sensitivity. We can cash out these phenomena in terms of the vagueness and context-sensitivity of the boundaries of the weighed set of norms in terms of which the activity in question can be understood. This is especially helpful given that what one could also conceive of as distinct formulations or interpretations of the same norm can also be distinguished as several specific norms. Ascriptions of know-how are vague and context-sensitive because the relevant set of weighted norms which is taken to be governing the performances in question is vague and context-dependent.

Consider again the continuum of cases of chess players with a complete ignoramus at the one end and an aspiring grandmaster at the other end. Depending on where one is inclined to draw the line between the minimal and the evaluative norms of chess, certain cases will not count as involving skills at all, others as borderline cases and still others as meeting these requirements minimally. Is the threshold at making a mistake every third

move, at making three mistakes for every five moves, or somewhere else? What degree of reliability in meeting these minimal norms one should take to be required here is not clear, but vague.

But this phenomenon is not limited to minimal norms. If we want to evaluate who is currently the best chess player, there are several criteria which one may plausibly apply and several ways in which one may weigh them – swiftness of winning, elegance of play, maintenance of momentum, and so forth. And even if we resort to the seeming simplicitly of sheer statistics, there are not always clear lines to be drawn. If one player has won five out of five matches against other grandmasters and a second player has won seven out of eight, it is not clear at all whether the first should be judged to be the better one since she has won 100% as opposed to 87.5% of her matches or whether the second should be judged to be better than the first since she has won seven as opposed to five matches.

I have argued that we can make sense of the vagueness and context-sensitivity of know-how in terms of the set of norms which constitute the activity in question, together with their relative weight. What I have said so far may seem to be concerned only with vagueness as opposed to context-sensitivity. But I take it to be clear that at least one form of context-sensitivity can easily be identified in the picture proposed so far: Given that there is no precise criterion for the membership of a norm in the set of norms constituting an activity and for its exact relative weight with respect to the other norms in that set, there is a certain amount of variation in the norms one may reasonably assume a performance to be governed by – that is, in the activities one may reasonably assume a performance to be engaging in. Depending on the perspective of the person assessing the performance and depending on the context in which the assessment takes place, one may reach differing, but equally reasonable verdicts.

For example, consider again two chess players and suppose that what they are engaging in cannot only be described as the activity of playing chess, but also in various more specific ways – just like the swimmers already discussed on page 19. Suppose that the chess players are also experimenting with a certain opening strategy, focusing on winning the game with the force of the rooks alone, and so forth. Any given sequence of performances of moving chess pieces can be understood as exercises of many of such finely-grained activities. These finely-grained activities are constituted by intersecting, but different sets of norms. Accordingly, when assessed as an exercise of one such activity, a given performance will count as particularly good, and when assessed as an exercise of another such activity, it may count as very bad and maybe not even a proper exercise of that activity.

in the first place. Of course, the assessment of a person's performances is guided by the principle of charity and it is important to focus on those activities in the context of which a performance makes sense. But even in these boundaries, there is a considerable amount of reasonable variation.

A further example for the context-sensitivity of know-how stems from Katherine Hawley. After spelling out the way in somebody's knowledge how to drive is assessed with regard to a limited range of situations – what I followed Ryle in calling “normal” situations in § 1.4 –, she writes:

Driving Contexts (Hawley 2003)

There is, however, no unique task or range of tasks which is always involved when we ask whether someone knows how to drive. Rather, different tasks are salient in different conversational contexts. For example, in a UK context, it would be reasonable to infer from Sarah's knowing how to drive that she knows how to drive a manual, stick shift car. In most US contexts, however, this would not be a reasonable inference. (Hawley 2003, 21)

Again, there are several equally reasonable ways to conceive of Sarah's competence. One may say that she knows how to drive, period, and bracket the question whether this is true in virtue of her skill to drive a manual shift or an automatic shift car. Or one may say that these are actually two different competences and demand a clarification of which skill is under discussion.

I have argued that many performances can be understood as exercises of a variety of different activities and that these activities can in turn be understood as constituted by the set of norms which govern their exercise, as well as by the relative weight of these norms. Further, I have characterized the vagueness and context-sensitivity of competences in these terms. However, there is more to the vagueness of know-how than only these phenomena. Everything I have said so far may in principle also pertain to mere abilities as opposed to full-blown skills. For I have only been talking about abilities to reliably meet the norms in question and left out the crucial element which distinguishes know-how, namely that one reliably meets the norms in question *because* one is guided by an understanding of them.

A further form of vagueness in conceiving of know-how stems from the vagueness of the boundary between an ability's being a mere ability and its being a full-blown competence. Ryle explicitly discusses this:

Animals & Infants (Ryle 1949, 121)

Our concern is only with a restricted class of dispositional terms, namely those appropriate only to [...] the characterization of such stretches of human behaviour as exhibit qualities of intellect and character. [...] Of course, the edges of this distinction are blurred. Dogs as well as infants are drilled to respond to words of

command, to pointing and to the ringing of dinner-bells; apes learn to use and even construct instruments; kittens are playful and parrots are imitative. If we like to say that the behaviour of animals is instinctive while part of the behaviour of human beings is rational, though we are drawing attention to an important difference or family of differences, it is a difference the edges of which are, in their turn, blurred. Exactly when does the instinctive imitativeness of the infant develop into rational histrionics? By which birthday has the child ceased ever to respond to the dinner-bell like a dog and begun always to respond to it like an angel?

Ryle is certainly correct that children are a wonderful example of the loose boundary between mere abilities and full-blown competences since it seems very clear that newborns only have mere abilities and adults possess genuine know-how. However, the details are far from trivial, and so is the question of non-human animals. I will come back to these points in § 6.7.

However, there is also a group of paradigm cases on the borderline between mere ability and know-how *within* the realm of normal human adults, namely, performances at the borderline between basic actions and non-basic actions. As I have argued in § 1.5, basic actions are not the things one may know how to do because they are not normative practices in the relevant sense. However, one may compete for the most finely-grained description available and argue over whether an allegedly basic action is indeed basic – that is, whether there is a more finely-grained description of the performance in virtue of which it can be described as a doing of something else. I do not need to take a stance on the general question if there is such a thing as ‘the’ or ‘the finest-grained’ description of a performance. But to the degree that we are unsure if an action is indeed basic, we will also be unsure if there is such a thing as knowing how to perform it. And since different people in different contexts draw the distinction between basic and non-basic actions differently, the distinction between know-how and mere ability is context-dependent.

For example, one may find it odd to say that I have the competence to walk or that I know how to walk. But as soon as the internal complexity of this everyday competence is brought into view, maybe in comparison with people who must re-acquire the capacity to walk after an accident, it is clear that I walk by moving my limbs and balancing my weight in many more simple and basic acts – acts in virtue of which I walk well. In this sense, it is also entirely innocuous to speak of the normativity of walking, of the standards which make the difference between somebody who walks elegantly or with ease or somebody who hardly even manages to shuffle along. Normally, such norms of the activity of walking are not at issue.

But as soon as one realizes how the complex arrangement of basic acts in an ordinary walk could go awry, it becomes intelligible, and natural, to see walking as a competences.

I have discussed forms of vagueness context-dependence of the concept of know-how which stem from the element which distinguishes know-how from mere ability – the fact that one reliably meets the norms in question because one is guided by an understanding of those norms. In chapter 2, I will go on to clarify this notion of understanding activities and of guidance by an understanding of an activity. One of the notions which plays a crucial role there is what Ryle calls ‘the intellect’. And in the discussion of this concept, he also admits its loose boundaries. He writes:

But after all, does it matter if all attempts at giving a hard-edged definition of ‘intellectual’ and ‘thought’ break down at some somewhere or other? We know well enough how to distinguish urban from rustic areas, games from work, and spring from summer, and are unembarrassed by the discovery of undecidable marginal cases. [...] Our daily use of the concepts of the intellect and of thought is unembarrassed by the discovery of a moderate number of borderline cases. (Ryle 1949, 267)

If the boundaries of the intellectual are vague, and if, as I will show, the concept of know-how partially depends on the intellect, it follows that the boundaries of the concept of know-how turn out to be vague, as well.

At the risk of stating the obvious, I would like to conclude by siding with Ryle’s assessment of this phenomenon. Vague concepts, as well as context-dependent concepts, are not therefore useless or otherwise problematic. There is no good reason to suppose that the characteristics of the concept of know-how which have been spelled out so far discredit this concept in any way. Likewise, it is not a flaw of a conception of know-how to call attention to these phenomena. Instead, it is a virtue of such a conception of these important facts about know-how are spelled out and explicitly account for in the proposed conception. And that was the project of this section.

Chapter 2

Ryle on Intelligence and Intellect

I have argued that know-how is an intelligent ability, an ability to do well in an activity in virtue of an understanding of what it takes to do so. In this chapter, I will continue to spell out this account by following Ryle's discussion of what he calls 'the intellect' and its role in intelligent practice.

I begin, in § 2.1, with the notion of exercising a skill and therein manifesting an understanding of the activity in question, distinguishing this phenomenon from further ways of manifesting such an understanding – say, by talking about an activity. In § 2.2, I discuss the interrelations between having know-how and having the competences to teach one's skills and to assess performances of them.

In § 2.3, I show how Ryle introduces what he calls 'the intellect' into the notion of know-how and discuss the crucial role of learning. The next section § 2.4 expands elaborates this in terms of the distinction between acquisition and improvement by learning. § 2.5 will show that such an understanding turns out to include propositional knowledge because this plays a crucial role in the acquisition of know-how. The concluding § 2.6 is devoted to a defense of this interpretation of Ryle's texts and to an account of *prima facie* contradicting passages.

§ 2.1 The Manifestation of Know-how

This section is concerned with the question what it means to manifest or exercise one's know-how. When a person knows how to do something, how does this competence manifest itself? There is one sense in which the answer to this question is "in a potentially infinite number of different ways", but there is another sense in which the answer is "in one way, and one way only". Let me explain.

When I exercise my competence to play squash, I do lots of different things. I observe the ball, I run through the court, I hit the ball, I watch the movements of my opponent, I feint to play the ball *here* and play it *there* instead, and so forth. Also, I do these different things in many different ways. For example, when I hit the ball, I may move my arm in a continuum of ever so slightly different ways. And similarly for all the other acts just mentioned and for the myriad further ways of behavior involved in squash. I take it that this is what Ryle has in mind when he says that competences are “dispositions the exercises of which are indefinitely heterogeneous” (Ryle 1949, 44). In this sense, my know-how manifests itself in a potentially infinite number of different ways.

By contrast, it is also true that I exercise my skill to play squash in one way, and in one way only. That is, I exercise my competence to play squash if and only if I play squash. All of the different things I have mentioned above are unified in that they are examples of a single activity – playing squash. What differs on a fine-grained level of distinguishing performances is nevertheless the same thing on a coarse-grained level of distinguishing performances.

This distinction between levels of granularity makes good on the intuitively contradictory ideas this section started out with. One exercises know-how by doing one thing and one thing only, but this single thing is realized in a potentially infinite number of ways.

This insight is more than a play on words. Among other things, it is crucial in understanding some otherwise rather puzzling remarks by Ryle. In the remainder of this section, I will elaborate and defend this position, and I will show how it is precisely what the most plausible interpretation of Ryle’s texts requires him to hold. To see this, consider first what he says about someone, indeed *anyone*, who exercises know-how:

[H]is observance of rules, principles, etc., must, if it is there at all, be realised in his performance of his tasks. It need not (though it can) be also advertised in an extra performance of paying some internal or external lip-service to those rules or principles. He must work judiciously; he may also propound judgments. For propounding judgments is just another special activity, which can itself be judiciously or injudiciously performed. (Ryle 1945a, 8–9)

Applied to a couple of examples, he puts this point as follows:

Cleverness at fighting is exhibited in the giving and parrying of blows, just as ability at reasoning is exhibited in the construction of valid arguments and the detection of fallacies, not in the avowal of logicians’ formulae. Nor does the surgeon’s skill function in his tongue uttering medical truths but only in his hands making the correct movements. (Ryle 1949, 48)

Thus, that the hallmark of my skill at squash are my performances of playing squash. Thinking or talking *about* squash would not be an exercise of my competence to play squash anymore. This is “another special activity” (Ryle 1945a, 9) which differs from the activity it is about.

On the other hand, when Ryle discusses the question whether somebody who happened to make a good shot actually knows how to shoot, he writes:

[W]e should take into account his subsequent shots, his past record, his explanations or excuses, the advice he gave to his neighbour and a host of other clues of various sorts. There is no one signal of a man’s knowing how to shoot, but a modest assemblage of heterogenous performances generally suffices to establish beyond reasonable doubt whether he knows how to shoot or not. (Ryle 1949, 45)

This suggests that his knowledge how to shoot can manifest itself not only when he shoots and shoots intelligently, but also when he gives advice about shooting, gives explanations, and so on. But this contradicts his view that talking about an activity is “another special activity” (Ryle 1945a, 9).

This tension is even more clear in the following pair of passages:

Principles of inference are not extra premisses and knowing these principles exhibits itself not in the recitation of formulæ but in the execution of valid inferences and in the avoidance, detection and correction of fallacies, etc. (Ryle 1945a, 7)

You exercise your knowledge how to tie a clove-hitch not only in acts of tying clove-hitches and in correcting your mistakes, but also in imagining tying them correctly, in instructing pupils, in criticizing the incorrect and clumsy movements and applauding the correct movements that they make, in inferring from a faulty result to the error which produced it, in predicting the outcomes of observed lapses, and so on indefinitely. (Ryle 1949, 54)

Suppose that – in line with the second passage – knowledge how to tie a clove-hitch can manifest itself in talking about how to tie a clove-hitch correctly. If this is true, why should it be impossible – as the first passage has it – that knowledge how to draw an inference manifests itself in talking about how to draw inferences correctly? Ryle appears to be inconsistent with regard to the question what counts as an exercise of know-how and what does not.¹ But given what I have laid out so far, it is possible to clarify these problems rather easily.

According to the Rylean view I shall defend, all we need in order to set the record straight is a distinction which Ryle evidently did not make clear enough: I have pointed out that knowing how to do something involves

¹ Part of the reason for this confusion on Ryle’s part may be the ambiguity or polysemy of the expression ‘knows how to’ which I shall discuss in detail in chapter 7.

an understanding of the norms which govern that activity. Thus, we can distinguish between an exercise of know-how and a *different* manifestation of the understanding which is part of that know-how.²

Let me take up the above examples in order to illustrate this distinction.

When I want to know whether somebody knows how to shoot, I can have a look at two different kinds of evidence. First, I can consider “his subsequent shots” and “his past record” (Ryle 1949, 45) and ask whether he has shot well so far and shoots well now. The result will be evidence for or against the hypothesis that he has a reliable ability to meet the norms of shooting. Since this is a proper part of what it takes to know how to shoot, this will also be evidence in deciding whether he has this know-how. Second, however, I can have a look at “his explanations or excuses, the advice he gave to his neighbour and a host of other clues of various sorts” (Ryle 1949, 45) and ask whether he has an apt understanding of shooting, whether he has a grasp of what it takes to shoot well. The result will be evidence for or against the hypothesis that he has an understanding of the norms of shooting well. Since this is a proper part of what it takes to know how to shoot, this will also be evidence in deciding whether he has this know-how.

Similarly in the case of drawing inferences. On the one hand, I can find out whether a student correctly draws inferences according to the rule of *modus ponens*, regardless of how she comments on this practice. And on the other hand, I can find out whether that student has an understanding of the rule of *modus ponens* by considering how she comments on such inferences or on formal or informal statements of the corresponding rule of inference. Both are part of what it takes to know how to draw inferences according to *modus ponens*. Therefore, both are legitimate pieces of evidence in order to decide if somebody has the competence to make *modus ponens* inferences.

The distinction between an exercise of know-how on the one hand and a *different* manifestation of the understanding which is part of that know-how on the other hand allows Ryle to maintain a claim which I have already presented – that every exercise of knowledge how to engage in some activity A is an A-ing. This view is very dear to Ryle. Most pointedly, he declares that “intelligence is exhibited by deeds, not by internal or external dicta.” (Ryle 1945a, 8) Thus, thinking, talking or theorizing about the standards of the activity of A-ing is *not* an exercise of the competence to A. Still, it *is* evidence for this competence because it is a manifestation of something which is necessary for it – of an understanding of the standards of A-ing.

But Ryle seems to contradict this view elsewhere:

² Benjamin Elzinga expresses a cognate distinction in terms of ‘first-order’ know-how and ‘second-order’ know-how (cf. Elzinga 2016, 7–10).

Knowing *how*, then, is a disposition, but not a single-track disposition like a reflex or a habit. Its exercises are observances of rules or canons or the applications of criteria [...]. Further, its exercises can be overt or covert, deeds performed or deeds imagined, words spoken aloud or words heard in one's head, pictures painted on canvas or pictures in the mind's eye. Or they can be amalgamations of the two. (Ryle 1949, 46)

It is natural to read this to entail that a competence can be exercised in *all* the ways Ryle mentions here. And Ryle's comments on the knowledge how to tie clove-hitches, as quoted on page 49, also sit well with this interpretation. This suggests that my thinking about my playing squash, my talking about squash and my picturing the playing of squash would *also* qualify as exercises of my knowledge how to play squash. But it is not clear that this passage should be read in this way. And given that this interpretation would lead to a self-contradiction in Ryle's view, we should look for an alternative.

Fortunately, such an alternative reading is easily available. For the scope of the quantifiers on the claim under discussion is far from precise. It is not clear that Ryle wants to say that *one and the same* competence can be exercised in all the ways he mentions. Instead, he could just as well want to say that *skills in general* can be exercised in all these ways. This, in turn, is compatible with the view that only some or even *no* individual competences can be exercised in *all* of those ways, as long as there is at least one competence for each. And clearly, there *are* competences where speaking, thinking or imagining are straightforward exercises of them.

Ryle is perfectly clear on this everywhere else. In addition to the above quotations, Ryle offers another elegant statement of his view:

When a person knows how to do things of a certain sort (e.g., cook omelettes, design dresses or persuade juries), his performance is in some way governed by principles, rules, canons, standards or criteria. [...] But his observance of rules, principles, etc., must, if it is there at all, be realised in his performance of his tasks. It need not (though it can) be also advertised in an extra performance of paying some internal or external lip-service to those rules or principles. (Ryle 1945a, 8–9)

In other words, the hallmark of somebody's knowledge how to A are her performances in A-ing. Since these rely on her understanding of what it takes to A well, she can also exhibit what Ryle calls “extra performances” of other activities such as thinking about the activity A or discussing what it takes to A well. But this does not show that such performances are also manifestations of the competence to A. They remain “extra performances”.

In the light of this account, it is also possible to make sense of an idea which Ryle expresses in passing in the following passages:

When a person knows how to do things of a certain sort (e.g., make good jokes, conduct battles or behave at funerals), his knowledge is actualised or exercised in what he does. It is not exercised (save *per accidens*) in the propounding of propositions or in saying “Yes” to those propounded by others. His intelligence is exhibited by deeds, not by internal or external dicta. A good experimentalist exercises his skill not in reciting maxims of technology but in making experiments. (Ryle 1945a, 8)

[T]he engineer’s schooling and workshop experience teach him to design bridges and not, save *per accidens*, to build or expound theories. (Ryle 1949, 298)

While Ryle here forcefully reiterates his core contention – that the competence to A is exercised in A-ing and in A-ing only –, he also speaks of a ‘*per accidens*’ exercise of know-how. The distinction between exercising know-how and otherwise manifesting one’s understanding of an activity clarifies this notion: The competence to conduct battles is exercised in conducting battles well and only in conducting battles well. But the understanding of warfare which is required for this skill can also manifest itself in theorizing about how to conduct a battle. From the perspective of the competence to conduct battles, these further manifestations of the understanding involved are accidental. They are not part of the essence of that competence.

In sum, there is a clear distinction between the exercises of somebody’s know-how, that is, her engaging in the activity in question, and other forms in which she may manifest her understanding of that activity.

§ 2.2 Teaching and Assessing

In § 2.1, I have only considered cases where somebody knows how to do something and therefore also has an understanding of this activity, and asked how the manifestations of her competence can differ from *other* manifestations of her understanding of the activity she has the skill to engage in. This possibility is built into the notion of competence because competence involves understanding. But there are several different such further ways in which an understanding of an activity can manifest itself.

In what follows, I shall comment on the two most important cases – the rather complex phenomenon of teaching an activity on the one hand and the more basic competence to assess performances of an activity on the other hand. Following Ryle, I shall argue that the capacity to teach a skill is independent from having that competence oneself, but that the capacity to assess performances of an activity is closer connected to that skill – indeed, it is necessary for it. The latter insight will prove to be crucial at various

later points in this book and play a crucial role in elucidating the notion of an understanding of an activity and its standards.

To begin with, I shall follow Ryle in considering a wide notion of teaching. While the paradigm case of teaching, say, how to speak a foreign language, is certainly a language lesson by a professional language teacher, there are also other things which fall under this broader notion. For example, a language learner can also be taught how to speak a language better than she already does simply by some small comments and corrections she receives in her conversations with native speakers. Such remarks certainly do not require the competence to give a whole training program in the language, but there is a good sense in which even these count as cases of teaching, however imperfect.

One of Ryle's most important points about teaching know-how is this:

Knowing how to behave is exhibited by correct behaviour, just as knowing how to cook is exhibited by palatable dishes. True, the conscientious man may be asked to instruct other agents how to behave, and then he will, if he knows how, publish maxims or specific prescriptions exemplifying maxims. But a man might know how to behave without knowing how to give good advice. [...] Knowing how to advise about behaviour is not the same thing as knowing how to behave. It requires at least three extra techniques: ability to abstract, ability to express and ability to impress. (Ryle 1945a, 13)

Ryle correctly states that the competence to do something does not entail the competence to convey one's know-how. But to draw a more nuanced picture, it seems that normal people with competences have something like a *minimal* capacity to convey their competences. Everybody with sufficient time, verbal and demonstrative capacities can try to convey their know-how by showing others how they engage in the activity in question and comment on it. However, it is crucial to see that having a skill does not entail any *substantial* competence at such teaching attempts. In other words, the way in which every skilled person seems to have a minimal capacity to teach that competence does not differ from the way in which an intelligent observer can learn from the example of another person, entirely independently of the question whether or not that person *intends* her performances as examples for the first one to learn from. I shall come back to this issue in § 2.3. In essence, then, Ryle is right. The competence to do something does not entail any substantial skill to convey one's know-how.

However, and much more importantly, the entailment in the other direction fails to hold, and very clearly so. The sentence left out in the middle of the passage just considered reads:

Sometimes a man might give good advice who did not know how to behave. (Ryle 1945a, 13)

Accordingly, it is possible to have an understanding of an activity, and to know how to teach others in this activity, but still fail to know how to engage in that activity oneself. The competence to teach something does not entail the skill to do it oneself.

This is not to deny that it is very helpful for a teacher to also possess the relevant competence herself, say, the knowledge how to perform a somersault. This way, a teacher has a first-personal experience of performing a somersault which certainly helps her to give advice on how to behave at certain stages of the performance. But, again, this does not require that a teacher must know how to perform a somersault herself, merely because she is good at teaching how to do it. The most common scenario which illustrates this point is one in which the teacher *used* to be able to perform a somersault, but has lost this competence in the meantime.

However, the clearest cases are those where a teacher *never* had this skill herself. Suppose that she is an expert in sports science, as well as an expert motivational psychologist, as well as an expert teacher, as well as a specialist for somersaulting. Such a person would have a competence to teach the somersault which borders on perfection. But all of this does not entail that she ever acquired the competence to perform a somersault herself. Ellen Fridland has illustrated this point as follows:

Bela Karoli (Fridland 2012, 9)

Bela Karoli is the world-famous coach of several gold-medal-winning women gymnasts. Bela Karoli knows the rules governing the skills that he teaches his athletes, but he is unable to perform these skills himself. It is not the case that he once could perform these skills, but is too fat or too old to perform them now. He never knew how to perform these skills and neither a good diet nor a good time machine could change that fact. Bela Karoli knows the rules governing the skillful performance of, e.g., a standing layout on beam, and he also knows how to express that knowledge in such a way that his gymnasts can apply it to their own learning and performance. He knows *about* how to perform a standing layout on beam. What he does not know, however, is how to perform a standing layout on beam.

I have argued that first-order skills are neither sufficient nor necessary for second-order skills to teach them. Ryle exemplifies this as well:

Critics & Writers (Ryle 1949, 49)

There have been thoughtful and original literary critics who have formulated admirable canons and prose style in execrable prose. There have been others who have employed brilliant English in the expression of the silliest theories of what constitutes good writing.

This example is also helpful in order to clarify the limits of the idea that competences entail a minimal capacity to teach them. I have already said that such a minimal capacity will often amount to nothing more than the possibility of practicing while somebody else observes one's deeds and tries to learn from what they perceive. The example of writing with literary style is instructive here because it makes vivid that such 'lessons' will sometimes be of very little use. In order to acquire a novelist's competences, it is not only very impractical to sit in on an esteemed novelist's writing sessions. It is probably also hopeless to do so.

Thus, teaching an activity is a possible manifestation of an understanding of an activity which is independent from having the knowledge how to engage in that activity oneself. However, there is a different manifestation of such an understanding which is indeed *necessary* for having the competence oneself – the capacity to assess performances of the activity in question.

Ryle expresses this insight in the slogan that "a person's appraisals of his own performances do not differ in kind from his appraisals of those of others" (Ryle 1949, 52). And given what I have laid out so far, it is easy to see why this is the case. To know how to do something is to have an ability to do so well in virtue of one's understanding of what it takes to do so. But when I understand what it takes to do something well, my understanding is not limited to *my* attempts to live up to these norms. Instead, it can also be applied to the acts of others. When I assess someone else's performances, I appeal to the same standards which also guide my own performances:

Of course, to execute an operation intelligently is not exactly the same thing as to follow its execution intelligently. The agent is originating, the spectator is only contemplating. But the rules which the agent observes and the criteria which he applies are one with those which govern the spectator's applause and jeers. (Ryle 1949, 53)

However, assessing performances of an activity does not necessarily involve the capacity to explicate the grounds for these assessments or to express them in a descriptively and explanatorily rich way. Ryle states that "the capacity to perform and to appreciate an operation does not necessarily involve the ability to formulate criticisms or lessons." (Ryle 1949, 55). And he gives the following example:

Knotting Sailor (Ryle 1949, 55)

A well-trained sailor boy can both tie complex knots and discern whether someone else is tying them correctly or incorrectly, deftly or clumsily. But he is probably incapable of the difficult task of describing in words how the knots should be tied.

The sailor boy is able to assess attempts to tie complex knots. This competence is part of his understanding of the tying of these knots, and without such an understanding, we could not coherently understand him as having the skill to tie them. But he lacks the appropriate conceptual capacities to describe or explain how these knots should be tied. He is surely able to refer *demonstratively* to a variety of examples and assess their respective and comparative qualities. But he cannot describe these procedures precisely.

However, while the capacity to assess a given performance of an activity is clearly *necessary* for having the skill oneself, Ryle is also led to assert that it is *sufficient* for having that skill. As I shall argue in the remainder of this section, this is where Ryle goes against his own declared view, and he is mistaken in this respect. However, this mistake can be explained in the light of the specific example he considered.

First, here is how Ryle asserts that skills are not only sufficient for having the capacity to assess performances of that skill, but also necessary for it:

The knowledge that is required for understanding intelligent performances of a specific kind is some degree of competence in performances of that kind. [...] [T]he one necessary condition is that he has some mastery of the art or procedure, examples of which he is to appraise. (Ryle 1949, 53)

The intelligent performer operates critically, the intelligent spectator follows critically. Roughly, execution and understanding are merely different exercises of knowledge of the tricks of the same trade. (Ryle 1949, 54)

The most blatant problem with this view lies in the possibility of being able to teach others how to do something without having the competence to do so oneself. This fact is incompatible with the idea that the capacity to assess a given performance of an activity is sufficient for having the competence to do so. The simple reason is that being able to teach others how to do something *already involves* the capacity to assess performances of this activity. It would therefore be inconsistent to go on to argue that this capacity is in turn sufficient for having the competence. For if it were, every competent teacher of a certain competence would necessarily – and miraculously – possess that competence herself.

Ryle wants his claim to be taken with a pinch of salt. He remarks that “the ability to appreciate a performance does not involve the same degree of competence as the ability to execute it.” (Ryle 1949, 54–55) This is a fair and important point. But it does not solve Ryle’s present problem. Even if the capacity to assess performances of a given activity does not entail a particularly *good* competence to engage in that activity, the idea that it entails such a competence *at all* remains very implausible. Or so I, following

Ryle himself, have argued earlier. Thus, Ryle is mistaken at this point. The capacity to assess the quality of performances of an activity is not sufficient for having the skill to engage in this activity, despite the fact that at least some degree of proficiency in assessing the quality of performances of an activity is indeed necessary for having the competence oneself.

But it should be noted that Ryle might well be correct *locally* as opposed to globally. That is, it is arguably true for some, but not for all activities that the capacity to assess performances entails the competence to perform. The example Ryle mentions in this context is very instructive:

The commentator on Plato's philosophy need not possess much philosophic originality, but if he cannot, as too many commentators cannot, appreciate the force, drift or motive of a philosophical argument, his comments will be worthless. If he can appreciate them, then he knows how to do part of what Plato knew how to do. (Ryle 1949, 53–54)

In short, Ryle holds that the competence to assess philosophy is not only necessary for knowing how to philosophize oneself. It is also sufficient. And in this case, I am very much inclined to agree. Philosophy seems to be a kind of activity in which being good at performing and being good at assessing performances are closely connected. But even if this is indeed the case for philosophy, and also for other activities, it is clearly still a local phenomenon. Globally, the capacity to assess performances does not entail the competence to produce these performances. In § 4.7, I shall come back to these phenomena and offer an explanation of why and where they occur.

To conclude this section, I would like to propose a further way to see the distinctions and connections between the competence to do something, the competence to assess performances of it and the competence to teach it. In § 1.2, I have proposed to individuate an activity in terms of the set of norms which govern it. Therefore, I can distinguish the skill to A from the skill to assess A-ing and from the skill to teach A-ing by distinguishing between the norms which govern A-ing, the norms which govern assessing A-ing, and the norms which govern teaching A-ing.

I take it to be clear that what it takes to do well at A-ing determines partially, but not fully, what it takes to do well at assessing A-ing. If one assesses, say, performances of playing football, then part of what makes such an assessment an accurate assessment is determined by the norms which govern football. The quality of these assessments is also determined by other factors – that is, not only by the facts about how well the relevant norms are met in these performances in question, but also by further elements such as the amount of time and information which is available to the one making

the assessment. Thus, A-ing and assessing A-ing are distinct activities, but they are related because what it takes to do well in the former partially determines what it takes to do well in the latter.

What it takes to do well at teaching A-ing is also partially, but not fully determined by what it takes to do well at A-ing. If one teaches skills at football, then part of what determines the quality of such lessons is determined by the norms of football because those who have been taught are supposed to actually live up to these norms later. Further elements also include the amount of time and information available to the teacher, as well as the pre-existing skills and health conditions of the students. Again, it follows that A-ing and teaching A-ing are distinct activities, but they are related because in the former partially, but not fully determines the latter.

As I have argued, both the competence to A and the competence to teach A-ing require the capacity to assess A-ing. In chapter 4, I shall come back to these points and discuss how exactly the capacity to assess A-ing plays a role both in the teaching of A-ing and in the activity of A-ing itself. Before this, however, I will now go on to follow Ryle's further discussion of the role of the intellect in intelligence.

§ 2.3 Learning and the Intellect

In sections § 2.1 and § 2.2, I have discussed the exercise of know-how and the way in which the underlying understanding of an activity can manifest itself in other ways. It has turned out that a necessary, but not sufficient condition for having a competence is the capacity to assess performances of the activity in question. And it has turned out that a sufficient, but not necessary condition for the knowledge how to make such assessments is the capacity to teach the competence in question to others. Now, I shall turn the tables and look at the acquisition of know-how not from the side of the teacher, but from the side of the learner. This, I will argue, sheds more light on the notion of understanding activities. The key to this argument is Ryle's discussion of what he calls 'the intellect'.

To begin with, Ryle is very careful to distinguish the intellect from intelligence and thereby from know-how. Most pointedly, he declares that “‘Intelligent’ cannot be defined in terms of ‘intellectual’ [...].” (Ryle 1949, 32). But already in chapter II of *The Concept of Mind*, the chapter on know-how, he draws a closer connection between know-how and ‘intellectual capacities’, one which is entirely compatible with the earlier insistence that the latter does not define the former. He writes:

It will be shown later (in Chapter IX), that the learning of all but the most unsophisticated of knacks requires some intellectual capacity. The ability to do things in accordance with instructions necessitates understanding these instructions. So some propositional competence is a condition of acquiring any of these competences. [...] I could not have learned the breast stroke, if I had not been able to understand the lessons given to me in that stroke [...]. (Ryle 1949, 48)

Thus, learning a skill from instruction involves the intellect. It involves understanding the instructions one is given and this involves something which Ryle calls, without much further elaboration, a “propositional competence”.

These promissory notes are indeed taken up in chapter IX of *The Concept of Mind*, but Ryle explicitly mentions how his discussion of the intellect bears on the problem on know-how only in chapter II. I shall argue that Ryle’s discussion of the intellect in chapter IX ties in well with his earlier elaboration of the concept of know-how and with his claim that instances of know-how “require some intellectual capacity” (Ryle 1949, 48).

I suggest that what Ryle here discusses under the label ‘intellectual capacity’ is nothing other but the understanding of an activity which is required for knowing how to engage in it (cf. § 1.5). If so, it is perfectly clear why “the most unsophisticated of knacks” do not require intellectual capacities (cf. Ryle 1949, 48). These are not competences, but mere abilities or dispositions. Intelligence requires the intellect.

On the face of it, the term ‘the intellect’ may appear way too specific and sophisticated in order to be the kind of thing that underlies the whole variety of activities one may have a competence to engage in. After all, skills include much more than the practices of intellectuals and Ryle is very aware of this (cf. *Ryle’s Range of Cases* on page 14). However, he is willing to use the terms ‘intellect’ and ‘intellectual’ even for those cases. In the passage just quoted from chapter II, he names even the activity of swimming the breast stroke as a skill which involves the intellect. And at the end of chapter IX, he repeats this assessment more generally:

It is easy to see that intellectual development is a condition of the existence of all but the most primitive occupations and interest. [...] We do not have to be scientists in order to solve anagrams, or play whist. But we have to be literate and be able to add and subtract. (Ryle 1949, 298)

But if ‘the intellect’ is supposed to cover *all* cases of understanding activities, how are we to make sense of this term? On Ryle’s view, we can understand the intellect in terms of the practice of teaching and learning. An intellectual feature turns out to be just *that* kind of feature which one paradigmatically acquires by learning it from somebody else who teaches it. He writes:

There is one idea not far from the forefront of most people's minds when they contrast intellectual powers and performances with other powers and performances, namely that of schooling. The intellectual powers are those or some of those which are developed by set lessons and tested by set examinations. Intellectual tasks are those or some of those which only the schooled can perform. (Ryle 1949, 268)

[I]t is in terms of didactic discourse that the concept of the intellect is being elucidated. At least an important part of what we mean by 'intellectual powers' is those specific capacities which are originally inculcated and developed predominantly by didactic discourse [...]. (Ryle 1949, 291)

It is important to stress that Ryle never *equates* the intellect with what one has been taught. He only says that the intellect is "elucidated" by learning and teaching and that "[a]t least an important part" of intellectual accomplishments are acquired by being taught (Ryle 1949, 291). But how should the intellect be understood more precisely?

I propose to clarify Ryle's view by leveraging the explanatory weight away from teaching and over to *learning*. On this view, a feature is an intellectual feature – an instance of understanding or knowledge – just in case it has been acquired by learning it – be it by being taught or in other ways. I take it that this interpretation is most faithful to the progression of Ryle's elaboration of the intellect over the course of the whole of chapter IX. To see why, consider the following idea with which Ryle starts:

Native or untutored knacks are not classed with intellectual proficiencies, and even arts learned mainly by sheer imitation, like skipping, playing Snap, and chatting, are not spoken of as intellectual accomplishments. This certificate is reserved for exploitations of lessons learned at least in part from books and lectures, or, in general, from didactic discourse. (Ryle 1949, 268)

Later, Ryle distinguishes these two claims more carefully, maintaining one and retracting the other. On the one hand, he defends the claim that native capacities are not intelligent, i.e. not guided by the intellect. And rightly so. Since competences involve an understanding of what activities require, they cannot come into existence in any way which does not involve the acquisition of such an understanding. Know-how requires *learning*.³

³ There are certainly borderline cases with respect to the criterion that know-how must be acquired by learning. But these borderline cases are at the vague and context-dependent borderline between basic and non-basic actions on which I have already commented in § 1.5 and § 1.7. For example, one may say that I learned how to wiggle my ears because I have tried doing so often enough, grimacing awkwardly in front of a mirror. At some point, I got the hang of it and acquired this ability. But there is only a very far-fetched sense in which I now *know* how to wiggle my ears. This is because there is only a far-fetched sense in which I have an understanding of ear-wiggling which guides my performances. I just do this as a basic action. I shall come back to this on page 151 and discuss a puzzle case involving ear-wiggling by Bengson & Moffett.

On the other hand, Ryle quickly steps back from the claim that learning competences necessarily requires being taught in ‘didactic discourse’:

Of course, not all teaching is done by talking didactically. Infants learn things by following examples which may, or may not, be deliberately set for their imitation. Some lessons are taught by deliberately setting examples and giving demonstrations. (Ryle 1949, 292)

Didactic influence can be exerted not only by one person upon another but also by one person upon himself. He can coach himself to say and do things which are not echoes of the words in which that coaching is given. [...] When we comment on a person’s intellectual proficiencies and limitations, the main things we have in mind are his efficiency and keenness in making such advances. (Ryle 1949, 294)

In the first of these passages, Ryle extends his conception of teaching from what he takes to be the paradigm cases of ‘talking didactically’ to other practices, including non-verbal ones where students gain insight from example and demonstration. And in the second passage just two pages later, it turns out that the core of the intellectual is not what teachers do, but what learners do. To learn, in this sense, is to continue to be open to improvement. And this crucially includes practices of teaching oneself.

This opens up the possibility of conceiving of even further situations as ‘teaching’ somebody something.⁴ In fact, it already entails this. A student can learn from a teacher’s example, but what matters in the end is not the fact that the example has been set by a teacher, but the fact that a student *takes* her teacher’s performances *as* an example she can learn from. But then it follows that a student can also learn from other examples which are instructive to her, even if these were not intended as demonstrations. Somebody who secretly observes a carpenter manufacture a piece of furniture can, given suitable background competences, learn from this example how to manufacture another such piece. Despite the fact that the carpenter’s work is not *set* as an example for others to learn, it nevertheless *is* an example from which others *can* learn. What counts is not what teachers *set as* examples for learners, but what learners *take as* examples.

This point can also be found explicitly earlier in Ryle’s discussion:

But it would be quite possible for a boy to learn chess without even hearing or reading the rules at all. By watching the moves made by others and by noticing which of his own moves were conceded and which were rejected, he could pick up the art of playing correctly while still quite unable to propound the regulations

⁴ Katherine Hawley has offered an illuminating discussion of teaching and learning of know-how (cf. Hawley 2010). She also poses a problem for anti-intellectualism concerning know-how from testimony. I comment on this topic in § 2.5 and § 6.3.

in terms of which ‘correct’ and ‘incorrect’ are defined. We all learned the rules of hunt-the-thimble and hide-and-seek and the elementary rules of grammar and logic in this way. We learn *how* by practice, schooled indeed by criticism and example, but often quite unaided by any lessons in the theory. (Ryle 1949, 41)

The general aim of this passage is to defend the view that competences do not require explicit knowledge of rules or norms, a point I have already stressed and discussed in § 1.5 and § 2.2. But Ryle also states – in passing, but very explicitly – that somebody can “pick up the art of playing correctly” – i.e. acquire the competence to play correctly – simply by keen observation.

Thus, the explanatory burden in Ryle’s elucidation of the concept of the intellect lies clearly on the idea of learning, and not on the specificities of teaching. Teaching is the paradigm source of learning, but learning is still a broader affair which also covers cases in which learners are not taught.

Still, there appears to be a sense in which teaching is more than the paradigm source of learning. It could very well be the *origin* of all instances of learning, even if those later instances do not then involve teaching. For example, the secret observer of the carpenter just imagined can only learn from her example if he is able to grasp her example *as* an example for himself to learn from. But maybe this is only possible if the observer has himself already received some explicit teaching. If this is true, one can learn without being taught, but one cannot learn without being taught unless one has been taught before. This clarification notwithstanding, I maintain that not every instance of learning involves being taught there and then.

§ 2.4 Acquisition and Improvement

I have argued that Ryle’s elucidation of the concept of the intellect in chapter IX of *The Concept of Mind* is centered around the capacity to learn and improve. In this section, I shall expand on this argument. On the one hand, I will tie this discussion back to the specific issue of competences. And on the other hand, I will take into account the double nature of the notion of learning, both as acquisition of a skill and as its improvement.

In fact, Ryle’s concern for the capacity to learn is not limited to his discussion of the intellect chapter IX. This is also what he uses to demarcate know-how in chapter II. An important passage which I already employed in § 1.5 ties the understanding of norms very closely to learning:

To be intelligent is not merely to satisfy criteria, but to apply them; to regulate one’s actions and not merely to be well-regulated. A person’s performance is described as careful or skilful, if in his operations he is ready to repeat and improve

on successes, to profit from the examples of others and so forth. He applies criteria in performing critically, that is, in trying to get things right. (Ryle 1949, 29)

Later, Ryle even stresses the connection between know-how and continuous improvement as something *essential* to know-how:

It is of the essence of intelligent practices that one performance is modified by its predecessors. The agent is still learning. (Ryle 1949, 42)

Thus, the core criterion for demarcating know-how as an ability to succeed in virtue of understanding has as its *essence* what Ryle otherwise calls an ‘intellectual’ feature – an understanding acquired and continuously refined by learning. Again, intelligence requires the intellect.

To be clear, the criterion that current exercises of one’s competences must always be used to improve on future such exercises should not be too strong. Ryle’s formulation that it is essential “that one performance is modified by its predecessors” (Ryle 1949, 42) is clearly too demanding. For example, this would make it difficult to conceive of an exercise of a skill which is known to be the very last – say, when somebody knows that they will die shortly afterwards. Instead, Ryle’s idea is that something counts as an exercise of a competence only if it is *possible* for this performance to figure in the future improvement of that competence. In Ryle’s more apt formulation, he speaks of *readiness* to continue to learn, of being “ready to repeat and improve on successes [...] and so forth” (Ryle 1949, 29).

But how is such a possibility and readiness to learn *essential* to a performance’s being an exercise of know-how? As I have laid out so far, it is essential to an ability’s being know-how that it involves an understanding of what it takes to succeed in the activity in question – a capacity to assess performances against the standards of this activity. On this basis, the possibility to *implement* these assessments in one’s further conduct already shows that know-how essentially involves the possibility to learn from past conduct. For now, I will have to leave this argument at this intuitive sketch. However, I will return to this crucial idea in § 4.4, where I spell out this notion of an implementation of assessments more fully.

I have argued that intelligence requires the intellect in the sense that know-how requires an understanding acquired and continuously refined by learning. Apart from the general argument for this claim I have laid out so far, there is also a further systematic case to be made.⁵

⁵ Will Small has recently suggested further arguments in this direction (cf. Small (2014), with particular emphasis on the social aspects of learning and teaching skills and on learning through guided practice. While I cannot discuss these considerations here in detail, I take it that the view presented here is very much in the same spirit.

As I have argued, to know how to do something is to be able to meet the norms of an activity in virtue of guidance of an understanding of them. But there is a spectrum of differences between individual activities with respect to the question how fixed or settled these norms are. At one end, there are activities the standards of which are entirely or almost entirely settled – such as calculating the sums of one-digit numbers in the decimal system. At the other end, there are activities the standards of which are entirely or almost entirely variable – such as entertaining an audience with jokes. But it does not matter where any specific activity and or competence lies in this spectrum, such as the ones I mentioned here. Given vagueness and context-dependence (cf. § 1.2 and § 1.7), this is complicated anyway. What matters is only that there is such a spectrum, with certain examples here and there.

Normally, activities lie somewhere between these extremes, with some aspects settled and others in flux. In particular, very general norms are often rather stable while their consequences for specific cases and situations are more variable. In any case, I shall formulate my argument for the view that intelligence involves continuous learning in terms of these extremes, but my considerations will also apply at any point in between.

On the one hand, if a normative standard is in flux, I take it to be very clear that guidance by such a norm requires the continuous awareness of these changes and appropriate adjustments. Suppose that somebody either fails to realize that it takes something else to surprise an audience with a new joke than it took an hour ago, or who does realize this, but fails to adjust to it. *Ceteris paribus*, such a person would not count as competent at surprising people with new jokes. Thus, guidance by norms in flux requires continuous learning because it requires continuous adjustment to changes.

On the other hand, consider normative standards which are completely settled. Arguably, nobody meets these normative demands *perfectly* – that is, always, in every respect and to a degree of 100%. Even if some performances are perfect in *some* respects, there is always room for improvement in *elsewhere* such as in reliability, efficiency, elegance or ease. Competences entail an awareness of this and the drive for improvement where possible.

Ryle understands moral knowledge as know-how rather than as propositional knowledge. On this basis, he beautifully illustrates the present point by generalizing from moral knowledge to know-how in general:

Moral imperatives and ought-statements have no place in the lives of saints or complete sinners. For saints are not still learning how to behave and complete sinners have not yet begun to learn. So neither experiences scruples. Neither considers maxims. Logical rules, tactical maxims and technical canons are in the same way helpful only to the half-trained. (Ryle 1945a, 14)

Maxims, imperatives and ought-statements about what to do are helpful only for those who still learn. But this is true of all of us. Just like every real person is somewhere in between being a sinner and being a saint, every real competence is somewhere in between perfection and failure.

As a side note, it is important to see that Ryle's talk of imperatives does not undermine the fact that genuine knowledge is involved here, despite the fact that Ryle attempts a quick argument on this direction elsewhere:

[W]hen we try to express these principles we find that they cannot easily be put in the indicative mood. They fall automatically into the imperative mood. Hence comes the awkwardness for the intellectualist theories of stating what are the truths, or facts which we acknowledge when we acknowledge a rule or maxim. We cannot call an imperative a truth or falsehood. (Ryle 1945a, 12)

Ryle makes this case only here. Elsewhere, he easily switches between imperative and declarative expressions, both in other passages of his Presidential address and in *The Concept of Mind*. Accordingly, I shall continue to disregard this distinction. For the purposes of this book then, imperatives which express the command to do certain things in order to, say, cook a certain dish, and can be reformulated in declarative statements which assert that doing so is what one should or ought to do in order to cook that dish.

In sum, understanding the normative requirements of an activity involves the insight that one's competence can always be improved. Every attempt is an opportunity to refine and improve one's competence, and essentially so. In this sense, exercising know-how is “performing critically” (Ryle 1949, 29) and making sure that “one performance is modified by its predecessors” because one is “still learning” (Ryle 1949, 42). It is in terms of this idea of learning – from teaching and practical experience – that we can elucidate what it is to understand an activity. My proposal for an account of understanding and guidance in chapter 4 will be in precisely this spirit.

§ 2.5 A Place for Propositional Knowledge

I have argued that the understanding of an activity and its normative demands which underlies the competence to engage in it is, in Ryle's word, something intellectual. It is something which is acquired by learning, the paradigm of which is learning by being taught. Given this background, I shall now go on to show that there also is a place for propositional knowledge in every competence. Understanding an activity requires at least some propositional knowledge. As I shall argue, propositional knowledge plays a crucial role in the acquisition of know-how.

This insight can also be found very clearly in Ryle's own texts, despite the fact that received interpretations of Ryle have him denying any role of propositional knowledge. While I will comment on these views later, in § 2.6, the present section is devoted to the systematic positive case for a place for propositional knowledge.

The key to this point lies in the close connections Ryle draws between two concepts: being able to talk didactically and having theories or plans. To begin with, Ryle has a very broad notion of 'theory' and 'theorizing':

[T]he word 'theory' has widely different senses. Sherlock Holmes's theories were not built by the same methods as those of Marx, nor were the uses or applications of them similar to those of Marx. But both were alike in delivering their theories in didactic prose. (Ryle 1949, 269)

[I]n talking of building theories I am not referring only to the classical examples of famous discoveries but to a class of tasks in which all people who have had any education participate to some degree on some occasions. [...] I am also using the word 'theory' to cover the results of any kind of systematic inquiry, whether or not the results make up a deductive system. (Ryle 1949, 272)

Thus, any piece of insight which can be expressed in what Ryle calls 'didactic prose' qualifies as a bit of theory in this broad sense. And at least in principle, theories or plans are the kind of thing which one can teach or tell:

[T]o have a theory or plan is to be prepared either to tell it or to apply it, if occasion arises to do so. [...] Having a theory or plan is not merely being able to tell what one's theory or plan is. Being able to tell a theory is, in fact, being able to make just one, namely the didactic exploitation of it. (Ryle 1949, 270)

We might say, therefore, that in theorizing the soul is, *inter alia*, preparing itself to talk or write didactically; and that the intended benefits to the recipient consist of acquired preparednesses to act and react in various new ways, only some of which will themselves be further didactic pronouncements. (Ryle 1949, 271)

Thus, theories or plans are what can be expressed didactically, which is to say that they can be told to somebody else. And at this point, Ryle finally seems to feel comfortable to introduce a notion he otherwise eschews:

It will not escape those who are familiar with the philosophical discussions of the nature and status of what are called 'propositions', that the predicates by which propositions are described are just those which do belong *ex officio* to the jobs of didactic discourse [...]. It is no accident that some theorists like to define 'intellectual operations' as operations with propositions, or that other theorists like to define 'propositions' as the product or implements of intellectual operations. Both are implicitly referring to our lesson giving, lesson taking and lesson using activities and powers, without, of course, explicitly mentioning such vulgar matters. (Ryle 1949, 293)

In his own explanations, Ryle largely avoids to use the term ‘proposition’ himself. But he does not strictly refuse to use this notion. After all, he explicitly states that the intellectual competence involved in understanding activities is a “propositional competence”. (Ryle 1949, 48) But here, finally, he explains how to understand this notion. Propositions are the *contents* of teachings and tellings, of rememberings, thinkings, and so forth.

It is useful to have a conception of such contents as propositions because what we teach and tell each other, what we listen to, repeat and eventually learn are things which it is useful to be able to identify and re-identify intersubjectively and intrasubjectively. For example, I can think the same thing on different occasions, forget it in between and be reminded of it by somebody else who explicitly asserts this thought. Ryle expresses this insight by saying that “the predicates by which propositions are described” include that they can be “accumulated, assembled, compared, sifted, and criticized” (Ryle 1949, 292–293). According to him, these characterizations are precisely “the predicates [...] which do belong *ex officio* to the jobs of didactic discourse” (Ryle 1949, 293).⁶ And of course, this view of the role of propositions is widely shared and usually expressed as the idea that we need to appeal to them in explaining propositional attitudes.⁷

Finally, Ryle claims that the natural goal of our intellectual endeavors turns out to be a specific propositional attitude – knowledge:

When we speak of the intellect or, better of the intellectual powers and performances of persons, we are referring primarily to that special class of operations which constitute theorizing. The goal of these operations is the knowledge of true propositions or facts. (Ryle 1949, 27)

Paradigmatically and in the case of success, what a student acquires by learning it from a teacher who tells her something is propositional knowledge by her teacher’s testimony.⁸ And such testimonial propositional knowledge is therefore an important part of what it means to understand an activity, and therefore also of what it means to be guided by such an understanding. Ryle brings up the concept of a proposition only in passing and only at the end of his discussion of the intellect because he is engaged in a larger project

⁶ Ryle defended a similar position in his early “Are There Propositions?” (cf. Ryle 1930)

⁷ There is a substantial debate about the nature of propositions and about the way they are involved in the explanation of propositional attitudes (cf. Brogaard 2008a; McGrath 2012; McKay & Nelson 2010; King 2011; Fitch & Nelson 2013). Nevertheless, I take it that Ryle’s explication of the place of propositionality is sufficiently uncontroversial.

⁸ Of course, there are many important and interesting problems in the epistemology of testimony on which I cannot comment here (cf. e.g. Adler 2012). My claim that the cases described constitute cases of testimony is intended to be neutral with respect to both additional and more specific questions.

trying to demystify this notion. But given his own account of the intellect and his explications of the point of these notions, it is now safe to use this concept in the way he offers himself.

This has merely been a sketch of the place of the notion of propositional knowledge and of Ryle's view of this notion.⁹ Still, this finally puts us in a position to make the case for a place for propositional knowledge in the acquisition and possession of *know-how*.

I have already claimed that propositional knowledge is involved in the understanding of what an activity requires. But why exactly should this be so? In what follows, I shall argue that propositional knowledge plays a crucial role in the process of learning a competence. To see this, it is instructive to consider the way in which Ryle describes swimming instructions as a paradigm example for learning by being taught:

Didactic talk is meant to instruct. The swimming instructor says things to his pupils, but he is not primarily intending to get the pupil to say those same things. He intends him now to make the required strokes with his arms and legs and later to make strokes like these without the accompaniment of spoken or silent instructions. Ultimately, perhaps, the pupil will teach other novices to swim, or at least teach himself to make new strokes or to make the old strokes in more difficult conditions. Learning the imparted lesson is becoming competent, not merely or primarily to parrot it, but to do a systematic variety of other things. The same holds good of more academic lessons [...]. (Ryle 1949, 293–294)

What a swimming instructor wants her pupils to acquire is the skill to swim. That is, she wants them to acquire an ability to swim well in virtue of an understanding of what it takes to swim. In this passage, Ryle says twice that such a teacher does “not primarily” want her student to be able to repeat the given instructions. And this is fair enough because the student is supposed to manifest his understanding of what it takes to swim *by swimming well* rather than in any other way (cf. § 2.1). But because what should become manifest is an *understanding* of what it takes to swim, the student thereby *also* acquires the ability to manifest this understanding in *other* ways. Teaching something “not primarily” is still teaching it – ‘secondarily’ or ‘*per accidens*’, to borrow Ryle’s own expression (cf. Ryle 1945a, 8; Ryle 1949, 298) on which I have commented in § 2.1.

Thus, a student who is verbally instructed in swimming does not merely have an opportunity to learn how to swim. He can also gain testimonial propositional knowledge about swimming. Afterwards, his understanding of swimming can manifest itself for example when he passes on the lesson he has received – i.e. when he expresses this propositional knowledge.

⁹ John Hyman (1999) and Michael Kremer (2016) discuss this topic more thoroughly.

But these developments can also be connected. The propositional knowledge which the student acquires often figures as a stepping stone in the process of acquiring an understanding of the activity in question and the competence to engage in that activity. Ryle writes:

It is, of course, true that when people can reason intelligently, logicians can then extract the nerve of a range of similar inferences and exhibit this nerve in a logicians' formula. And they can teach it in lessons to novices who first learn the formula by heart and later find out how to detect the presence of a common nerve in a variety of formally similar but materially different arguments. (Ryle 1945a, 7)

In such a cases, some of the rules or norms which givern the activity in question are expressed by a teacher in propositional form, and then conveyed to students who can come to acquire testimonial propositional knowledge about the activity in question.

Of course, and as discussed in § 2.3 and § 2.4, competences are not always learned by being taught. But all of them are acquired by *some* process of learning. How is propositional knowledge involved in these other cases?

One important family of cases involves learning from trial and error and by coincidence. Ruth Millikan puts this as follows:

Pink Ice Cream (Millikan 2000, 64)

[O]bserving only once that you have a certain capacity can immediately turn it into an ability. Anything that you find out you can effect immediately becomes an ability. Having observed that stirring the red strawberries into the vanilla ice cream turns it pink, the child knows how to make pink ice cream.

In such a case, the child has gained propositional knowledge through her experiences of playing around with food. She has acquired a competence – knowlede how to make pink ice cream – because she has learned that stirring the red strawberries into the vanilla ice cream turns it pink and because she was antecedently able to do just that. The child already had a sufficiently rich understanding of what it takes to color ice cream such that she was able to realize that stirring something into it may do the job.

In this and other ways, crucial pieces of propositional knowledge can play the role of a stepping stone in gaining or deepening the understanding of an activity. Millikan's *Pink Ice Cream* brings out that it is less important whether the relevant propositions were learned from a teacher or from experience or, indeed, from any other source. What counts in the end is their role in making the understanding of the relevant activity sufficiently rich such that the person is then able to guide herself in the light of this understanding and actually proceed to engage in the activity herself.

One can even generalize from examples like these. As discussed in § 2.3 and § 2.4, there are many ways in which somebody may acquire or improve a competence in learning – whether by being taught explicitly, or by observing somebody else, or even by observing their own acts and seeing the potential for improvement. In all of these cases, there is a clear sense in which such a learner gains knowledge about the quality of specific token performances. Whatever the relevant example was, the learner comes to know that this example was a good or a bad exercise of the relevant activity, or an efficient or a subtle one, or indeed any other relevant kind of normative qualification. Thus, every case of competence-learning involves such propositional knowledge, even if this knowledge is merely demonstrative.

I will elaborate on this crucial point in chapter 4, and particularly in § 4.3. As I shall argue there, there is even an additional case to be made why every competence involves such propositional knowledge. As of now, however, I have merely argued that propositional knowledge plays a crucial role in the acquisition of know-how, among other things as a steppig stone in gaining the understanding which is involved in competence.

But it should be noted that this does *not* show that learning a competence necessarily involves propositional *knowledge* rather than merely true belief or merely justified true belief. However, as I will argue with respect to the problem of epistemic luck in § 6.3, this is a strength rather than a weakness. While Ryle's paradigm case of learning by being taught is certainly such that the true beliefs in question do amount to knowledge *qua* expert testimony, many other cases can also be construed in such a way that no full-blown knowledge is involved. But since *successful* cases of testimony nevertheless form the core of the practice of explicit teaching, it remains true that propositional knowledge is crucial for acquiring know-how.

And the same holds for other cases of learning. The observational beliefs about specific instances of exercising a given activity may be nothing more than true beliefs. But if all goes well, they, too, will amount to genuine propositional knowledge.

§ 2.6 Ryle on Knowledge and the Mind

In the preceding sections, I have offered an interpretation of Ryle's texts, and the beginnings of an independently developed Rylean account. The resulting views include a number of elements which may seem alien to Ryle. In this final section of this chapter, I would like to point out where the account offered in this book departs from Ryle. I will mainly comment on the role

of propositional knowledge as a necessary condition of competence and on the question of primacy between know-how and propositional knowledge. At the end, I will close with some remarks on the place of know-how in the whole of Ryle's *The Concept of Mind*, and on Ryle's stance on behaviorism.

As for the first point, it is fair to say that the received view of Ryle's philosophical outlook is one which does not reserve any substantial role for propositional knowledge. And there are indeed passages in which Ryle explicitly states that know-how does not require *any* propositional knowledge. However, I hope to have shown that the best account of Ryle's discussion of know-how and of his own view of the place and point of the notion of propositional knowledge entails that these dissenting statements should be seen as less decisive. In order to save the coherence of the Rylean view defended here, some of Ryle's own words must be explained away.

In order to spell out how this can be accomplished, I shall comment on the clearest passage in question. Ryle declares:

Efficient practice precedes the theory of it; methodologies presuppose the application of the methods, of the critical investigation of which they are products. It was because Aristotle found himself and others reasoning now intelligently and now stupidly and it was because Izaak Walton found himself and others angling sometimes effectively and sometimes ineffectively that both were able to give to their pupils the maxims and prescriptions of their arts. It is therefore possible for people intelligently to perform some sorts of operations when they are not yet able to consider any propositions enjoining how they should be performed. (Ryle 1949, 31)

The last sentence of this passage clearly denies that know-how could possibly require any propositional attitude, let alone knowledge, about how a given activity should be performed. But it is crucial to see that the sentences which precede this assertion do not constitute an argument for this view, despite the fact that Ryle uses "therefore" to indicate this.

He begins with the observation that people who have developed influential theories of activities such as arguing and inferring in the case of Aristotle and fishing in the case of Izaak Walton could not possibly have done what they have done were it not for the fact that these practices were already functioning well. In this sense, intelligent practice clearly precedes theory. Elsewhere, Ryle makes this point with the aid of elegant analogies:

In short the propositional acknowledgement of rules, reasons or principles is not the parent of the intelligent application of them; it is a step-child of that application. (Ryle 1945a, 9)

Rules, like birds, must live before they can be stuffed. (Ryle 1945a, 11)

But unfortunately, this insight does not entail what Ryle takes it to entail. Given his own clarification of the concepts of theory, plan, and proposition (cf. § 2.5), there is lots of propositional knowledge to be had which does not amount to advanced theory. The fact that there must already *be* competences before their advanced theory can come into being does not entail the possibility of competences without any propositional knowledge whatsoever.

I propose to diagnose this shortcoming as a lack of oversight over the whole of *The Concept of Mind*. In the passage just considered, which stems from page 31 of this monograph, Ryle uses the term “theory” in a narrow sense, referring to advanced and explicit theory involving sufficiently rich descriptions and explanations. In § 2.5, I have argued that Ryle liberalizes the use of the notion of theory in chapter IX so as to include all the mundane things which “all people who have had any education” can come to know (Ryle 1949, 272).

This can be illustrated with an example which Ryle discusses just before the passage under consideration:

The wit, when challenged to cite the maxims, or canons, by which he constructs and appreciates jokes, is unable to answer. He knows how to make good jokes and how to detect bad ones, but he cannot tell us or himself any recipes for them. So the practice of humour is not a client of its theory. (Ryle 1949, 30)

Know-how certainly does not require the ability to give a “recipe” for performing well – a complete or near-complete statement of all and only those steps necessary in order to succeed. However, it does require the competence to detect failure, which is part of the general requirement of the capacity to assess performances in general. This, however, is intelligible only if we can understand the person in question as having some propositional knowledge.

First, as discussed in § 2.5, detecting failure requires knowledge of necessary conditions which must be satisfied for a performance to qualify as meeting the norms of an activity. In the case of joking, these necessary conditions might be hard to state or even very trivial – à la “A joke can be good only if at least some people can get it”. But this difficulty is beside the current point. However the relevant knowledge is to be construed precisely, somebody who has no such propositional knowledge whatsoever cannot be understood as knowing how to make jokes.

Second, as I will spell out in more detail in chapter 4, it is not even necessary to be able to *describe* such necessary conditions in any substantial way. For the capacity to “detect” successes and failures, as Ryle puts it in the passage just quoted, it is already sufficient to be able to refer *demonstratively* to certain cases. It is already full-blown propositional knowledge to know *de*

re, of a given performance, that *it* is a failure, or a success, or how else it may be evaluated with respect to the relevant activity. Having such propositional knowledge is necessary for know-how.

I take it that a large part of Ryle's criticism of the idea that propositional knowledge plays an important role for know-how stems from the fact that focused on propositions of the form of universally quantified rules or maxims where “the reason, or maxim, is inevitably a proposition of some generality” (Ryle 1949, 31). If this arbitrary restriction is lifted, it becomes even clearer why propositional knowledge is involved in know-how.

At the end of his discussion of the intellect, Ryle is much clearer about the relationship between know-how and propositional knowledge, or between intelligence and the intellect. There is even a subsection with the telling title “The Primacy of the Intellect”, which starts as follows:

It is now easy to distinguish the sense in which intellectual operations are higher than, and do ‘govern’, the exercises of other mental capacities, from the sense in which I have denied that the occurrence of intellectual operations is implied in all those descriptions we give of people's actions and reactions which embody mental concepts. (Ryle 1949, 295–296)

Accordingly, it is one thing to say that *having* know-how requires *having* propositional knowledge about the activity in question. It is quite another thing to say that such (or any) propositional knowledge must be the object of explicit thought whenever one *exercises* a competence or that it single-handedly *explains* such an intelligent performance. The latter claim leads to what Ryle calls ‘the intellectualist legend’, a view which he dismisses with a forceful regress argument on which I shall comment in chapter 9. Still, it is important to stress already at this point that Ryle sees no contradiction between his rejection of intellectualism and his affirmation of the view that there is, after all, *some* sense in which the intellect is primary. Competences are abilities to achieve success in virtue of an understanding of what it takes to do so – an understanding which involves propositional knowledge.

This naturally leads to the question of priority in the relationship between know-how and propositional knowledge. At one point, Ryle explicitly sets out to show that “knowledge-how is a concept logically prior to the concept of knowledge-that” (Ryle 1945a, 4–5) and explains:

[K]nowing-that presupposes knowing-how. [...] To know a truth, I must have discovered or established it. But discovering and establishing are intelligent operations, requiring rules of method, checks, tests, criteria, etc. A scientist or an historian is primarily a man who knows how to decide certain sorts of questions. Only secondarily is he a man who has discovered a lot of facts, i.e., has achieved successes in his application of these rules, etc. (Ryle 1945a, 15–16)

Thus, know-how and propositional knowledge presuppose each other – these kinds of knowledge are essentially interdependent. This interpretation of Ryle, and the plausibility of view, has been spelled out most clearly by David Wiggins (2012) and Michael Kremer (2016). Since this topic will continue to play a role at various points over the course of this book, I shall here only offer some initial considerations.¹⁰

On the one hand, know-how requires propositional knowledge because know-how requires an understanding of the activity in question, a necessary ingredient of which is at least some relevant propositional knowledge about that activity – e.g. in the form of specific assessments of examples (cf. § 2.5 and § 4.3). On the other hand, Ryle is equally right when he says that propositional knowledge requires know-how because one cannot know that something is the case without having exercised one's competences in order to find out that it is the case. Of course, it is difficult to name and distinguish the relevant epistemic competences for every kind of proposition one may come to know. But one does not have to settle this question in order to see the present point. For propositional knowledge certainly requires *conceptual* capacities, capacities which allow somebody to understand the content of the proposition in the first place. But, as I will discuss in more detail in chapter 4, conceptual capacities are themselves competences.

Thus, know-how and propositional knowledge presuppose each other. They are interdependent. To possess knowledge *at all* always means to possess *two kinds* of knowledge states. But none of these kinds of knowledge is more fundamental than the other.

In sum, the best interpretation of Ryle's view includes the claim that competences require an understanding of activities and this requires at least a minimum of propositional knowledge about this activity. But this claim is compatible with Ryle's rejection of the view that competences require anything over and above such a minimum of propositional knowledge, let alone theories or 'recipes' of what to do.

To conclude this section, and the present chapter, I would also like to say a few words about a topic otherwise notoriously associated with Ryle's *The Concept of Mind* – behaviorism. I take it that the fact that I have been able to tell my Rylean story up to this point without touching on this issue already shows that Ryle's thinking about know-how can be seen as sufficiently independent from the general question of the ontology of the mind. Instead, the ontological view of behaviorism becomes important only in the context of the rival ontological claim of Cartesian dualism at

¹⁰ For a recent overview over the debate about the relationship between know-how and propositional knowledge, see Pavese (2016a; 2016b).

which many of the arguments (and most of the polemic) of Ryle's book is directed. For all I have quoted Ryle as saying, and indeed for all he should have said, he can reject dualism on epistemological grounds, but remain neutral, perhaps even quietistic, at the ontological level.

One may complain that this is a non-standard interpretation of Ryle, but this point is unsuccessful as an objection to my project. True, Ryle concludes this book by remarking that “[t]he general trend of this book will undoubtedly, and harmlessly, be stigmatized as ‘behaviorist.’” (Ryle 1949, 308) But this characterization is indeed harmless because Ryle sees himself in the same boat as behaviorists in psychology only with respect to methodological and epistemological issues. That is, he also claims that “theories should be based upon repeatable and publicly checkable observations and experiments.” (Ryle 1949, 308–309) This methodology, however, is entirely neutral with regard to the ontological question if all there is to the mental *just is* what is overtly available.

Such an interpretation of Ryle's texts may be less frequent than the standard assumption. But at the same time, there already is a growing literature on Ryle which argues that he should not be regarded as a behaviorist, at all.¹¹ The account developed here is in very good company.

With these findings, I shall now depart from a closer engagement with Ryle's texts and proceed to a more systematic development of my own Rylean position in chapters 3 and 4. In Part Two, this Rylean account will be defended against a number of problems surrounding the relation between the Rylean concept of know-how as intelligent ability on the one hand and the English expression ‘knows how to’ as well as the role of ability on the other hand. Some of Ryle's views will continue to come up in several places along the way. For example, I will comment on what Ryle may have made of the semantics of ‘knows how to’ in § 7.5, and I will discuss the relationship between know-how and dispositions, and Ryle's position on this question, in § 5.6. Most importantly, I will discuss Ryle's famous regress argument against what he calls ‘the intellectualist legend’ in chapter 9.

¹¹ This includes Park (1994), Tanney (2009), and Stout (2003), among others. For dissent, see Stanley (2011b, 7–11).

Chapter 3

Exercising Know-how

According to the Rylean view I have spelled out so far, the concept of know-how is the concept of a reliable ability to succeed at meeting the norms of an activity in virtue of being guided by an understanding of these norms. This raises the question what it is to understand the norms of an activity and how it is that these intellectual aspects of know-how guide the conduct of competent actors. In chapter 4, I shall offer answers to these questions. But before I can do this, I will have discuss something more fundamental, the question what it is to exercise know-how in the first place. This notion already played a crucial role in § 2.1, where I discussed Ryle's considerations about this issue. But this will have to be substantially refined.

I begin with the notion of automaticity in the exercise of know-how. In § 3.1, I argue that not every exercise of know-how is an intentional action since there are entirely automatic and therefore non-intentional performances which nevertheless qualify as genuine exercises of know-how. Then, in § 3.2, I turn to those exercises of know-how which *are* intentional actions and discuss what it is about an intentional action that one explains when one cites the agent's competence. In § 3.3, I argue that, despite the fact that not every exercise of know-how is an intentional action, the possibility of intentional action is nonetheless crucial for the concept of competence.

In § 3.4, I propose an answer to the pressing question how a performance can be completely automatic and at the same time, *qua* exercise of competence, be guided by norms. This will lead to some fundamental considerations concerning norms, normativity and normative guidance.

Finally, § 3.5 comments on consciousness, awareness, and knowledge of action in the exercise of competences. The closing § 3.6 deals with on the alleged phenomenological support which the automaticity of competences is supposed to lend to anti-intellectualism.

§ 3.1 Automatic Exercises of Competences

Knowing how to do something is knowing how to engage in an activity. But engaging in an activity need not consist in executing an intentional action. I have flagged this point already in § 1.1 and so far, this has not made any difference to the view I have been spelling out. Now, however, I shall elaborate this point more thoroughly and defend it against some objections.

To begin with, I take it to be phenomenologically clear that at least some exercises of skills are not intentional actions. Here is a paradigm case:

Unwelcome Sign

I know how to read and I execute this competence very often. But I do not only execute it intentionally, say, when I read a book which I *want* to read and which read *for reasons*. I also exercise my skill to read when I walk down the street and happen to see a sign with words written on it. I do not need desires or reasons in order to read what it says on such a sign. I might even have clear desires and reasons *not* to read what the sign says – say, because I do not want to pollute my mind with advertisements. But I still *do* read the sign and I read it correctly. I exercise my know-how completely automatically.

As another case in point, take the drawing of inferences. People often infer propositions from other propositions without paying any explicit attention to this. It would be very strange to assume that people always *intentionally* infer, say, that there will be thunder soon, from seeing lightning now. Sometimes, such an inference more or less happens to people, rather than that it is performed as an intentional action. On the other hand, Gilbert Ryle acted intentionally when he argued that what he calls “Descartes’ Myth” is false. His inference from the premises he offers to the conclusion that the view in question is mistaken is something which he draws *voluntarily* and which he has *reasons* for drawing. At least to me, this seems to be a clear case of the intentional drawing of an inference.

These observations about reading and drawing inferences can be supported by considering analogous activities which one may know how to engage in, e.g. reading a map and calculating, which are also paradigmatic examples of know-how (cf. *Ryle’s Range of Cases* on page 14). Sometimes these things are done voluntarily and intentionally. Other times they are performed completely automatically (cf. Schneider 1999). Still, all of these performances are genuine exercises of the respective competences.¹

¹ These cases may appear strange because they seem to involve motionless performances. But this worry is unfounded. It is simply false that every intentional action involves motion. For one thing, Ryle characterizes an “exercise of intelligence” as an “overt or inner performance” (Ryle 1949, 44) and thereby explicitly includes performances

I have deliberately proposed these considerations without presenting a detailed account of the notions of intentionality, intention, and intentional action, all of which are subject to philosophical debate.² Perhaps too boldly, I maintain that *whatever* one's philosophical view of intentional action may be, these cases of *completely* automatic performances do not amount to intentional actions on *any* such view. Such completely automatic acts are not done with the intention and on the basis of the intention of doing so. They are also not performed with a reason and for a reason. But intentional action clearly seems to require at least some criterion along these lines. I shall come back to these points in § 3.2. The phenomenology of practice clearly suggests that not every exercise of know-how is intentional, and that many cases are instead completely automatic.

To clarify, the fact that completely automatic performances are not *done for* reasons is entirely compatible with the fact that there still *are* reasons for performing them. In *Unwelcome Sign* on the facing page, the obvious reason to read the sign in the way I do is that the symbols on the sign mean what they mean. But, again, when I read this sign completely automatically, I do not read it *for* this reason. This is merely a *pro tanto* reason for me to read the sign in this way, if I read it at all, and not a reason to do so in the first place and all things considered. As I construed this case, there are overriding reasons *not* to read the sign.

This allows me to point out that the norms which govern an activity *always* provide *pro tanto* reasons to exercise one's competence in a certain way, provided one has reason to exercise them at all. Just like the norms of reading the language in question always provide a reason to read its symbols in a certain way, the norms of every activity always provide a reason to engage in them in this rather than in another way. Ryle already sees this when he continues the passage already quoted on page 18 as follows:

When a person knows how to do things of a certain sort (e.g., cook omelettes, design dresses or persuade juries), his performance is in some way governed by principles, rules, canons, standards or criteria. (For most purposes it does not matter which we say.) It is always possible in principle, if not in practice, to explain why he tends to succeed, that is, to state the reasons for his actions. It is tautology to say that there is a method in his cleverness. (Ryle 1945a, 8)

without any directly discernible overt manifestations. Examples for such 'inner performances' include coming to believe something on the basis of reasoning or calculating (cf. page 51). For another thing, there are also motionless overt performances. I can keep a door shut by pressing my hand against it, I can stand still, I can refuse to answer a question by keeping quiet, I can stare at something intensely, and so forth. All of these can be performed, both automatically and intentionally, despite the fact that they are motionless (cf. von Wright 1988, 97; Sandis 2012, 326).

² For an overview of these problems, see Wilson & Shpall (2012) and Setiya (2014).

As a further clarification, in calling the latter kinds of performances ‘completely automatic’, I do not intend to suggest a mechanistic view of these aspects of human practice. Instead of the term ‘automatic’, I will also talk of *sheer routine* acts and performances. But the mechanistic metaphor implicit in the term ‘automatic’ is helpful at least in stating very clearly, and better that the term ‘routine’, that the performances in question are to be conceived as *completely* devoid of intentional action.

Furthermore, there is certainly a continuum between completely automatic performances at the one end and maximally reflective and intentional acts at the other end. When I discuss completely automatic performances here, I do not wish to deny this. Methodologically, I try to show that my account holds for *all* exercises of know-how, regardless of the exact degree to which they are automatic or not, by discussing those cases that are indeed completely automatic. *Mutatis mutandis*, my conclusions will then also apply to all other cases which involve a lower degree of automaticity.

To illustrate this point, let me briefly discuss the view of automaticity and habituation proposed by Julia Annas (2011). She argues that there is a crucial distinction between two kinds of automatic or habitualized abilities, and that only one of these deserves the title “practical expertise”, while the other one can only be described as “routine” (Annas 2011, 103). What she has in mind is illustrated by her examples. A person’s drive from home to work, which she performs almost every day, is a case of “habituation that results in mere habit and routine”, where “the reaction to the relevant situation is always the same” (Annas 2011, 102). On the other hand, a skilled pianist playing a piece she has learned to play and practiced very well is a different affair. According to Annas, this is a case of practical expertise which involves “reactions that differentiate among, and are appropriate to, different situations” (Annas 2011, 102).

This notion of practical expertise has many of the features of the notion of know-how or competence that I defend here. For example, it involves an understanding of what Annas calls “the point” of the activity in question and the drive to improve one’s capacity in this respect (Annas 2011, 104–105, 108). This makes for a clear contrast with mere routine and what I call a mere automatized ability.

However, it is not clear in which sense all the individual performances of driving to work are indeed “always the same”, while all the individual acts of playing a certain piece on the piano “differentiate among, and are appropriate to, different situations” (Annas 2011, 102). I take it to be much more plausible to credit both abilities with *some* degree of differential responsiveness to novel situations and with *some* degree of routinely doing

the same all over again. Even on a routine commute, a driver has to respond to today's traffic in a way which is different from the way she responded to yesterday's traffic. And even when playing a very demanding piece for the first time, a pianist can rely on the fact that at least some of her acts are routine responses. It is precisely on the basis of these routines that she can strive to play as best as she can and even to improve her capacities further. As I have just pointed out, there is a continuum of cases with completely automatic performances at the one extreme and less and less automatic and more and more reflective and intentional acts from there on.

Further, Annas thinks that the possibility of the articulation of reasons and explanations for one's practice makes for a further distinction between routine and practical expertise. In cases of routine, she writes, "I have to detach myself from my ongoing practice" in order to give reasons and explanations, while in cases of practical expertise, the person's "thoughts have effaced themselves, but they have not entirely evaporated" and "they will become available without her having to detach herself from the activity" (Annas 2011, 111). Again, however, I contend that this distinction is better described as a matter of degree than as a substantial distinction marking off entirely different phenomena. And this is what Annas says for herself when she mentions the corresponding question of the ability to convey one's reasons, explanations and skills to somebody else (cf. Annas 2011, 111).

Accordingly, I take Annas' examples to illustrate the spectrum of the degrees of automaticity and routine within the broader notion of know-how or competence rather than to mark off this concept from something else entirely. As long as a person is able to exert direct control, or even a form of indirect control that I will discuss in § 3.4, her performances can be understood as exercises of genuine competences rather than of mere abilities.

In spite of the unequivocal phenomenological datum of completely automatic exercises of know-how, some philosophers explicitly or implicitly hold that know-how is a capacity for intentional action and *only* for intentional action. For example, Paul Snowdon comments on Ryle's example of the skill to draw correct inferences as a case of know-how as follows:

[T]alk of knowing how to F is appropriate only where F represents a genuine action. [...] [F]orming a belief (by inference) is not an action, but is, rather, a cognitive development or change. [...] If inference is not an action, then there is no such thing as knowing how to do it. (Snowdon 2011, 70)

Similarly, Katherine Hawley holds:

[S]uccess cannot amount to knowledge-how unless intentional action is involved. (Hawley 2003, 26)

And Jason Stanley and Timothy Williamson also restrict the exercise of know-how to intentional action alone. Discussing Ryle's famous anti-intellectualist argument which I shall defend in chapter 9, they write:

(1) If one Fs, one employs knowledge how to F.

[...] Premise (1) is true only when the range of actions is restricted to *intentional actions*. (Stanley & Williamson 2001, 413, 415)

In short, the view expressed here is that every exercise of know-how must be an intentional action. But on the basis of my observations and arguments above, I would claim to have shown that this view is false – indeed, that it *must* be false if our philosophical views are to be faithful to the phenomena.

But if the situation is indeed as obvious as I would like to have it, why should this mistaken view become so widespread? A closer consideration of the passages I have just quoted reveals that these philosophers are led to believe that every exercise of know-how is intentional because they want to distinguish know-how from what I have called a mere ability:

[T]alk of knowing how to F is appropriate only where F represents a genuine action. We do not know how to digest food, faint or sweat. We simply are able to and do these things (in a nonactional sense of 'do'). (Snowdon 2011, 70)

Digesting food is not something done intentionally, and that is why it is not a manifestation of knowledge-how. Similarly, Hannah did not intentionally win the lottery, although she doubtless hoped to win it. So, for premise (1) to be true, the range of actions under consideration must be restricted to intentional actions, or perhaps even a proper subset thereof. (Stanley & Williamson 2001, 415)

[S]uccess cannot amount to knowledge-how unless intentional action is involved. We don't describe ourselves as knowing how to produce white blood cells. (Hawley 2003, 26)

These cases can be classified along the lines I have proposed in chapter 1 – indeed, I have already mentioned some of them. *Digestion* on page 34 is a mere ability – reliable yet unintelligent success. So is the ability to produce white blood cells. Maybe sweating and fainting can be classified similarly, provided one sweats or faints with sufficient reliability in situations of some specified type. If not, they will go into the same box as *Lottery* on page 25 and cases of beginner's luck. These are examples of things one *can* do, but in the sense of mere possibility as opposed to reliable ability.

Thus, Snowdon, Hawley and Stanley & Williamson argue that if mere abilities are not cases of know-how, and if mere abilities are not cases of intentional action, then every exercise of know-how must involve intentional action. But this can now be revealed as a fallacy. Given the premises in

play, it remains perfectly possible to account for the distinction between know-how and mere ability in terms of something other than intentional action. And indeed, this is what I supported in detail in § 1.5. Unlike mere ability, know-how involves success *in virtue of normative guidance* – an understanding of and a guidance by the norms of the activity in question. So far, nothing whatsoever must be said about intentionality.

I conclude that Snowdon, Hawley and Stanley & Williamson are mistaken. It is false that every exercise of skill is an intentional action. However, I shall come back to the relationship between know-how and intentional action in § 3.4 and add a more friendly amendment to these criticisms.

Unfortunately, the view defended here gives rise to a problem – a problem which the philosophers just criticized would be likely to point out now. If normative guidance is what distinguishes know-how from mere ability, and if exercising know-how can be automatic, how can one be automatically guided by norms? It seems rather clear how I can be in control of my living up to certain norms when I intentionally try to act in accordance with these norms. But how can I be in control of this when I do not intentionally try to act in accordance with these norms? When my conformity to certain norms is an automatism, how can it be a competence?

This worry is a serious and important. But before giving a direct answer in § 3.4, I would like to show how this tension shows up in Ryle's texts.

Ryle initially makes the same mistake as Snowdon, Hawley, and Stanley & Williamson. When he tries to distinguish exercising know-how from “doing something by pure or blind habit”, he appeals to the idea that the latter means doing something “automatically” (Ryle 1949, 42). This suggests that exercising know-how cannot be automatic at all. Ryle goes on:

The distinction between habits and intelligent capacities can be illustrated by reference to the parallel distinction between the methods used for inculcating the two sorts of second nature. We build up habits by drill, but we build up intelligent capacities by training. Drill (or conditioning) consists in the imposition of repetitions. [...] The practices are not learned until the pupil's responses to his cues are automatic [...]. Training, on the other hand, though it embodies plenty of sheer drill, does not consist of drill. It involves the stimulation by criticism and example of the pupil's own judgement. [...] Drill dispenses with intelligence, training develops it. (Ryle 1949, 42)

Ryle thinks that while blind, automatic habits are acquired by sheer drill and *independently* of an understanding of the norms of the activities in question, intelligent capacities or competences are acquired by training, which *does* involve the development of such an understanding (cf. Ryle 1945a, 15). This makes intelligence and automaticity appear nearly incompatible.

However, the passage just quoted already hints at a way to resolve this tension. This idea can be traced back to Ryle's claim that training "does not consist of drill", but that it "embodies plenty of sheer drill" (Ryle 1949, 42). If this is true, then know-how involves automatic habits. If training is what generates competences, if training "embodies plenty of sheer drill" (Ryle 1949, 42) and if drill is what generates automatic habits, it clearly follows that competences can 'embody plenty of' automatic habit, as well. And indeed, this is what Ryle claims explicitly. In an earlier statement of these considerations, he says of a soldier in training:

He acquires not a habit but a skill (though naturally skills contain habits). (Ryle 1945a, 15)

Thus, and despite the apparent tension, Ryle acknowledges that there must be *some* sense in which automatic habits and intelligent capacities are compatible. Indeed, they *must* be compatible if automatic habits are supposed to be part of skills. Given his careful choice of words, Ryle appears to be aware of this problem. For whenever he contrasts intelligent skill with mere unintelligent ability, he avoids to use the plain notion of 'habit', but only talks of 'pure' habit or 'blind' habit. And he avoids to use the plain notion of 'drill', but only talks of 'sheer' drill (cf. Ryle 1949, 42). However, Ryle does not say how exactly habits *can* be part of competences. In § 3.4, I will present a solution to this problem that is very much in Ryle's spirit.

§ 3.2 How Know-how Explains Action

In § 1.5 and § 3.1, I have argued that the distinction between know-how and mere ability is independent of the question of intentional action. What distinguishes know-how from mere ability is normative guidance. But what, then, *is* the relationship between somebody's competence in an activity on the one hand and her intentional actions of performing that activity on the other hand? How, in short, are know-how and intentionality related?

In this section, and in § 3.3, I will develop answers to this question. First, I shall reconsider what it is exactly that we explain about a performance when we explain it in terms of the person's know-how. I shall disentangle different explananda and point out what kind of explanation the concept of know-how is aimed at. Then, in § 3.3, I discuss the crucial connection which still remains between know-how and intentional action.

In order to address the issue of how know-how explains action, I must come back to something very important – the nature of the explanandum for which know-how was introduced as an explanans. I have argued that

the concept of know-how helps us understand a crucial feature of normative practice (cf. § 1.1). If somebody knows how to do something, then this is what explains her *intelligence* in doing it – the fact that her reliable success (cf. § 1.4) is due to her understanding of and guidance by the norms against which the successes in the activity in question are measured (cf. § 1.5).

What *kind* of explanation does the concept of know-how offer? Suppose that a skilled archer shoots a couple of arrows at a target and does extremely well: She hits the bull's eye with every single arrow. If we seek to explain these performances, we may ask a very general question such as this:

- (1) How come that the archer shot her arrows well?

I take it that this question is the most general way to express a request for an explanation of these performances in ordinary English. But we can distinguish two more specific ways of asking for such an explanation, which differ in the kind of explanation they are after:

- (2) (a) Why *did* the archer shoot her arrows well?
(b) Why was it even *possible* for the archer to shoot her arrows well?

I take it to be obvious that an explanation which cites the archer's know-how, her skill in archery, does nothing to answer question (2 a), but provides a perfect answer to question (2 b). Know-how, in short, is not what explains why an act *happens*. Rather, it explains why a certain *kind* of act – an intelligent, that is, successful because norm-guided act – is possible in the first place. According to that explanation, the act is possible because it is an exercise of the archer's competence – her reliable ability to succeed in virtue of her understanding of the relevant norms.³

By contrast, when talking about intentional action and its explanation, what we have in mind is typically question (2 a), the question *why* the archer shot her arrows well. Again, I cannot discuss the concepts of intention and intentional action here in detail (cf. footnote 2 on page 79). Still, I contend that, on any plausible account of these notions, the explanation of an act as an intentional action would concern (2 a) rather than (2 b). Such an explanation would refer to the archer's desires or intentions – say, to practice her archery skills, to impress her friend by performing well, or to win a competition. And it would also refer to her reasons or beliefs – say, that this is a good method to practice, that this is indeed a means to impress her friend, or that the competition actually is about hitting the bull's eye.

³ This explanation is probably closely related to Aristotle's notion of a formal cause (cf. e.g. *Physics* 195a). But I cannot address these questions here in proper detail (cf. Falcon 2015; Henning 2009). For a discussion of the deep Aristotelianism within Ryle's work, particularly in *The Concept of Mind*, see Stout (2003) and Wiggins (2012).

To be clear, there is also a way of answering question (2 a) which does *not* lead to an explanation of the performance in question as an intentional action. This question can also be understood to ask for different kinds of explananda. For example, Constantine Sandis distinguishes between three elements which might be the targets of questions of the form “Why did A move her body?” like (2 a) (Sandis 2012, 333):

- (3) (a) Why A’s body moved
- (b) Why A’s action of moving her body occurred
- (c) Why A moved her body

Subtle differences notwithstanding, I take it to be very clear that explanations of (3 a) or (3 b) would be merely explanations of why an event occurred which happens to have been a moving of a body. By contrast, only an explanation of (3 c) would be an explanation of the performance in question *as an intentional action*. While the explanans of (3 a) or (3 b) can be a purely causal story on the sub-personal level, the explanans of (3 c) will have to include elements from the personal level – desires, intentions, beliefs or reasons. Of course, it is highly contested how these four elements are connected to one another and how they relate to the notion of causation which also plays a role in explaining (3 a) or (3 b) (cf. Wilson & Shpall 2012). However, these complications between (3 a), (3 b) and (3 c) all only concern the first question I have distinguished above, question (2 a).

I have argued that the concept of know-how is involved in question (2 b), but not in answering question (2 a). It does not explain why a performance occurred, in any of the senses just considered. Instead, it explains why it was possible for such a norm-guided performance to occur in the first place.

However, it is important to see that the explanandum is not the *mere* possibility of a performance, but rather the possibility of an *intelligent* performance. To cite an archer’s know-how in order to explain her successful shots is not merely to say that it was possible for her to make these shots. Given that this has already happened, it is clear that it must have been possible. The interesting part of the explanation is not that this that it was an intelligent performance. This involves two elements.

First, the explanation includes that the act was performed with sufficient reliability (cf. § 1.4). This rules out lucky coincidences because it entails that the performance would have been equally or close-to equally successful if the circumstances had been slightly different. In other words, the explanation entails that the performance was not merely possible, but that it was, in some relevant sense, *probable*. Given the archer’s competence, her success at shooting the arrows was not merely to a possibility. It was to be *expected*.

Second, the explanation includes that the act was not performed as an exercise of a *mere* ability. It was not to be expected for *any* reason – such as a secretly installed magnet which guided the arrowheads⁴ – but for the *right* reason. That is, it was to be expected in virtue of the person's *understanding* of the relevant activity which *guided* her in her practical conduct.

In sum, to explain a performance with appeal to know-how is *not* to explain *why* the person performs this act, answering question (2 a) with an intentional action explanation. Rather, it is to explain why it was even possible for that person to perform this act *as an intelligent performance*.

This is faithful to Ryle who explicitly states that different kinds of explanation are required in order to explain intelligent practice:

There is the one activity, but it is one susceptible of and requiring more than one kind of explanatory description. (Ryle 1949, 50)

And earlier, he expresses the specific distinction I have laid out in this section with the following analogy:

The explanation is not of the type 'the glass broke because a stone hit it', but more nearly of the different type 'the glass broke when the stone hit it, because it was brittle'. (Ryle 1949, 49)

In this case, the explanation under consideration does not give a cause of the breaking of the glass. Analogously, the explanation of the archer's success in terms of her competence does not answer the first question, by giving reasons or causes. Instead, the explanation in Ryle's example explains how the glass' being hit by a stone *could* be a cause for the breaking of the glass in the first place. It explains the glass' breaking by appealing to its brittleness – that is, as an actualization of its disposition to break when hit by a stone. Analogously, the competence-based explanation of the archer's success explains how these things could lead to the archer's intelligent success in the first place.

I will come back to this important tie between know-how and dispositions in § 5.5 and § 5.6. For now, I conclude that the concept of know-how does a very different explanatory job than the concept of intentional action.

Furthermore, there are also explanatory questions which neither of these notions manages to answer. Consider the example that Lionel Messi scores an amazing goal. One part of the explanation of this performance will be that he intended to do so, that he had reasons and desires to do so, and that he had beliefs about how he can manage to do it. A second part of

⁴ I will discuss cases like these at length in § 5.4, where I categorize them as 'practical luck' in order to distinguish them clearly from other families of puzzle cases.

the explanation of Messi's scoring that amazing goal will be that he had the competence to do so in the first place, in the absence of which this performance would not have been possible at all. But even given these elements, it may remain an open question why, given all of this, Messi managed to perform so exceptionally well right here. After all, even he misses shots, makes mistakes, and so forth. In such a case, I take it that we are looking for further explanatory elements such as favorable circumstances, recent learning and adjusting to things like the play of his opponents or to the weather, and even sheer luck.

Thus, while the notion of know-how does a different explanatory job than the notion of intentionality, even when these are combined, there are further aspects of somebody's performances requiring further circumstantial explanation. The general demand for an explanation of some action, as in (1) on page 85, can be answered with an explanation appealing to intentional action – answering question (2 a) – and it can, alternatively or additionally, be answered with an explanation appealing to competence – answering question (2 b). But even if one has given perfect answers to these questions, the initial question (1) may still require further elaboration.

§ 3.3 Why Know-how Needs Intentional Action

So far, I have had only negative things to say about the relationship between know-how and intentional action. In particular, § 3.1 has argued that it is simply false that an exercise of know-how is always an intentional action. And § 3.2 has shown that the concept of know-how and the concept of intentionality explain entirely different things about an intelligent action. One might take this to imply that intentionality and know-how are completely independent from one another. But this would be a mistake. Instead, I will now show how competence and intentional action are importantly intertwined, and why and where know-how needs intentional action, after all.

To begin with, there are many kinds of intentional actions which only the skilled can perform (cf. *Ryle's Range of Cases* on page 14). Without the relevant know-how, it is impossible to do things like calculating, arguing, and speaking. All intentional actions which *constitute* exercises of intelligence can be pursued in practice and explained in theory only on the basis of know-how. Unlike scoring a bull's eye, which may or may not be an intelligent action, performances like calculating or inferring are essentially intelligent. Without competences to perform them, nothing will allow one to perform such an essentially intelligent act.

These considerations also permit to comment on the view that the concept of intentional action *itself* has what is often called a ‘skill’ or ‘control’-sense. For example, results from empirical psychological studies support a close correlation between the description of people’s actions as intentional on the one hand and as skilled, controlled or at least not accidental on the other hand.⁵ Furthermore, there is a somewhat neglected philosophical explication of intentional action by Harry Frankfurt (1978)⁶ which appeals to, roughly, guidance through actual and possible voluntary interference rather than to the causal history of the performance in question. This idea is closely related to the notion of responsible control which will play a crucial role in my explanation of intellectual guidance in § 4.4.⁷ But I cannot comment on all of these subtleties of the everyday usage of ‘intentional’ and of the debate on the concept of intentional action. And in any case, the notion of skilled control which is relevant for the concept of know-how is more specific than whatever is involved in the explanation of intentional action. I will come back to this point at the end of this section.

Still, the account of know-how developed in this book directly predicts a close connection with intentional action. When we are dealing with intentional actions that are intelligent performances – that is, acts guided by an understanding of what it takes to do well in a given activity –, then skill, control and non-accidentality are all necessary.

There is an even more important connection between competence and intentional action. Despite the fact that some exercises of skills are unintentional and completely automatic performances, it is still true that *every* competence *can* be exercised intentionally. Conversely, no competence can *only* be performed automatically and *never* intentionally.

Let me point out that this constitutes an important concession in my dismissal of the claim that every exercise of know-how is an intentional action. In § 3.1, I have discussed the phenomenology of automatic exercises of know-how and criticized Snowdon (2011), Hawley (2003) and Stanley & Williamson (2001) for denying these phenomena in a failed attempt to distinguish intelligent know-how from mere ability. I can now add that these philosophers are certainly correct to insist on a closer conceptual relationship between know-how and intentional action. However, it is not true

⁵ See, for example, Malle & Knobe (1997), Cova *et al.* (2012) and Nadelhoffer (2005).

⁶ Kent Bach seems to have proposed a similar view at the same time (cf. Bach 1978).

⁷ Ezio Di Nucci has presented an interesting discussion and defense of this view (2011a; 2011b; 2013). Di Nucci argues that this is the best way to account for involuntary, unreflective and automatic acts, and that it supersedes causalist views of action. I am very sympathetic with this position, but I shall bracket it, at least partly, as my present topic is know-how and its exercise rather than intentional action in general.

that *every* exercise of any given competence must be an intentional action. Rather, the possibilities to exercise of any given skill must *include* intentional actions. Along these lines, it turns out that intentional action plays an important role in distinguishing competence from mere ability after all.

The fact that every competence *can* be exercised intentionally, even if some of its exercises may be entirely automatic, also allows me to respond to a more subtle way of connecting know-how and intentional action, proposed by John Bengson and Marc Moffett. They write:

[K]now-how becomes a possibility only in cases of intentional, though possibly rote, action. For instance, one can know how to shift gears on a car, but one cannot know how to jerk one's knee reflexively[.] (Bengson & Moffett 2007, 45)

This statement can be taken to express precisely what I have been advocating here. It is possible for somebody to know-how to do something only if it is also possible for them to intentionally engage in this activity. However, one may also read this statement as another way of arguing that every exercise of know-how is an intentional action, while accomodating the fact that some intentional actions are unreflective routines. This idea fits in nicely with my discussion of the spectrum of degrees of reflectiveness or automaticity, where even clearly intentional action may involve a lot of automaticity (cf. § 3.1). Still, some exercises of competences are *entirely* automatic and therefore *not* intentional at all. This also explains why it is correct to say that “one cannot know how to jerk one's knee reflexively” (Bengson & Moffett 2007, 45). Maybe one can know how to jerk one's knee, but one cannot know how to do so *reflexively* since this description already involves the fact that the action is completely automatic.

But so far, I have merely stated the crucial claim that every competence can be exercised intentionally. I shall now turn to the task of justifying it.

I shall prove that there is no skill for which an intentional exercise is impossible by showing that the existence of all automatic routines in the exercise of competences depends on the possibility to intentionally exercise them. In other words, I shall argue that no competence can *only* be exercised automatically because the very existence of such automatisms depends on the fact that these competences can also be exercised intentionally.

My argument for this view starts with the simple observation that every routine to perform in a certain way – just like every automatism even of automata – comes into existence at some point in time. In particular, automatisms in the performance of intelligent acts – that is, routines in exercising skills – are acquired by learning. As I have argued in § 2.3, every competence is learned and only mere abilities can be innate.

But how is it that a routine in exercising a skill is acquired? I contend that the only way to do so is to start to *practice* engaging in the activity in question, and to do so *intentionally*. After all, one of the aims of the performances in question just *is* to improve one's capacities and to develop routines. I take it to be a commonplace datum of phenomenology, and indeed a well-established result of cognitive science,⁸ that skills are improved by practice. But at least in the beginning, the engaging in an activity cannot be understood as a *practicing* of it unless it is an intentional action – an action intentionally aimed at meeting the norms of the activity in question and at acquiring automatic routines of doing so.

Could this be true and still fail to entail my point? Could the acquisition of automatic skills require their intentional exercise, while at least some people who have acquired an automatic skill thereby *lose* the capacity to exercise it intentionally? This seems to be possible in principle. But it is hard to imagine an example which fits this description.

Fortunately, there is also a systematic reason which rules out such cases: The fact that the acquisition of a competence requires its intentional exercise need only be considered in tandem with the further insight that any given competence can always be improved – by practice and in other ways – and that the distinction between acquiring a new skill and improving an existing one depends on the exact conception of the activity. In other words, as I have spelled out in § 1.7, there is no precise boundary as to whether an improvement is described as the acquisition of a whole new piece of know-how – say, the competence to beat a chess grandmaster – or as an improvement of an existing competence – the competence to play chess. It follows that if an automatism can only be *acquired* by intentional rather than purely automatic exercise, then it can also only be *improved* by intentional exercise. And if it is true that no competence is ever perfect, i.e. that every competence can be improved further as suggested in § 2.3, then this already shows that every competence can be exercised intentionally. One may improve a competence by later assessing an entirely automatic performance, but one can also intentionally exercise an otherwise automatized competence just in order to improve it. Wanting to improve a competence and exercising it because of this desire makes for a clearly intentional action.

Thus, the claim that every skill can be exercised intentionally, even if not all exercises are intentional since some are completely automatic, is a direct consequence of my account of know-how. I will come back to this in my discussion of the opacity of know-how ascriptions in § 6.1 and § 6.2.

⁸ In *The MIT Encyclopedia of the Cognitive Sciences*, it says that “[a]utomatic processes generally develop slowly, with practice over hundreds of trials” (1999, 63).

To conclude my account of the relationship between intentional action and know-how, I would like to comment on one last question. While I have argued that not all exercises of know-how are intentional actions, even if every piece of know-how can be exercised intentionally, we should also look at the reverse of this statement – the claim that if somebody does something intentionally, then she exercises her knowledge how to do so. Many philosophers subscribe to this claim, including Stanley & Williamson (2001) (cf. page 82). However, I think that this claim is false.

As I have argued in § 1.5, basic actions such as opening one's eyes are not activities which one may know how to perform. Instead, capacities to perform basic actions are what I called mere abilities as opposed to competences because such acts do not rely on guidance by an understanding of what it takes to do them well. Instead, one simply acquires such capacities at some point and is then able to exercise them. The boundary between the basic and the non-basic is admittedly difficult to draw (cf. § 1.7). Still, performing basic actions is clearly something one can do intentionally. I can certainly open my eyes intentionally – that is, with reasons for doing so, with the desire to do so, with the explicit plan to do so at the right moment, and so on. And I can also be in control of such basic actions in Frankfurt's sense that, roughly, I oversee and voluntarily influence what I do (cf. Frankfurt 1978). Thus, it is false that if one does something intentionally, one exercises one's know-how to do so. If one intentionally performs an automatic action, one exercises a mere ability.

What I have shown is this. Not every exercise of a competence is an intentional action. But every competence can be exercised intentionally. And not every intentional action is an exercise of a competence.

§ 3.4 What Ought to Happen

In § 3.1, I have argued that it is a criterion of adequacy for any view of know-how that it should be able to account for the fact that people sometimes exercise their skills automatically and unintentionally. However, this seems to create a tension with the view that know-how involves normative guidance. How can something which I do completely automatically be guided by understanding of the norms which govern such acts?

In this section, I will propose an answer this question. My core idea is that the norms which one must understand and be guided by in order to exercise know-how are norms of two different forms. These norms are not only about what one ought to *do*, but also about what one ought to *be* like.

Consider again the case of *Unwelcome Sign* on page 78. Despite the fact that my automatic exercises of my competence to read are not intentional actions and therefore not guided by my understanding of how I ought to *act*, these performances are still guided by my understanding of norms of what I ought to *be* like. In particular, the orthography of German requires that a sequence of letters such as “Parken verboten” be read to mean that one may not park at the corresponding location. There are norms for what one ought to come to believe about the significance of written signs. One can read in conformity with these norms without reading intentionally. And this can nevertheless be a case of guidance by these norms as opposed to mere conformity to them. If one has acquired this automatism in the light of an understanding that *these* are the norms to comply with, then even an automatic performance counts as genuinely guided by these norms. It is because one ought to *be* somebody who reads German in this way that I have acquired the automatic competence to follow these rules.

With this, I have sketched the view I shall now defend in more detail. This defense will consist of two steps. First, I discuss the distinction between the two kinds of norms just outlined – between how one ought to act and what one ought to be like. Second, I show how guidance by an understanding of what one ought to be like can be manifest in automatic acts.

To begin with, however, I should finally reveal that the distinction behind this account is not an invention of mine. While several philosophers have proposed similar distinctions, I shall mainly rely on Wilfrid Sellars. In his words, we must distinguish between “rules of action” and “rules of criticism” or, equivalently, between “ought-to-be” and “ought-to-do” (cf. Sellars 1968, 75–77; Sellars 1969, 506–513).

Sellars mentions this distinction after presenting essentially the same considerations I have developed in § 1.5. That is, Sellars distinguishes mere conformity to norms from conformity in virtue of normative guidance.

In order to introduce the distinction between “ought-to-do” and “ought-to-be” norms, Sellars mentions the following cases of “ought-to-be” norms:

[O]ne ought to feel gratitude for benefits received, though feeling gratitude is not something which one *does*, save in that broad sense in which anything expressed in the active voice is a doing. (Sellars 1968, 76)

One ought to feel sympathy for bereaved people. (Sellars 1969, 509)

These cases should make the distinction in question pretty clear. Feeling gratitude for benefits received or feeling sympathy for bereaved people is not something which one ought to *do*, simply because they are not the kind of thing one *can* possibly do – they are not actions. However, one ought

to be somebody who feels gratitude for benefits received, and one ought to be somebody who feels sympathy for bereaved people. These norms differ from related norms such as the norm to say “Thank you” or to offer one’s condolences in these situations. The former are “ought-to-be”, the latter are “ought-to-do”, and both are clearly species of the genus of norms.

But this clarification only pushes the problem back one step. True, one might say, now we have pointed out that there are “ought-to-be” norms that one does not follow by acting intentionally. But the question was how this should be compatible with the claim that knowing how to meet such norms involves that one’s performances are guided by those norms in the sense that one possesses an *understanding* of the norms and meets them in virtue of that understanding. Labelling a phenomenon is not explaining it.

This is where the second step of my Sellarsian proposal comes in. In a characteristically dense passage already quoted partially above, Sellars gives the following answer to this problem:

[O]ught-to-be’s imply (with additional premises) ought-to-do’s, and ought-to-do’s imply ought-to-be’s. [...] [O]ne ought to feel gratitude for benefits received, though feeling gratitude is not something which one *does*, save in that broad sense in which anything expressed in the active voice is a doing. [...] One ought, however, to criticize (an action proper) oneself for not feeling gratitude and to take steps (again an action proper) to improve one’s character. (Sellars 1968, 76)

Thus, “ought-to-be” norms and “ought-to-do” norms are intimately interrelated in both directions. I shall unpack this idea a little more slowly.

On the one hand, the “ought-to-do” norm that one ought to express gratitude after receiving benefits implies the “ought-to-be” norm that one ought to be such that one expresses gratitude after receiving a benefit. It would be strange to say that somebody ought to do something, and to deny that she ought to be such that she does so. Generally, if it is true that one ought to A in circumstances C, one ought to be a person who As in C.

On the other hand, the “ought-to-be” norm that one ought to feel gratitude for benefits received implies several other and more interesting “ought-to-do” norms – that one ought to criticize oneself for not feeling gratitude for benefits received, that one ought to shape one’s character such that one *does* feel gratitude for benefits received, that one ought to criticize and shape other people’s characters in the same way, and so forth. Generally, if it is true that one ought to be a person who As in C, then one ought to make it the case that one becomes or remains a person who As in C.⁹

⁹ I shall come back to these principles later in this section, on page 97.

Along these lines, we can finally come to grips with the idea that the act of feeling gratitude after receiving benefits is guided by a relevant norm. Somebody can conform to the norm in virtue of being guided by an understanding of it. Such a guidance by an “ought-to-be” norm consists in the fact that she has been guided by at least some of the “ought-to-do” norms which are entailed by that “ought-to-be” norm. She has shaped or participated in shaping her character in such a way that she now often automatically meets the norm in question. Thus, being guided by “ought-to-be” norms is perfectly intelligible. It requires something which merely conforming to them without this normative guidance does not involve. That is, it requires understanding the “ought-to-be” norm and being guided by at least some of those “ought-to-do” norms which the former norm entails.

So far, it is not entirely clear how this applies to the case of know-how. Intuitively, it appears obvious that feeling gratitude and feeling sympathy are not the kinds of things one may know how to do.¹⁰ And indeed, this is what the account defended so far predicts. As shown in § 3.3, it is necessary for competence that it be possible to intentionally exercise it, even if some of its exercises are automatic (cf. § 3.1). Feeling gratitude is not the sort of thing which one may know how to do because one cannot *intentionally* feel gratitude. At most, one may have the mere ability to do so.¹¹

Thus, it remains to be seen how the idea of normative guidance by “ought-to-be” norms applies to cases of genuine know-how. I began with examples where somebody possesses a genuine competence, but exercises it automatically. In § 3.1, I argued that knowing how to read and how to draw inferences are at least sometimes nonintentional, but rather entirely automatic. Sellars himself mentions a further example which parallels these Rylean cases, namely the following “ought-to-be” norm:

(Ceteris paribus) one ought to respond to red objects in sunlight by uttering or being disposed to utter ‘this is red.’ (Sellars 1969, 511)

For Sellars, norms like these play an important role in explaining the very contents and the normativity of thought and talk in general. In particular, they are crucial for his project to explain the nature of thought in terms of overt linguistic behavior (cf. e.g. Sellars 1968; Sellars 1969). By contrast,

¹⁰ There may be philosophers who would welcome the view that feeling gratitude is something one can be skilled at, arguing that experiencing emotions is itself a form of agency. For example, Nico Frijda (1986) and Brian Parkinson (1995) have proposed conceptions of emotions which tie them more closely to agency. But these questions should be settled independently from the account of know-how I am advertising here.

¹¹ I say ‘at most’ because I can leave open the question whether things like feeling gratitude allow for genuine abilities or only for something weaker, say, a mere disposition.

my own use of this example will be entirely independent from these issues. What I draw from Sellars' view is only one element – the fact that the norm in question is an “ought-to-be” norm governing the use of the term ‘red’.

But this is already enough to establish an example of know-how which involves guidance by an “ought-to-be” norm. Using the term ‘red’ is an activity governed by norms – among other things, by the norm Sellars quotes. That norm – that, *ceteris paribus*, one ought to respond to red objects in sunlight by uttering or being disposed to utter ‘this is red’ – is an “ought-to-be” norm since becoming disposed to utter ‘this is red’, or coming to believe that this is so,¹² is not something one does intentionally. Further, successfully engaging in the activity of using the term ‘red’ requires knowledge how to use this term. And *qua* know-how, this involves the understanding of, and the normative guidance by, the norms governing that activity. And among these is the “ought-to-be” norm mentioned above.

Thus, employing the term ‘red’ is an instance of know-how which involves the understanding of, and the normative guidance by, “ought-to-be” norms.

I propose to tell the same story whenever a piece of know-how is executed automatically and unintentionally. Let me return to *Unwelcome Sign* on page 78 again. When I automatically read a sign in the street, I do not act intentionally. But my reading performance is guided by the same norms which also play a role when I do read intentionally. I am guided by the “ought-to-be” norms of language and orthography. For example, one ought to be somebody who responds to the written letters “Parken verboten” by coming to believe that one may not park there. My automatic response is still guided by these norms because I understand them and because I participate in the continuing process of becoming and remaining somebody who automatically conforms to them.

I conclude that the exercises of know-how can be both automatic and guided by an understanding of norms. People are in control of such automatic performances because they can intentionally shape these automatisms. Indirect control of individual automatic acts in virtue of direct control of the automatisms which produce these acts is a sufficiently rich form of control to allow for normative guidance. In chapter 4, I will build on these considerations and offer a full account of the notion of normative guidance.

It may seem that these considerations still remain somewhat alien to Ryle’s view. But even if Ryle does not talk about the distinction between “ought-to-be” and “ought-to-do” norms, he clearly endorses Sellars’ expla-

¹² I shall remain neutral on the question if a belief may just be a disposition which manifests itself in such utterances and in other ways (cf. Schwitzgebel 2002). However, this option will also play a role later, in § 9.5.

nation of normative guidance as, among other things, following the norms to correct mistakes and shortcomings and to shape one's character appropriately. This has already been foreshadowed in § 2.4, where I discussed how continuous learning is essential for know-how. The point is further substantiated in a passage already quoted on page 62 and earlier:

To be intelligent is not merely to satisfy criteria, but to apply them; to regulate one's actions and not merely to be well-regulated. A person's performance is described as careful or skilful, if in his operations he is ready to detect and correct lapses, to repeat and improve on successes, to profit from the examples of others and so forth. He applies criteria in performing critically, that is, in trying to get things right. (Ryle 1949, 29)

This characterization resonates with Sellars' explication of how one may be guided by an "ought-to-be" norm *in virtue of* being guided by "ought-to-do" norms. When one performs completely automatically, but is nevertheless "ready to detect and correct lapses, to repeat and improve on successes, to profit from the examples of others and so forth" (Ryle 1949, 29), then it is clearly the "ought-to-be" norms of one's automatic acts which one tries to meet by acting in accordance with the "ought-to-do" norms these entail.

To conclude this section, I would now like to apply these considerations to my account of the identity criteria of activities presented in § 1.2. According to my proposal, an activity can be individuated in terms of the weighed set of norms which govern it. Now, I have discussed two different kinds of norms – "ought-to-do" norms and "ought-to-be" norms – and their interrelations. The result of these considerations can be stated as follows:

- (A) For any activity A, there is a set of norms which govern A-ing, where this set includes, for each relevant specific act of type X_i and for each relevant specific set of circumstances C_i , norms of the forms (A 1–3).
 - (A 1) In A-ing, one ought to perform X_i in C_i .
 - (A 2) In A-ing, one ought be such that one performs X_i in C_i .
 - (A 3) In A-ing, one ought to make it the case that one performs X_i in C_i .

As I have already indicated, these distinct kinds of norms entail each other. First, each norm of the form (A 1) entails the corresponding norm of the form (A 2) – each norm according to which one ought to do something in a certain situation entails the norm that one ought to be such that one does so in that situation. Second, each norm of the form (A 2) entails the corresponding norm of the form (A 3) – each norm according to which one ought to be such that one performs a certain act in a certain situation entails the norm that one ought to make it the case that one does so.

When it comes to completely automatic acts, norms of the form (A 3) play the most important role. These “ought-to-do” norms account for the notion of guidance by the underlying “ought-to-be” norms of the form (A 2). However, whenever the relevant acts can be performed intentionally, it is also important to see that each norm of the form (A 3) entails the corresponding norm of the form (A 1) – each norm according to which one ought to make it the case that one performs a certain act in a certain situation entails the norm that one ought to perform this act in that situation. For intentionally doing something is simply the most straightforward way for one to make it the case that one does this.

These relations of entailment between individual norms rely only on considerations about their form. They are conceptually necessary. Accordingly, a weighted set of norms which governs an activity must be, to put it technically, closed under this entailment relation. That is, any set of norms which includes a norm of the form (A 1) thereby also includes the corresponding norm of the form (A 2). Any set of norms which includes a norm of the form (A 2) thereby also includes the corresponding norm of the form (A 3). And any set of norms which includes a norm of the form (A 3) thereby also includes the corresponding norm of the form (A 1), even if this norm will be cancelled wherever the relevant act can never be performed intentionally, on grounds of the principle that ‘ought’ entails ‘can’.

Thus, every activity one may know how to engage in relies on a set of norms, where each relevant pair of specific acts X_i and specific sets of circumstances C_i figures in three distinct, but conceptually interrelated norms (A 1–3). Therefore, this is the set of norms with which somebody who is able to engage in that activity conforms. By the same token, this is the set of norms which guides the practical conduct of somebody who exercises not only her mere ability, but her knowledge how to engage in that activity. This will be spelled out in detail in chapter 4, particularly in § 4.4.

§ 3.5 Consciousness and Knowledge of Action

In § 3.1, I have argued that the exercise of know-how is not necessarily an intentional action but sometimes an entirely automatic performance. Nevertheless, § 3.3 has shown that a competence is a kind of ability which cannot be exercised *only* automatically, but which can always also be exercised intentionally. This raises the question to what extent the exercise of a competence is a conscious act, something one is aware of doing. In this section, I will provide an answer to this question and to some related problems.

To introduce the relevant notion of consciousness, a natural starting point is the idea of intentional action on which I have already commented. After all, a paradigmatic case of something I do consciously is something I explicitly intend to do because I consciously decide to do it and then, in fact, go on to do so. However, it is important to realize that one might have lots of different things in mind when calling an act or a state ‘conscious’. At the very least, pertinent survey articles distinguish between a decent number of established senses of this term (cf. Van Gulick 2014; Menary 2009). What I am interested in here is not a performance’s being conscious in the sense that it has phenomenal qualities or that there is something it is like to perform that act. These and other important senses of the term ‘conscious’ notwithstanding, my topic concerns the consciousness of an act in the sense of being *aware* of it or of *attending* to it.¹³

In a survey article on consciousness and control of action in *The Cambridge Handbook of Consciousness*, Carlo Umiltà writes:

It is important to point out, however, that the terms “conscious” and “intentional” should not be used interchangeably when referring to action. In fact, consciously performed actions are not necessarily actions that are performed intentionally. One can be conscious of performing an action that is non-intentional, or automatic, in nature. [...] [I]ntentional actions engage processes different from those engaged in automatic actions (e.g. Prochazka et. al., 2000). Here the fundamental principle is that consciousness is necessary for intentional action. In some cases, this principle is explicitly stated in one of a variety of different forms, whereas in some other cases it is simply implied. (Umiltà 2007, 327–329)

Thus, every intentional action is conscious in the sense of awareness, but not every conscious act is therefore an intentional action. If every intentional action is conscious, it follows that those exercises of competences which are intentional actions are conscious acts. But since it is false that all automatic acts are unconscious, it remains an entirely open question whether none, some, or all of the automatic exercises of competences are conscious, too.

One might think that conscious awareness is indeed necessary for every exercise of a competence – be it intentional or automatic. In particular, this seems to be essential to the account I gave in § 3.3 and § 3.4. I have argued that the hallmark of normative guidance is, if not an intentional action guided by the norms of the activity in question, then at least the indirect control of automatic acts via a direct intentional influence on the automatic processes which produce it, intentionally bringing it into accord with the norms in question. For this to be possible, one must be aware of one’s performances.

¹³ For an overview of the attention and its significance, see Watzl (2011a; 2011b).

These considerations imply that consciousness in the sense of awareness is absolutely essential to the exercise of know-how and I am entirely in agreement with this line of thought. However, it is important to see that this argument only shows that some, indeed most and the paradigmatic cases of exercising competences must be conceived of as conscious. It does not require that there cannot be any non-conscious exercise of a competence whatsoever. For nothing I have said so far requires that every single exercise of a competence needs to play a role in the further improvement of this competence. Instead, there may very well be cases which do not.

I think that this is a welcome result. There are clear cases of unconscious exercises of competences – in sleepwalking. Both anecdotally and from scientific studies (cf. e.g. Luce & Segal 1966), we know that there are people who do remarkable things in their sleep. Some do not only walk around, but climb up the roof or get into the car and drive off. Others do not merely babble in their sleep, but talk clearly and coherently – sometimes even in dialogue with a wake person.¹⁴ I take it that such people are simply exercising the very competence they also exercise while being awake. Being asleep is certainly an impediment in exercising a skill where one is not as reliable as when one is awake. But even so, it is the very same skills that people exercise when awake and when asleep. After all, the reason why people can do what they do in their sleep, say speak a certain language, is simply that they have learned how to do so during waking hours.¹⁵

It would make for an interesting thought experiment to imagine a person who has a certain competence to engage in some activity which she is *only* able to exercise in her sleep. I, for one, am not sure if I can actually conceive of such a case.¹⁶ But to the extent that I can imagine something along these lines, it seems clear that we could not describe such a person as having a competence, but only as exercising a mere ability. This intuition is neatly captured by the account I have been advertizing. It follows immediately as soon as one acknowledges that some of the necessary ingredients for skills are missing here. Without intentional and conscious guidance by an understanding of norms, and without the possibility for an intentional exercise of the ability in question, it cannot be a case of know-how.

¹⁴ I know from testimony that I can express my gratitude very eloquently when I am asleep, and that I even respond politely to mumbled requests for clarification.

¹⁵ For more on sleepwalking, see Wallis (2008) and Brownstein & Michaelson (2016).

¹⁶ One may, however, think of somebody who cannot exercise certain competences because of panic attacks due to a severe trauma, but who is perfectly able to exercise these competences when sleepwalking or under the influence of hypnosis. In such a case, however, it is only accidentally true that the competence cannot be exercised when the person is awake. For a discussion of such examples, see § 5.2 and § 5.5.

However, exercising a competence in one's sleep is still a very peculiar and entirely derivative way to do so. It is entirely derivative because it depends essentially on the skill's having been learned during waking hours. And it is very peculiar because it is not only a wholly automatic performance, but one of which the person is entirely unaware. Therefore, these experiences are not available for the further shaping of the automatic processes involved. I can *speak* Spanish in my sleep, but I cannot *practice* my Spanish in my sleep, not even in the minimal sense that, say, I happen to notice my mispronunciations and try to remember to do better next time.

To some, this peculiarity of unconscious exercises of a competence will probably suggest the view that these performances are not intelligent and therefore not exercises of know-how. One philosopher who seems to subscribe to this assessment is Gilbert Ryle. He writes:

In judging that someone's performance is or is not intelligent, we have, as has been said, in a certain manner to look beyond the performance itself. For there is no particular overt or inner performance which could not have been accidentally or 'mechanically' executed by an idiot, a sleepwalker, a man in panic, absence of mind or delirium or even, sometimes, by a parrot. (Ryle 1949, 44–45)

Ryle, explicitly referring to sleepwalking, seems to hold that unconscious performances cannot be exercises of know-how. This is not surprising since unconscious performances are automatic acts and since Ryle often asserts that automatic performances are not exercises of know-how. In § 3.1, I have already discussed this topic and argued that Ryle's view is best understood as allowing for automatic exercises of competences. Along the same lines, I can now add that Ryle should be happy to accept unconscious exercises of know-how as well. In other words, Ryle's failure to accept unconscious exercises of competences is a direct result of his failure to accept automatic exercises of competences in general. If my arguments against the dismissal of automatic exercises of know-how in general are sound, so are my arguments against the dismissal of unconscious exercises of know-how.

In sum, I suggest that the exercise of know-how – be it intentional or automatic – is paradigmatically a conscious performance – an act which the person is aware of performing. However, to the extent that a skill is habituated or automatized, it can also be exercised unconsciously – paradigm examples being cases of sleepwalking. The sufficiently automatized processes are precisely what is responsible for the production of entirely unconscious acts. Despite other peculiarities, I hold that there is nothing mysterious in calling these performances exercises of the very same competences which one can also exercise consciously.

These considerations also allow me to comment on one of the senses of ‘practical knowledge’ which I have distinguished in the Introduction – the knowledge of what one is currently doing. Following the seminal work of Elizabeth Anscombe (1957), there has been a lively debate about the nature of such knowledge, focusing, among other things, on the idea that it does not require observation (cf. Schwenkler 2012). Sadly, I cannot give sufficient credit to the intricacies of this debate in the context of this book. However, I take it to be obvious that there is a close relation between what I have been concerned with, consciousness in the sense of awareness of what I am currently doing, and ‘practical knowledge’ in the sense of knowledge that what I am currently doing is such-and-such.

Gilbert Ryle also seems to suggest such a close relation when he takes the term ‘intelligent’ to involve the agent’s ‘thinking what she is doing’ (cf. Ryle 1949, 32). And there is also a systematic argument to be made for this.¹⁷ Intuitively, being aware of what I am currently doing already *is* being aware of it *as* engaging in some particular activity – say, *as* dancing or *as* talking. Thus, such awareness must already be understood as involving the concept of what I am currently doing and the knowledge that I am currently doing such-and-such.¹⁸ This goes hand in hand with the assessment of one’s own performances. The capacity to make such assessments was already introduced in § 2.2 and it will play a crucial role in chapter 4.

§ 3.6 The Phenomenological Fallacy

I have argued that the exercise of know-how can proceed both completely automatically and fully intentionally, and that it involves consciousness in the sense of self-awareness or practical knowledge of one is doing. To conclude this account, I will now comment on the role of explicit thinking in the exercise of know-how. This will also include the alleged phenomenological support which unreflective, automatic performances without occurrent thought are often claimed to provide for anti-intellectualism. As I shall argue, however, this argument for anti-intellectualism fails. It constitutes what I shall call a ‘phenomenological fallacy’.

¹⁷ Further considerations on this relationship have been proposed by Setiya (2012), Brownstein (2014) and Hornsby (2016), among others.

¹⁸ One might object that mere awareness of what I am currently doing may only consist in a true belief as opposed to full-blown knowledge. But as Adrian Haddock (2010) and others have suggested, the justification of such a true belief may already be grounded in the fact that I am in intentional control of my performances. Similarly, in § 4.1, I will propose an argument to the effect that assessments of an activity which stem from an understanding of this capacity amount to knowledge for precisely this reason.

The question of explicit thinking and can be put in terms of the distinction between an actor's beliefs about what to do as remaining merely *dispositional* beliefs or as becoming *occurrent* at the time of the exercise of the competence (cf. Schwitzgebel 2010, sect. 2.1). To begin with, I take it to be clear that, sometimes, explicit thinking indeed occurs in the exercise of a competence. For example, my practical knowledge that what I am currently doing is writing is not merely a dispositional belief but just became occurrent. In § 2.5, I have followed Ryle in showing how propositional knowledge plays an important role in learning and improving competences. As I have already quoted on page 64, Ryle thinks that propositions like “[l]ogical rules, tactical maxims and technical canons are [...] helpful only to the half-trained.” (Ryle 1945a, 14) Here, it may also be the case that not only dispositional beliefs are in play, but that an actor exercising a competence explicitly and reflectively considers relevant propositional knowledge. Paradigm examples of this include people who cook with the help of a recipe they have memorized and need to recall explicitly from time to time. And some people may only be able to screw or unscrew a bolt after they recite “Righty, tighty; lefty, loosey” (cf. Bengson & Moffett 2011b, 15–16).

But experts certainly do not or only seldom explicitly consider rules when they are engaged in their practice. Hubert Dreyfus famously argues:

[P]henomenology suggests that, although many forms of expertise pass through a stage in which one needs reasons to guide action, after much involved experience, the learner develops a way of coping in which reasons play no role. After responding to an estimated million specific chess positions in the process of becoming a chess master, the master, confronted with a new position, spontaneously does something similar to what has previously worked and, lo and behold, it usually works. In general, instead of relying on rules and standards to decide on or to justify her actions, the expert immediately responds to the current concrete situation. (Dreyfus 2005, 53)

However, the phenomenological fact that experts do not explicitly and consciously reflect does not entail what Dreyfus takes it to entail, namely that when an expert performs, there is no “relying on rules and standards to decide on or to justify her actions”. Dreyfus and a number of other philosophers prominently, and notoriously, defend this inference. One of the boldest statements of this conclusion reads:

[T]he expert is simply not following any rules! (Dreyfus & Dreyfus 1986, 108).

But if this were indeed the consequence of phenomenological facts, then experts would not have genuine know-how, after all precisely because they would not be guided by norms anymore. This would threaten to collapse

the whole distinction between know-how and mere ability and make it impossible to explain the core facts – ironically, *phenomenological* facts – with which the Rylean project of understanding competence begins (cf. § 1.5).¹⁹

Fortunately, however, this consequence can be avoided. As have argued extensively in § 3.4, it may be true that one “immediately responds to the current concrete situation”, and this entirely automatically, while one is still guided by norms in the full sense of the term. In chapter 4, I shall go on to substantiate this insight with a full account of this notion of guidance.

Thus, the premise that automaticity is widespread and crucial for the exercise of competences fails to entail the conclusion that automaticity is essential to competence to the extent that every form of occurrent thought or deliberate intentionality is an impediment to expertise. To cling to this inference is to overstate a mere part of the phenomenological evidence over the rest of the evidence. In a word, it is a phenomenological fallacy.

Bracketing this fallacy, however, Dreyfus’ phenomenological evidence is certainly very important. As he reports, explicit and conscious reflection may *impede* rather than support the exercise of a skill. Following John McDowell and Hans-Georg Gadamer in calling such explicit and conscious reflection a “free, distanced orientation” towards one’s environment (cf. McDowell 2007, 346; Gadamer 1992, 445), Dreyfus presents this case:

Chuck Knoblauch (Dreyfus 2007, 354)

As second baseman for the New York Yankees, Knoblauch was so successful he was voted best infielder of the year, but one day, rather than simply fielding a hit and throwing the ball to first base, it seems he stepped back and took up a “free, distanced orientation” towards the ball and how he was throwing it—to the mechanics of it, as he put it. After that, he couldn’t recover his former absorption and often—though not always—threw the ball to first base erratically—once into the face of a spectator. Interestingly, even after he seemed unable to resist stepping back and being mindful, Knoblauch could still play brilliant baseball in difficult situations—catching a hard-hit ground ball and throwing it to first faster than thought. What he couldn’t do was field an easy routine grounder directly to second base, because that gave him time to think before throwing to first. I’m told that in some replays of such easy throws one could actually see Knoblauch looking with puzzlement at his hand trying to figure out the mechanics of throwing the ball. There was nothing wrong with Knoblauch’s body; he could still exercise his skill as long as the situation required that he act before he had time to think.

¹⁹ There is certainly much more to be said about the status of this inference (cf. Dreyfus 2002; Dreyfus 2005; Dreyfus 2007; McDowell 2007; Rietveld 2010; van Dijk & Bongers 2013). However, I contend that, as far as the big picture is concerned, my argument is entirely accurate. An argument along the same lines has been proposed in more detail by Ellen Fridland (2014b, sect. 3), as well as in cognate criticisms (cf. e.g. Rey 2002; Sutton *et al.* 2011; Gottlieb 2011).

Dreyfus continues: “In this case we can see precisely that the enemy of expertise is thought.” (Dreyfus 2007, 354) “The same phenomenon must occur in all high-speed, skilled activities.” (Dreyfus 2007, 365 fn. 8)

As I have already pointed out, it may sometimes be true that “the enemy of expertise is thought” – at least “in all high-speed, skilled activities” – while expert performances are nevertheless guided by an understanding of the relevant norms. The phenomenon that conscious reflection can impede expertise in certain high-speed activities is well-supported by empirical work in cognitive science (cf. e.g. Beilock *et al.* 2002; Flegal & Anderson 2008). This is also known as ‘Steve Blass Disease’, with reference to a further affected athlete next to Chuck Knoblauch. But this phenomenon is perfectly compatible with the account I have been offering so far.

Such “high-speed, skilled activities” crucially depend on the automatization of many, if not all performances, simply because there is normally not enough time to perform them after prior reflection. At the same time, such automatisms may very well be fragile and vulnerable to disturbance by conscious reflection. Thus, if a highly specialized and automatized skill such as catching and throwing the ball quickly is once brought out of balance by the tendency to reflect, it seems plausible that, in extreme cases such as *Chuck Knoblauch*, the exercise of the competence will remain very reliable in cases with no time for the tendency to reflect to interfere, but become rather unreliable where there is enough time for this to happen.

These phenomena concern the way in which explicit reflection can both sometimes impede and support the exercise of a skill. Conversely, however, it is also true that completely unreflective automaticity can also both sometimes be a supporting factor and sometimes interfere with a competence’s quality. Its supporting function has already been discussed extensively here and in § 3.1, for example in the case of high-speed athletic performances. In addition to these considerations, unreflective automaticity in more basic activities appears to be necessary for proficiency in higher-order activities – for example, one cannot have the competence to give a good philosophical talk if one is not sufficiently competent at speaking the relevant language in an automatic way which does not require too much explicit reflection.

Likewise, unreflective automaticity can also constitute an impediment for the quality and the exercise of a skill. For example, a boxer may realize that her unreflective and automatic attempts to parry her opponent’s blows are bound to fail, but still continue to automatically make such attempts, despite the fact that she actually knows better and is trying in vain to improve her habits. Another case in point concerns a widespread phenomenon in second-language acquisition, the fact that some speakers fail to improve

their capacities beyond a certain point because they have learned that they can get by with, say, an only close-to correct pronunciation, and maybe even fail to perceive the relevant phonetical differences at all. Linguists call this problem the ‘fossilization’ of second-language competence (cf. e.g. Han 2004; Han 2012). I take it to be clear that the problem with ‘fossilized’ language competence consists in imperfect automatisms which cannot be improved any further, or would require extreme effort to do so. And the same may very well also occur in other activities.

However, it should be noted that the phenomenological evidence is not univocal. Dreyfus argues that occurrent thought may impede the exercise of expert competences, but there are also many examples which show how crucial, indeed essential it is not only to think automatically or implicitly, but explicitly and occurrently. In fact, there is even strong phenomenological evidence that this is the case even in the kinds of activities which Dreyfus cites as paradigm cases – in elite sports and other high-speed physical activities. Barbara Montero, for example, shows that this is the case in professional ballet dancing (cf. Montero 2010; Montero 2013). Thus, the relationship between conscious reflection and efficient practice is much more intricate than anti-intellectualists often assume.

What I have discussed here can be called *the dialectics of automaticity and reflection*, of habit and intentionality. Explicit reflection on what to do both sometimes supports and sometimes impedes the exercise of a competence, and unreflective automatization can equally both support and impede a skill’s exercise or development. Thus, an account of the core notion I propose in order to account for know-how, the notion of normative guidance, must be open to and dependent upon *both* explicit reflection and unreflective automaticity. This will be the topic of chapter 4.

Chapter 4

Intellectual Guidance

In chapters 1 and 2, I have followed Ryle in arguing that know-how is an intelligent ability, and that intelligence requires the intellect. In this chapter, I will come back to these intellectual aspects and I shall build on my discussion of the exercise of competences in chapter 3 in order to complete my positive proposal and offer an account of how an understanding of an activity can be seen to guide the conduct of a skilled person.

In § 4.1, I propose an account of what it is to understand an activity in terms of the capacity to assess performances of it. On this basis, an understanding of an activity will turn out to require the possession of the concept of this activity, as shown in § 4.2. In § 4.3, I proceed to discuss the form of the assessments reached by exercising this conceptual competence. As I shall argue, these assessments constitute states of propositional knowledge.

In § 4.4, I spell out how the understanding of what an activity demands can play the role of guiding the exercises of one's competence to engage in that activity. While this proposal may be seen to create a problem of vicious regress, § 4.5 and § 4.6 will discuss two different versions of this problem and show in detail that both worries are unfounded. Finally, § 4.7 will defend my account, among other things by discussing the obvious problem that it appears viciously circular because conceptual assessment capacities must themselves be understood as competences. This will also lead me to general considerations about the self-reflexivity of concepts.

§ 4.1 Understanding as Capacity to Assess

I have argued for two necessary conditions of know-how. In this section, I will show that these conditions are one. *Mutatis mutandis*, my arguments for these necessary conditions will be revealed as supporting the same claim.

In § 1.5, I have argued that an understanding of the normative structure of an activity is a necessary condition for the knowledge how to engage in it. In § 2.2, I have argued that know-how requires the capacity to assess the performances of the activity in question. Now, I shall go on to claim that a conception of an activity, in the sense in which it is involved in know-how, simply *is* the capacity to assess performances of that activity against the very standards which govern it.

To make this case, I shall start with the concept of understanding. Of course, what understanding amounts to is itself far from trivial and there is an intricate debate surrounding this notion (cf. e.g. Grimm 2012). However, I shall only try to give an account of the particular kind of understanding which is at issue here, without undertaking too serious commitments concerning understanding in general. Following a classic distinction (cf. e.g. Kvanvig 2003, 189), there are two main varieties of understanding: propositional understanding – that is, understanding that or why something is the case – and objectual understanding – that is, understanding of some object in the broadest sense of the term. The understanding of an activity, then, is a kind of objectual understanding. It is just *that* kind of objectual understanding where the object is an activity governed by norms.

Obviously, I cannot comment on the debate about the concept of understanding here at length. However, I shall discuss four characteristic properties of understanding which seem to be specific to this epistemic concept.¹ I shall argue that these aspects also pertain to the understanding I am interested in – the understanding of normative activities – and that these can be explained very well by the assumption that such an understanding consists in the competence to assess performances of this activity.

First, understanding admits of degrees. As I have repeatedly argued, competences also admit of degrees. Therefore, it is attractive to identify the understanding of a normative activity with the competence to assess performances of that activity against these norms. This way, the graduality of the understanding just is the graduality of proficiency in assessing.

Second, while it is a matter of controversy whether or not understanding can somehow be *reduced* to propositional knowledge, it is common ground that understanding *entails* some propositional knowledge, or at least *some* true beliefs about the subject at hand. For example, whether or not my understanding of computer programming is exhausted by the things I know about computer programming, I can only *have* this understanding if I know or truly believe at least *some* things about computer programming.

¹ These properties are often contrasted with those of propositional knowledge. One of the most prominent issues is the phenomenon of epistemic luck. I discuss this in § 6.3.

This constitutes a further *prima facie* reason to identify the understanding of a normative activity with the capacity to assess performances of it. As foreshadowed in § 2.5, and as I will discuss in more detail in § 4.3, there is a direct relation between understanding, assessment and propositional knowledge. By exercising the capacity to assess performances of a given activity, such as computer programming, one comes to believe that certain propositions are true, say, the proposition that entering this particular line into the program code will fix the current problem. And crucially, the propositions which play a role here are exactly the propositions which would also play a role in what is entailed by my understanding of computer programming. Therefore, the identification of the understanding of a normative activity with the capacity to assess performances of it makes perfect sense of the relation between the understanding of the activity and the propositional knowledge about it. It is just is the relation between the competence to assess performances of the activity in question and the knowledge it entails.

Third, and relatedly, it has been argued that understanding is not wholly ‘factive’, unlike propositional knowledge.² While my knowledge that the heater is broken entails the truth of my belief that the heater is broken, my understanding of the heater and the way it works does not entail that I believe *only* true propositions about this subject matter. Instead, my objectual understanding is compatible with my having false beliefs about the thing I understand. At least, this is Jonathan Kvanvig’s view. He writes:

Objectual understanding is, of course, not straightforwardly factive, for only propositions can be true or false. Still, the uses I wish to focus on are ones in which factivity is in the background. For example, to understand politics is to have beliefs about it, and for this objectual understanding to be the kind of interest here requires that these beliefs be true. (Kvanvig 2003, 191)

[I]t is hard to resist the view that understanding may be correctly ascribed even in the presence of some false beliefs concerning the subject matter [...]. When the falsehoods are peripheral, we can ascribe understanding based on the rest of the information grasped that is true and contains no falsehoods. In such a case, the false beliefs are not a part of the understanding the person has, even though they concern the very material regarding which the person has understanding. So in this way, the factive character of understanding can be preserved without having to say that a person with false beliefs about a subject matter can have no understanding of it. (Kvanvig 2003, 201–202)

² The notion of factivity is somewhat problematic in this context since it expresses the thought that something is, in fact, true. But an understanding of something cannot be true or false, only more or less accurate. Talk of ‘factivity’ in this context should therefore be understood as a shorthand for a more general notion of objective correctness, including the truth of a proposition and the accuracy of an understanding.

It has been questioned whether this is too restrictive. Some philosophers argue that understanding does indeed survive large amounts of error, even error about matters at the core of what is understood (cf. e.g. Zagzebski 2001; Elgin 2009; Riggs 2009). While I cannot assess this issue at length, I think that Kvanvig's view can be defended, at least for the purposes of the project of this book. There may be cases which involve significant proportions of false beliefs about the subject matter and even false beliefs about its most important aspects. But it is still part of the concept of understanding that the things about which the person who has that understanding is *correct* are even *more* significant, where this assessment may include a comparison in number, reliability, relative weight or other respects.

If this view is correct, then there is further reason to identify the understanding of a normative activity with the capacity to assess performances of it. As I have argued, somebody who possesses a competence is hardly ever perfect at the activity in question (cf. § 2.3) and also does not have to be absolutely reliable (cf. § 1.4). Understanding admits of such degrees as well. As already noted, by exercising the capacity to assess a given performance, one comes to believe that certain propositions are true, say, the proposition that this particular move in our game of chess will upset my plans. Likewise, false beliefs stem from imperfect or failed performances of assessing chess. The fact that understanding requires sufficient propositional knowledge is mirrored in the fact that a competence requires sufficient reliability.

Let me take stock. I have presented three *prima facie* reasons to identify the understanding of a normative activity with the skill to assess performances of it. First, both are gradable, and the graduality of the understanding of an activity can be cashed out in terms of the graduality of the assessment capacity. Second, both entail propositional knowledge, and the way in which the understanding of an activity entails propositional knowledge can be explained in terms of the way the assessment capacity does. Third, both are compatible with having false beliefs about the activity in question, and the way in which this holds for the understanding of an activity can again be cashed out in terms of the way in which it holds for the assessment competence. Given my initial observation that know-how entails both an understanding of the activity in question and the capacity to assess performances of it, this even amounts to a total of four facts which can be explained very neatly if the understanding of a normative activity simply is the same thing as the capacity to assess performances of it.

Of course, some of these facts only support the view that such an understanding consists in *some* competence, but remain neutral on the question *which* competence it is supposed to be. This is particularly clear in the

case of the first point concerning the graduality of both competence and understanding. However, I contend that the other three facts can only be explained if the understanding of a normative activity is identified with the specific competence to assess performances of it. As I have argued, the kind of propositional knowledge which an understanding of an activity requires just is the kind of propositional knowledge which the capacity to assess performances of that activity requires. And the false propositions which one may believe about an activity, and still count as understanding it, correspond to the false propositions which one may come to believe on the grounds of an imperfect capacity to assess performances of that activity.

My argument therefore proceeds in a two-step inference to the best explanation. If the identification of the understanding of a normative activity with *some* skill already constitutes a good explanation of *certain* facts, and if its identification with the *specific* skill to assess performances of it constitutes an even more general explanation of these and *further* facts, then this idea is indeed sufficiently justified. True, *qua* inference to the best explanation, this is not a knock-down argument. Its merits and flaws are interwoven with the general explanatory power of the account I am advertising.

Again, my proposal should not be understood as an account of the concept of understanding in general, but only of the specific kind of understanding which is objectual and directed at a normative activity. Whether this proposal can be generalized into an account of other cases of objectual understanding or even of other kinds of understanding is an interesting question, but one independent from the problems discussed in this book.

I would also like to mention a *prima facie* problem for my account. If a competence involves guidance by an understanding of what the activity in question demands, and if this understanding is the capacity to correctly assess performances of that activity – *itself* a *further* competence –, how should we conceive of the understanding which guides the exercises of *this* capacity? It seems as if this requires a second capacity to assess the performances of the first, and so on *ad finitum*. This is indeed a severe problem which I will have to solve. But I will first need to lay out the details of my account of normative guidance before I can return to this problem in § 4.6.

§ 4.2 Assessment Capacities as Conceptual

I have argued that the understanding required by competence can be identified with the capacity to assess individual performances of the relevant activity. This, however, paves the way for a further important issue – the

way in which know-how involves concepts and conceptual capacities. This is the topic of the present section. Here, I shall not discuss other accounts of know-how which also include an appeal to concepts and conceptual capacities.³ The most important of these, the objectualist intellectualism defended by John Bengson and Marc Moffett (2007; 2011c), will later be addressed in detail, especially in chapter 8.

To begin with, I think that there is a very straightforward argument for the role of concepts and conceptual capacities already implicit in the account as I have spelled it out so far. I have argued that know-how entails true beliefs, some of which even amount to propositional knowledge. § 2.5 has shown which role these beliefs play in learning a competence, and the current argument has shown that the exercise of an understanding of an activity – i.e. of the capacity to assess performances of it – *generates* further beliefs in the form of particular assessments. However, one can only have a belief and entertain a proposition if one possesses the concepts which are involved in this proposition.⁴ And clearly, the one concept which plays a role in all of these beliefs which are entailed by a competence is the concept of the relevant activity itself. In § 4.3, I will offer further support for this point by spelling out the general form of the propositional assessments required for understanding an activity.

Thus, these considerations already commit me to the view that the capacity to assess performances of an activity must involve the concept of this activity. But it has remained unclear how exactly this should be spelled out in detail. I contend that it is very natural to say that what one assesses when one assesses an act as a performance of a given activity *just is* whether or not it falls under one's concept of that activity. Of course, this is not as simple as saying whether, say, my computer is switched on. To assess an act as a performance of a given activity is a qualitative matter which admits of degrees of normative adequacy. But, crucially, so does the corresponding concept. For example, in assessing an act as a performance of dancing, one wonders whether or not the concept of dancing *simpliciter* applies to it at all. But this is only part of the assessment. One also wonders to what degree the *qualitative* concept of dancing *well* applies to that act. The degrees to which the act meets the relevant norms are mirrored in the degrees in which the concept of a good, proper or elegant dance applies to it.

³ This also includes Bartels & May (2015a), who also claim that this view ties in well with the way in which empirical cognitive science draws on concepts. For discussion, see Glauer (2015) and Bartels & May (2015b). I discuss other core notions from cognitive science, particularly on the concept of procedural knowledge, in § 6.5.

⁴ At the very least, this is true on the most natural and most widely held view of what propositions are. On this debate, see footnote 7 on page 67.

Of course, I should note that there is an important and substantial debate about the notion of a concept on which I cannot even begin to comment here (cf. e.g. Laurence & Margolis 1999; Margolis & Laurence 2011). But while it is a difficult question what a concept *is*, there is one thing which I take to be, at least nearly enough, common ground: One *possesses* a concept if, and only if, one has the capacity to employ it correctly. One may also defend the related stronger claim that possession of the concept *just is* possession of the capacity for its correct employment. But the first and weaker claim suffices for the purposes of this book.

To clarify, this view requires a notion of correctness which applies to the employments of the relevant concepts. But this is nevertheless neutral with respect to several substantively different explications of concepts. To illustrate this point in an overly simplified way, my view is neutral between, on the one hand, representational accounts, according to which a concept is applied correctly just in case a corresponding property actually applies to the object(s) in question, and, on the other hand, inferential accounts, according to which it is applied correctly just in case the inferential rules which govern the use of the concept allow for this application. These accounts are substantially different, but both clearly support a notion of correct use. At this level of abstraction, we can bracket the question whether this notion is explained in terms of referential or inferential correctness.⁵

Given that the *possession* of a concept is essentially correlated with a capacity, it has been argued that concepts can *themselves* be understood as the capacities of their use (cf. Margolis & Laurence 2011, sect. 1.2). For example, Anthony Kenny writes:

Concepts are best understood as a particular kind of human ability: a person who has mastered the use of a word for F in some language possesses the concept of F. Abilities are individuated by their possessors and their exercises, though they are not to be identified with either. (Kenny 2010, 105)

I take it that Kenny is right to point out that an ability is not identical with its exercises, just as no potentiality is identical with its actualizations (e.g. my window's fragility is not identical with the instances of its breaking), and that an ability is not identical with its possessor just like no potentiality is identical with its bearer (e.g. my window's fragility is not identical with my window). But his claim that an ability can nevertheless be *individuated* by its possessor and its exercises strikes me as problematic. For this obviously only leads to a conception of *token* abilities as opposed to types of abilities.

⁵ This point is nothing new, of course. It is part of the common ground of the debate about rule-following and the normativity of meaning (cf. § 1.6). And on page 37, I have already quoted a crucial passage on this question by Paul Boghossian (1989).

In particular, this view makes it impossible to conceive of *shared* abilities and competences. If the competence to play chess is individuated in terms of its possessor, then I have my token competence to play chess and you have yours. This view renders it unclear how all of these token competences can be token competences of the same type – competences to play chess. Even worse, if such competences are supposed to explain concepts, then there are only subjective tokens of concepts rather than shared types of concepts.

But this problem can be solved on the basis of § 1.2, where I have suggested that an activity can be individuated in terms of the set of norms which govern it. Along these lines, we can see that two people can possess the same competence – know how to engage in the same activity – because they understand and are guided by the same set of norms – the norms which constitute the activity in question. On this basis, we should not individuate concepts by their possessors and their exercises alone, but instead by the activities of their use, that is, by the specific norms which govern their use.

Crucially, however, this point about the *individuation* of concepts does not entail an answer to the question what a concept *is*. As I have already stressed, I cannot attempt to answer this big question here. One may go on to argue that concepts cannot merely be individuated in terms of the activities of their use, i.e. the norms of their correct use, but that concepts are *identical* with these activities, i.e. that activities *are* these sets of norms. However, one may also go on to treat this point about the individuation of concepts merely as a criterion of adequacy for any substantial account of concepts. Given that I have already sketched the reasons why referentialist and inferentialist accounts of concepts both support a notion of correct use, it is plausible to assume that there is still much room for different views.

Returning to my main argument, Kenny's remark on the preceding page makes clear that the easiest way to appreciate the fact that possessing a concept is a skill rather than a mere ability is via the expression of concepts in language. Mastery of a natural language is one of the most salient examples of know-how. Gilbert Ryle, among many others, cites this example frequently (cf. *Ryle's Range of Cases* on page 14) and it underwrites the debate about rule-following and the normativity of meaning (cf. § 1.6). Accordingly, if the mastery of a word in a natural language is an example of know-how, so is the mastery of the concept which that word expresses.

However, it is only a sufficient and not a necessary condition for the possession of a concept F that one has mastered the use of a *specific* word 'F' in some language (cf. e.g. Kenny 2010, 106). That is, I may possess and be competent in the application of a concept for which I do only have multiple words, complex descriptions or even only demonstrative conceptions.

I conclude that the capacity to assess an act as a performance of an activity already involves the use of the concept of that activity. In this sense, know-how requires a conception of the activity in question, and it requires the possession of the concept of this activity.⁶

§ 4.3 Assessments as Propositional Knowledge

I have argued that know-how requires an understanding of the relevant activity, i.e. the conceptual capacity to assess the performances of it (cf. § 1.5, § 4.1 and § 4.2). In this section, I would like to show in detail how having such an assessment capacity entails having propositional knowledge about the activity in question in the form of the correct assessments reached by exercising this competence.

In doing so, I expand on Ryle's argument from § 2.5, where I showed that propositional knowledge can play an important role in the acquisition of competences. Now, I will support the additional claim that having an assessment capacity for a given activity also requires having and coming to possess propositional knowledge.⁷

I proceed in two steps. First, I shall propose an argument for this conclusion, establishing two interrelated kinds of correct assessments which somebody with an understanding of an activity must make and be capable of making. Second, I shall show why these correct assessments amount to full-blown propositional knowledge as opposed to merely true beliefs.

Let me begin by introducing the idea that there are two kinds of assessments involved here. The capacity to assess the performances of an activity is a capacity to make correct judgments about individual cases. Crucially, however, these individual cases are always specific and situated. A specific act is performed in a specific situation. One cannot assess the quality of an individual act without taking into account its circumstances. The very same behavior will be assessed very differently in different circumstances.

On this basis, the very same capacity to assess acts in their circumstances can also be employed to focus on these circumstances themselves, assessing them with regard to the question which acts would, in these circumstances, amount to a good performance of some activity. In § 1.4, I have discussed

⁶ Of course, one of the main points of dissent is whether or not one may possess concepts without *any* linguistic capacities whatsoever. This problem is closely connected to the problem whether or not non-linguistic beings are capable of thinking, of knowing facts and of knowing how to do things. I will discuss this briefly in § 6.7.

⁷ For independent considerations which point in this direction, see Michael Luntley's arguments for the view that the mark of an expert is not what she already knows or how she knows it, but her capacity to learn (cf. Luntley 2009).

Ryle's notion of a normal situation with regard to the a given activity, and I have introduced the notions of an option or an affordance for doing so. In these terms, assessing circumstances in this way can also be described as perceiving options or affordances for exercising a given competence.

In sum, assessing performances in their circumstances and assessing circumstances with regard to the options they provide are two sides of the same coin – both are exercises of the single capacity to assess the activity in question. Of course, it is a delicate matter to specify the precise absolute and relative quality of a performance as an exercise of some activity of A-ing and the exact quality of an option as an option for A-ing. But as discussed in § 1.4 and § 1.7, I can leave open these details in the present context.

This argument establishes the following claim:

- (B) The exercise of a capacity to assess A-ing leads to assessments of individual acts and situations, i.e. judgments of the forms (B 1–2).
- (B 1) Token act x in token circumstances c is a performance of A-ing of quality Q.
- (B 2) Token circumstances c provide the option for performing token act x, which is a performance of A-ing of quality Q.

In other words, possessing an understanding of A-ing entails possessing and coming to possess propositional knowledge of the forms (B 1–2).

Such assessments do not require sophisticated conceptions of individual acts and their circumstances. In § 1.4, I have already argued that it is impossible to explicate all the enabling and defeating conditions of doing something well – to explicate what Ryle called a 'normal situation' for the exercise of a given competence. Likewise, it is overly demanding to assume that possessing a conception of an activity requires a complete grasp of all the relevant aspects of an individual act and an individual situation.

However, this does not constitute an argument against my claim that knowledge of the forms (B 1–2) is nevertheless necessary for possessing an understanding of an activity. For nothing I have said entails that one needs to possess a *descriptive* conception of these acts and situations. It is entirely sufficient to possess a *demonstrative* conception of them in order to possess knowledge of the forms (B 1–2). As I shall discuss later, in § 8.4, this point has already been stressed by intellectualists about know-how. One way to illustrate this idea is with respect to paradigm examples.

Being competent at swimming or otherwise possessing an understanding of swimming requires knowledge of the form (B 1) because it requires the correct assessment of paradigm cases of good swimming and of paradigm

cases of not swimming at all. Somebody who carefully observes professional swimmers in a breast stroke championship race and fails to know that what they do qualifies as swimming the breast stroke rather well would hardly count as being able to assess performances of swimming the breast stroke. At the very least, we would have to make substantial adjustments to what we take that person to believe and how reliable we take her perception to be in order to maintain the idea that she is actually still competent to assess what she so blatantly fails to judge correctly. Generally speaking, somebody who systematically fails to recognize *all* paradigm examples of particularly good exercises of an activity does not have the capacity to assess performances of that activity. And the same holds for performances which clearly do not amount to swimming the breast stroke at all.

Likewise, being competent at swimming or otherwise possessing an understanding of swimming requires knowledge of the form (B 2) because it requires the correct assessment of paradigm situations which offer the option to swim and of paradigm situations which do not offer this option at all. To take simple examples, this involves the knowledge that swimming requires a certain minimum amount of water and the absence of other swimmers or objects who are constantly in one's way, and the knowledge that a perfectly calm swimming pool without any people or objects floating around offers a perfect opportunity for swimming. Again, if somebody fails to possess such propositional knowledge, we would, at the very least, have to substantially adjust our interpretation of this person in order to maintain the idea that she is actually still competent to assess what she so blatantly misconceives.

As should have become clear from my characterization of these examples, my claim (B) does not require the people in question to possess anything like a rich descriptive conception of the performances and situations in question. It is entirely sufficient to just know demonstratively, *of* a certain act or situation, that it is assessed *thus*.⁸ However, it is hard to imagine a case where somebody who is genuinely competent at swimming or who otherwise possesses a genuine understanding of swimming has *only* and *purely* demonstrative conceptions of the relevant acts and situations.

⁸ John McDowell has prominently made this point in a discussion of color concepts, pointing out that “one can give linguistic expression to a concept that is exactly as fine-grained as the experience, by uttering the phrase ‘that shade’, in which the demonstrative exploits the presence of the sample.” (McDowell 1994, 57) But McDowell’s insight is a general one and can also apply elsewhere. This also includes his famous debate with Hubert Dreyfus which has already played a role in § 3.6 (cf. McDowell 2007; 2013; Dreyfus 2005; 2007; 2013; Schear 2013). But since this exchange is concerned with the nature of mind and action in general, rather than specifically with know-how, I shall not address this debate in detail. For further discussion of demonstrative concepts with respect to know-how, see § 8.4.

First, even if one starts out with such purely demonstrative conceptions, they remain genuinely conceptual. And if things go well enough, it will therefore be possible at least to demonstratively re-identify numerically distinct acts as acts of ‘doing the same’ as earlier performances and numerically distinct situations as situations ‘offering the same option’ as earlier ones. Thus, demonstrative conceptions are nevertheless sufficiently general. They can, at least in principle, also be applied elsewhere.

Second, it is hard to imagine that in such a process the relevant concept remains *purely* demonstrative and does not, precisely *through* such re-identifications, gain at least *some* descriptive content. For example, somebody who first has a merely demonstrative conception of a token instance of paradigmatically good breaststroke swimming will, other things being equal, be able to realize, in her re-application of this demonstrative concept to other token instances of swimming the breaststroke, that such swimming performances are *ceteris paribus* better to the extent that both pairs of limbs are moved synchronously. In this way, a demonstrative conception of a paradigmatically good example of swimming the breast stroke easily gains at least some descriptive aspects, even if its demonstrative nature remains.

I take it to be uncontroversial that a conception of A-ing often involves relatively rich descriptive elements in cases of learning from explicit teaching (cf. § 2.3). What my present argument suggests is that at least some descriptive elements are directly involved, or evolve very easily, in *every* conception, whatever its source, even if the degree of descriptive specificity and sophistication varies very strongly and arguably even reduces to trivial matters at the extreme end. On the basis of (B), then, I can also formulate a further way in which propositional knowledge is involved in possessing assessment capacities – as assessment *principles*:

- (C) Every capacity to assess A-ing involves knowledge of at least some principles of assessment of acts and situations of the forms (C1–2).
- (C1) An act of the type X in circumstances of the type C is a performance of A-ing of quality Q.
- (C2) Circumstances of the type C provide the option for performing an act of the type X, which is a performance of A-ing of quality Q.

It should be noted that it is an entirely open question how fine-grained the types of acts abbreviated as ‘X’ are individuated and how much sophistication a competent actor possesses in understanding the interrelations between these ways – say, as mutually exclusive, as entailing each other, as hierarchically ordered, and so forth. As discussed in § 1.7, even the level

of granularity of individuating the *whole* of the activity somebody knows how to engage in can differ in different contexts. The same holds for the granularity of the types of acts which count as engaging in this activity.

One may worry that such a view is too far from Ryle's position. But given my considerations in Chapter 2, Ryle can happily endorse this view. All forms of assessment principles, however demonstrative or descriptive they may be, still clearly depend on particular assessments of individual cases of intelligent conduct. (C) depends on (B), which, in turn, depends on the existence and practice of the activity in question. This is in keeping with Ryle's view that practice is prior to descriptive theory (cf. § 2.6).

I have argued that understanding an activity necessarily requires propositional knowledge. But this is not to say that understanding an activity requires knowing *a lot* about that activity. I do not deny the fact that many competent practitioners are not particularly good theorists of their practices.⁹ As Ryle has correctly stressed, being able to abstract away from particular instances and to express the relevant differences descriptively is distinct from the competence to produce these instances oneself (cf. § 2.1). But I have argued that at least a minimal degree of these capacities is involved in every understanding of an activity and hence in every competence, vast differences between allegedly 'pure' practitioners and expert theorists notwithstanding. Further, with respect to the notion of guidance as responsible control which I will propose in § 4.4, the crucial kind of knowledge is only knowledge of the forms (B 1-2) – i.e. assessments of individual performances and situations, which may well be demonstrative. Thus, even when somebody *misconceives* what they do and attend to on a *descriptive* level – as apparently many professional athletes do when it comes to the skill to catch balls (cf. Reed *et al.* 2010) –, they may still be able to point to particular performances and correctly assess them as good or bad. This, too, is propositional knowledge which stems from a genuine assessment capacity, a genuine understanding of how to catch balls.¹⁰

So far, however, all of this propositional knowledge is independent from the question whether somebody actually possesses the relevant competence or if she merely possesses an understanding of this activity. But in fact, there is one kind of propositional knowledge which one can only come to possess when one possesses the relevant competence oneself.

⁹ This phenomenon has been studied extensively in cognitice science (cf. e.g. Berry & Broadbent 1984; Reed *et al.* 2010).

¹⁰ Cases like these are still problematic since they involve a tension between what a person explicitly declares and what she practically demonstrates. For a detailed discussion of such problems with respect to know-how, see Brownstein & Michaelson (2016).

As I shall argue in § 4.4, exercising one's know-how relies in part on assessing one's own acts and on assessing the situation in which one finds oneself. But this is possible only if one conceives of them in the right way, namely as *what I am doing here and now*. The general form of the individual assessments in (B 1-2) is more specific here, which can be expressed thus:

- (D) If one can engage in the activity A oneself, then the exercise of the capacity to assess A-ing leads to self-assessments of one's own individual acts and the individual situations in which one finds oneself, i.e. judgments of the forms (D 1-2).
- (D 1) My token act x in the circumstances here and now is a performance of A-ing of quality Q.
- (D 2) The circumstances here and now provide the option to perform my token act x, which is a performance of A-ing of quality Q.

This point relies on the observation that some demonstrative or indexical concepts are, in John Perry's term, '*essential indexicals*' (cf. Perry 1979). That is, the content of the relevant proposition depends constitutively on the involvement of indexical concepts.

My assessments can only play the right kind of role in my guiding of my performances if I conceive of what I am doing right now as *what I am doing right now* – rather than, say, as what that guy I currently see on the monitor happens to do at the time he was filmed. For even if what the guy I currently see on the monitor happens to do at the time he is filmed is what I am doing right now, because the monitor shows a live picture of myself, broadcast by a hidden camera, these assessments of mine would merely *happen* to be assessments of myself, and I could and would not take them into account in guiding my acts. In order to do so, I must come to realize that this guy is *me*, and that these things happen *here and now*.

Such essentially indexical conceptions of oneself have also come to be known under the name *de se*. Thus, I have shown that conceiving of oneself as oneself – that is, conceiving of oneself *de se* – is a crucial element in the explanation of the concept of know-how because an exercise of know-how requires guidance by assessments of one's own acts as one's own acts. Likewise, conceiving of the relevant circumstances as *the situation I am facing here and now* is equally crucial.¹¹ However, I will only be able to present a full account of the role of such essentially indexical assessments of self and situation in the guidance of the exercise of a competence in § 4.4.

¹¹ However, I cannot offer a detailed discussion of these notions in the context of this book. For an overview, compare Ninan (2010). For seminal contributions to the debate, compare Perry (1977; 1979), as well as Lewis (1979).

This concludes my argument for the claim that an understanding of an activity, i.e. a suitable assessment capacity, always involves propositional knowledge, and my elaboration of the kinds of propositions involved.

However, there is a final worry which I promised to address in this section. Given everything discussed said so far, one may object that I have not yet shown that a competence requires that the examples I have established actually qualify as propositional *knowledge* rather than only as true beliefs. However, the considerations discussed in § 4.1 suggest an answer. Of course, I cannot provide a detailed account of propositional knowledge here and comment on the debate surrounding this core concept. But I can give a rather straightforward argument in terms which are widely shared even despite further controversies.

My argument is, in fact, simple. If the understanding of an activity consists in the capacity to assess performances of it, and if this capacity is itself a competence which is exercised without substantial interference, then this exercise will be a non-accidentally correct act of assessing a performance. By the same token, the assessment which is thereby made will be non-accidentally correct as well, i.e. non-accidentally *true*. Such assessments are non-accidentally true because they stem from the exercise of one's reliable competence to make such assessments, thereby providing justification for its truth. Of course, this is not to say that every assessment competence is always exercised perfectly well without mistakes or interference. All my current argument requires is that somebody who possesses a competence has or can come to believe at least *some* such propositional knowledge – that is, that she is at least sometimes correct in making these assessments in her exercise of her competence to do so. And this, I take it, is ensured for conceptual reasons. As shown in § 1.4, without being at least sometimes correct under normal circumstances, it would be inconsistent to describe someone as being competent in making such assessments in the first place.

Of course, the argument just proposed is a paradigm example of the reasoning employed by defenders of virtue epistemology, particularly of *virtue reliabilism* (cf. e.g. Battaly 2008; Greco & Turri 2011). John Greco expresses one of the main strands in this family of views as follows:

S knows p if and only if S's believes the truth (with respect to p) because S's belief that p is produced by intellectual ability. (Greco 2009, 18)

While my argument is of precisely this kind, I can leave open the question if *all* propositional knowledge can be explained in this way,¹² or if this is only true in some cases. At the same time, I contend that whatever

¹² Michael Kremer argues that this is indeed Ryle's view (cf. Kremer 2016, 12–13).

plausibility this line of argument has already gained in the debate about virtue epistemology should also be granted to my reasoning.

I would like to conclude by pointing out that it is common within virtue epistemology to explain intellectual competence as merely *one* example of many other things one may be competent at doing. John Greco writes:

The idea is that knowledge is an instance of a more general and familiar normative kind—that of success through ability (or success through excellence, or success through virtue). Thus in any realm of human activity where success is possible, we make a distinction between success from ability and mere lucky success. For example, in soccer one can score a goal through ability, as when Pele unleashes a forceful, well-placed shot, and one can score a lucky goal, as when an errant pass finds its way into the net. The idea, then, is that knowledge is a species in this familiar genus. (Greco 2009, 17)

This raises a number of interesting further questions and arguably even promises a mutually illuminating discussion of the relationship between the epistemology of know-how and virtue epistemology, very much like with the debate about linguistic competences and rule-following on which I commented in § 1.6. However, I cannot pursue these interesting connections between my project and virtue epistemology any further at this point.

§ 4.4 Guidance as Responsible Control

In § 4.1, have argued that the understanding of an activity and its normative demands which underlies the competence to engage in it should be understood as the capacity to assess performances of the activity in question. Earlier, in § 2.4, I have followed Ryle in arguing that such an understanding allows a person to continue to improve her skills. Obviously, the key to this option is the fact that the capacity to assess an act as a performance of some activity involves the capacity to thereby assess one's *own* acts and options. This is what I have just discussed in § 4.3 and stated in (D) on page 120.

In this section, I will argue that such assessments of self and situation are a key element in the explanation of how an understanding of an activity can guide one's acts. It is in the light of such assessments that competent actors are in responsible control of their acts. I will begin with a first way of formulating this proposal and then go on to offer further refinements.

The proposal is this. Guidance by an understanding of an activity can be understood as *responsible control of one's acts in the light of assessments of self and situation*. More fully, to be guided by an understanding of the norms of an activity just is to *exercise* one's capacity to assess one's own

performances and one's options for performing and¹³ to *act in the light of these assessments* in order to meet the norms of the activity in question.

I have said that a competence is a special case of a reliable ability since a competence's reliability is explained by the guidance of an understanding of what it takes to do well at the activity in question. Now I can add that this requires the exercise of one's assessment capacity in order to establish, maintain and improve one's reliability. I will now go on to explain this view in detail. This is the core of what I have called *Rylean responsibilism*.¹⁴

Rylean responsibilism and its core notion of guidance as responsible control is very close to what Ryle suggests when he writes about acquiring competences by learning (cf. § 2.3). And there are several passages in Ryle's work which clearly suggest that responsible control is not only important for learning, but involved in every exercise of a competence. Ryle writes:

Trying to comply with the teaching is part of trying to do the thing, and as the child learns to do the things, he also learns to understand better and apply better the lessons in doing the thing. Hence he learns, too, to double the roles of instructor and pupil; he learns to coach himself and to heed his own coaching, i.e., to suit his deeds to his own words. [...] We are all trained in some degree to be our own referees, and though we are not, all or most of the time blowing our whistles, we are most of the time ready or half-ready to blow them, if the situation requires it, and to comply with them, when they are blown. (Ryle 1949, 148)

This suggests that the understanding which one acquires when learning how to do something – that is, the understanding of what it takes to do well in that activity – allows one to be 'one's own coach', to assess one's performances and to act on the basis of these assessments. However, all the considerations about learning discussed in § 2.3 have turned out not to be specific to those cases where one learns by being taught. Instead, they are entirely general. Thus, we can also generalize the present point in this passage of Ryle's. That is, independently of the specific way a competence

¹³ It is very important that this 'and' is only a logical conjunction. As I shall clarify later, the exercise of one's assessment capacity does not always antecede or even occur in tandem with one's being guided by this understanding.

¹⁴ I use the term 'responsibilism' with a nod to an important paper by Michael Williams (2008) which spells out precisely such an interplay between reliable performance and responsible maintenance of reliable procedures with respect to perceptual abilities and propositional knowledge. My employment of the Sellarsian notion of "ought-to-be" norms discussed in § 3.4 is very much in the spirit of Williams' use of this concept with respect to perceptual belief formation. But since the scope of my project in this book is much broader than only perceptual capacity, a topic which would also introduce a number of problems of its own, I have not developed my view closer to Williams' considerations. However, as mentioned at the end of § 4.3 with respect to virtue reliabilism, it promises to be fruitful to pursue this further.

has been acquired by learning, everybody who possesses and exercises a competence can and does ‘double the role’ of the practitioner and the critic.

This combination of practicing and self-assessing is also present in another very important passage already discussed on page 97 and earlier:

To be intelligent is not merely to satisfy criteria, but to apply them; to regulate one’s actions and not merely to be well-regulated. A person’s performance is described as careful or skilful, if in his operations he is ready to detect and correct lapses, to repeat and improve on successes, to profit from the examples of others and so forth. He applies criteria in performing critically, that is, in trying to get things right. (Ryle 1949, 29)

Clearly, being ‘ready to detect and correct lapses, to repeat and improve on successes, to profit from the examples of others and so forth’ clearly requires to ‘double one’s role’ and to combine one’s practice of doing something with one’s self-assessments in doing so. And there are several further pertinent passages where Ryle explicitly says, for example, that somebody’s exercising a competence involves acting “with care, self-control, attention to the conditions and thought of his instructions” (Ryle 1949, 45).

Thus, there is clear evidence that the core idea of responsible control is already present on Ryle’s own discussion of know-how. However, Ryle does not give a clear account of how we should spell out this view of more clearly.¹⁵ Rylean responsibilism will hopefully complete Ryle’s own account and spell out his suggestions in a coherent and plausible way.¹⁶

To introduce the notion of control, I would like to rely on a wonderfully clear and accurate characterization by Ellen Fridland. She writes:

[I]t is by an agent’s ability to respond to both expected and unpredictable environmental circumstances and to revise her strategy accordingly, that we measure skill. As such, control is evident not only in the smooth, elegant execution of an

¹⁵ In an insightful paper, Richard Parry has argued that Ryle fails to develop a clear stance on what he calls ‘self-monitoring’ because he is too occupied with the problem of Cartesian dualism. Parry develops an interpretation of Ryle’s account of know-how which is very much in line with my proposal: “Ryle’s theory of action aimed to maintain that there is a difference in kind between intelligent and nonintelligent action and that the difference is based on the claim that in intelligent action the agent monitors his action.” (Parry 1980, 390) However, Parry suggests that Ryle failed to ask the right question about this notion: “Ryle’s anti-Cartesian program led him to mislocate the logical problem. The problem is not whether the second tier monitoring activity is occult or not. Rather it is in what way the second tier monitoring activity is related to the first tier activity.” (Parry 1980, 391) I hope to do better – by spelling out how the intelligence of performances can be explained in part in terms of ‘second tier monitoring’ – i.e. guidance by an understanding of the ‘first tier activity’.

¹⁶ Very similar considerations have recently been proposed by Michael Kremer (2016, 13–16), with explicit reference to Ryle, and by Benjamin Elzinga (2016, 17–19). Unfortunately, I cannot spell out these parallels here in detail.

uninterrupted action but in the appropriate responses and recovery to variable factors, as well. As such, I submit that it is the controlled part of skilled action; that is, that part of an action that accounts for the exact, nuanced ways in which a skilled performer modifies, adjusts, revises, and guides her performance, which we must give an account of, if we are to have an adequate, philosophical account of skill. (Fridland 2014b, 3)

What distinguishes genuine competence from mere ability is the control which competent actors have over their conduct. And as I have suggested in chapter 1, such control lies in the fact that the actor has a grasp of what it takes to do well in the activity in question and attempts to live up to these norms on this basis. In this sense, competent actors are answerable to the norms of the activity they are engaging in, and they take responsibility for meeting them. In a word, *the mark of know-how is responsible control*.¹⁷

Crucially, competences are situated. Every exercise of know-how takes place as a performance of a specific person in a specific situation. Therefore, the key to understanding responsible control lies in understanding how competent actors grasp the merits and flaws of their performances and the opportunities and impediments of the situation they are facing. My proposal is that this is precisely the role which the understanding of an activity in the sense of the capacity to assess self and situation is able to play.

Rylean responsibilism holds that responsible control is exerted by *exercising* this capacity in order to reach *de se* assessments of oneself and the options here and now (cf. (D) on page 120), and by acting responsibly in the light of this propositional knowledge in order to do as well as possible.¹⁸

In what follows, I would like to spell out in more detail what this notion of responsible control comes down to. To do so, I shall rely extensively on my arguments in chapter 3. Since the capacity to assess an activity is *itself* a competence, my general considerations about the exercise of competences in general can also be applied to the exercises of this competence. And on this basis, I can show that the relation between one's assessments and one's performances is at its core a relation of conceptual foundation and explanatory priority, and not necessarily one of temporal or even causal antecedence. I shall spell out this point a little more slowly.

I take it that there are two paradigmatic ways in which one may perform in the light of one's assessments of one's own acts, i.e. of knowledge of

¹⁷ For further discussion of control, see Buskell (2015) and Fridland (2015; 2016).

¹⁸ In fact, this is well-supported even with respect to the kind of high-speed physical activities which motivate the anti-intellectualism of Hubert Dreyfus (cf. § 3.6). This has been shown most clearly by Barbara Montero (cf. Montero 2010; Montero 2013) and Ellen Fridland (cf. Fridland 2014a, 2015, 2016). For further discussion, see also Sutton *et al.* (2011), Toner (2014), Boutin *et al.* (2014) and Christensen *et al.* (2016).

the form (D 1), and in the light of one's assessments of the options in a given situation, i.e. of knowledge of the form (D 2). These ways rely on the paradigms of fully intentional actions on the one hand and completely automatic performances on the other hand, which I have discussed before, and of which I have already said that they merely mark the clearest ends of the continuum of possible degrees.

For example, I can begin by intentionally assessing my own performances as well as the relevant situation, then go on to consciously judge that such-and-such would be the best thing to do now, determining the best option in my knowledge of the form (D 2), and intentionally perform this act.

In the same way, one can also assess one's performances *in action* or *during* their performance. One may realize that one is on the way to complete a very good performance and therefore make sure to keep on track, relying first on one's knowledge about how well one is currently doing of the form (D 1). In this way, one may see how one's present course of action could be improved, comparing one's knowledge about one's current options of the form (D 2), and see to taking advantage of the best option. Likewise, one may perceive mistakes in one's performances and try to correct them.

Finally, one can also assess one's own acts *after* performing them, for example when they are completely automatic. When one performs a quick routine act and only later realizes to what extent one has succeeded in meeting the norms of the activity in question, this act may still be understood as guided by the capacity for such assessment, particularly when the later assessment leads to attempts for improvement in the future.

Crucially, the very same considerations apply to completely automatic *assessments* of self and situation rather than to one's automatically acting in the light of these. When the assessment one reaches during or after performing is also an automatic reaction to the situation and one becomes aware of it only later, it may still be a genuine exercise of the competence of assessing oneself and one's situation. This is arguably one of the most important aspects of true competence. Part of what distinguishes genuine expertise from lesser ability is that it does not always require occurrent thinking and explicit attention to certain aspects of one's performances and options. Such assessments are routinely and automatically made, thereby freeing the agent's limited cognitive capacities to attending to more finely-grained detail.¹⁹ However, as discussed in § 3.4, § 3.5 and § 3.6, these exercises are nevertheless guided by the relevant norms in the full sense.

¹⁹ This crucial role of automatic self-assessments resonates with many observations in cognitive science (cf. e.g. Bargh & Chartrand 1999). For an insightful discussion of this point with respect to the concept of skill in general, see Fridland (2014b).

Part of this explanation, the idea of unreflective automatic assessment and guidance, has also been discussed by Erik Rietveld in his work on the notion of situated normativity. He writes:

[T]he skillful individual's responsiveness to *relevant* affordances forms the core of the normative aspect of unreflective action. I will refer to such responsiveness as 'being moved to improve' by a relevant affordance or object. (Rietveld 2008, 977)

Rietveld explains this notion of an affordance's or an option's *moving* somebody to take advantage of it in terms of a specific form of assessment capacity.²⁰ He takes up examples of craftsmanship and design and follows Ludwig Wittgenstein's suggestion that experts in such activities experience their assessments as what he calls 'directed discontent'. Rietveld writes:

Directed *discontent* is related to the craftsman's ability to make all kinds of subtle discriminations instinctively and immediately in unreflective action. This enables him to see what should be done to improve the current situation or solve the problem; to perceive and act on possibilities for action (affordances). An example of directed discontent could be a situation in which the architect notices that the door in its architectural context is incorrect and immediately senses two relevant alternatives (make the door more narrow or make it higher), and, what is more, he responds immediately to the best of these [...]. (Rietveld 2008, 980)

Directed discontent immediately orientates and 'draws' the craftsman's action towards improvement or correctness. This is an example of being moved to improve by an incorrect object (in context). More specifically, directed discontent expresses the appreciation of an object (in context) as not (yet) correct. (Rietveld 2008, 982–983)

As these passages make clear, the capacity to experience directed discontent is what provides the grounds of Rietveld's notion of 'being moved to improve by perceiving affordances'. Such cases exemplify the ubiquitous importance of unreflective automaticity in the exercise of competences.

And this discussion fits in well with the proposal I am advertising. I contend that a capacity to experience directed discontent with regard to a certain object to be produced or worked upon already is a capacity to normatively assess this object as to whether or not and to what extent it satisfies the relevant norms. And we can even generalize Rietveld's characterization of affordances to improve the quality of an *object* to *any* affordances to exercise one's skills. After all, Rietveld explicitly mentions that the object of directed discontent "can be a thing [...] but also an event or a person" (2008, 967 fn. 4). It is therefore no great leap to add that the object

²⁰ As Rietveld stresses, this includes embodied concerns and an affective, emotional involvement of the person in question. But I will bracket these elements here.

of someone's directed discontent can also be her own doing. And in many other activities, this is even the only way to assess quality, simply because no objects are created – e.g. in dancing, singing, and in sports.

At this point, one might object that experiencing something as incorrect is less demanding than *conceptually* assessing it and gaining genuine propositional *knowledge*, as I maintain.²¹ Instead, the way in which competent actors adjust their behavior may rest on non-propositional affective states or image-like mental representations. Proposals along these lines have been prominently defended by Garry Young (2004; 2009; 2011) and Eva-Maria Jung (2012).²² While I regret that cannot discuss these proposals here in more detail, my general reaction to them is twofold. On the one hand, I am happy to admit that the exact nature of the relevant assessments by competent actors is still an open question and just remain neutral here, as long as these assessments play the crucial explanatory role I have reserved for them. But on the other hand, however *exactly* such assessments are represented in the mind, they *must* remain genuinely conceptual and propositional. There may be a notion of *control* which does not require concepts and propositional knowledge states, but such control would not be *responsible* control. Responsible control requires that the relevant assessments are intelligible as genuinely conceiving *of* some act *as* a good performance of doing something, or *of* some situation *as* a good opportunity for doing so. And this structure, I contend, is possible only with genuine concepts and propositional states, even if these may only be demonstrative (cf. § 4.2 and § 4.3).

Thus, I have argued that the capacity to assess one's own performances and options is what guides intelligent practice because exercising this capacity allows for the steering of one's current or future conduct such as to meet the norms of the activity in question. Such guidance can consist in the intentional choosing of a course of action, in the on-the-fly adjustments to the specificity of the situation, and in the intentional shaping of one's automatisms and routines. This is the Rylean *responsibilist* view of being guided by the norms of an activity in one's attempts to live up to them.

To conclude this section, I would like to briefly comment on the somewhat loose way in which I have been talking about guidance over the course of this book, particularly on the many kinds of *relata* which I have mentioned as involved in this relation. In the light of the account offered in this

²¹ Rietveld may prefer this view since he has no qualms employing the notion of perceiving and acting on affordances also for creatures which, intuitively, are clearly not capable of concepts, such as earthworms. For a brief discussion of the boundary between human and non-human animals with respect to know-how, see § 6.7.

²² For discussion, see also Jung & Newen (2010; 2011) and Demmerling (2013).

section, I hope that it has become clear how these ways of speaking should be understood. In particular, I have said that who or what is guided in the exercise of a competence can be expressed in at least three ways – that *I* am guided, that *my A-ing* are guided, and that, *in my A-ing*, *I* am guided. Evidently, however, the last of these three ways of talking is the crucial one, and the other two should be seen as derived from it.

As for what it is that *guides*, rather than is guided, in the exercise of a competence, I have variously said that it is my understanding of A-ing, or my understanding of the norms of A-ing, or my capacity to assess A-ing. Given my arguments in § 1.1, the first pair of these is equivalent, and given my arguments in § 4.1 and § 4.2, the third one is also equivalent with these.

Likewise, I have argued at length that the exercises of this competence lead to propositional knowledge in the form of correct assessments of one's own acts and of the options provided by the situation, and that exercising a competence involves acting in the light of this propositional knowledge. In this vein, one may also say that these options or affordances are what guides the exercise of a competence. Similarly, one may also say that the norms of an activity guide intelligent practice, as long as this is simply meant to convey that the guiding understanding of these norms is sufficiently correct.

In the end, however, the core notion of guidance in my account of responsible control identifies *actors themselves* as the ones who guide themselves in their A-ing. Responsible control ultimately rests on the fact that it is *my exercise* of my understanding of A-ing or of my grasp on the norms governing A-ing and that it is *my acting* in the light of my assessments of self and situation and of my knowledge of my options.

§ 4.5 Guidance without Regress

I have argued that competences involve the responsible control of one's act in the light of assessments of self and situation. But this seems to create a problem of vicious regress.²³ As a consequence of a Rylean account, this is deeply ironic since Ryle himself employs a famous regress objection against intellectualism. In chapter 9, I will go on to discuss and defend this argument. In this section, however, I shall elaborate how this problem arises for Rylean responsibilism. But I shall offer a number of arguments against this objection. The aim of these considerations is to show that Rylean responsibilism provides an account of guidance without regress.

²³ For a general account of the nature and structure of regress problems, see Löwenstein (2016). But I will leave formal reconstructions of the regresses discussed here implicit.

The objection can be presented as follows. Given the account on offer, it looks as if the competence to A involves an ability to A on the one hand, and an ability to act in the light of one's assessments of self and situation in the other hand. But if so, then acting in the light of these assessments would itself be something one would not only need an ability to engage in, but a full-blown further *competence*. Guiding the exercises of one's ability in this way is certainly something which can be done better or worse, efficiently or not, and so forth. It is an activity governed by norms in the sense identified in § 1.1. And being able to guide oneself in this way could not be something a competent person is *merely* able to do, but would itself have to be an intelligent performance in the sense identified in § 1.5. Thus, it looks as if the competence to A consists of an ability to A, a competence to assess A-ing, and a *further* competence to guide one's A-ings on the basis of one's assessments of one's A-ings. Clearly, this would fail to account for the concept of know-how. But my proposal does not lead to this consequence.

To begin with, Ryle is fully aware of this difficulty for his own view. In a passage the beginning of which I have already quoted on page 73, he writes:

It is now easy to distinguish the sense in which intellectual operations are higher than, and do 'govern', the exercises of other mental capacities, from the sense in which I have denied that the occurrence of intellectual operations is implied in all those descriptions we give of people's actions and reactions which embody mental concepts. Intellectual work has a cultural primacy, since it is the work of those who have received and can give a higher education, education, namely, by didactic discourse. It is what constitutes, or is a *sine qua non* of, culture. [...] [T]o describe someone as doing something which he could not have done without formerly having had a certain education does not entail saying that he must have recited all or any of these lessons just before he acted. I could not now read a Greek sentence, if I had not formerly learned Greek grammar, but I do not ordinarily remind myself of any rules of Greek grammar, before I construe a Greek sentence. I construe according to those rules, but I do not think of them. I bear them in mind, but I do not appeal to them, unless I get into difficulties. (Ryle 1949, 295–296)

I take it that the Sellarsian account I proposed in § 3.4 and the Rylean responsibilist view of guidance as responsible control in the light of assessments of self and situation I have just laid out are very much in the spirit of these remarks. My proposal explains what it means to say that guiding norms are such that "I bear them in mind, but I do not appeal to them, unless I get into difficulties". (Ryle 1949, 296)

I have promised to show why my proposal does not lead to the problem of regress. The crucial point in the line of thought just spelled out is the

premise that responsibly controlling the exercises of one's ability to A in the light of one's assessments of self and situation is an activity in its own right, i.e. a *further* activity. But this premise is false. Instead, this activity is identical with the activity A itself. Therefore, the competence to A does not consist of three elements – the ability to A, the competence to assess A-ing, and the capacity to guide the former on the basis of the latter. Instead, the competence to A already *is* the ability to responsibly control one's A-ing in light of one's assessments of self and situation in A-ing.

I take it that this claim is already plausible on intuitive grounds. Arguably then, it can already stand for itself. Furthermore, the line of thought just presented constitutes a powerful argument for this claim – a *reductio* of its negation, the claim that the allegedly different activities in play are indeed different. Nevertheless, I would like to propose two further arguments in its support, both of which are independent from these considerations.

My first argument relies on § 1.2 where I have argued that an activity can be individuated in terms of the set of norms which govern it. On this basis, I can justify the claim that the activity of A-ing is identical with the activity of responsibly controlling one's A-ing by showing that these allegedly distinct activities are governed by the very same set of norms. This, however, is revealed by the considerations presented in § 3.4, where I argued that such a set of norms always includes all three of the distinct, but interrelated norms of the forms (A 1–3) on page 97.

At a first glance, the norms which govern the activity of A-ing have the form given in (A 1) – that is, the form “In A-ing, one ought to perform X_i in C_i .” Likewise, the norms which govern the activity of responsibly controlling one's A-ing seem to have the form given in (A 3) – that is, the form “In A-ing, one ought to make it the case that one performs X_i in C_i .” But this difference in the form of the most salient norms does not amount to a difference in the activity itself. As shown in § 3.4, any set of norms which governs an activity which one may know how to do involves a norm of the form (A 1) just in case it also involves the corresponding norm of the form (A 3). Thus, one engages in an activity which is constituted of norms of the form (A 1) just in case the very same activity is also constituted by the corresponding norms of the form (A 3). Accordingly, the activity of A-ing is identical with the activity of responsibly controlling one's A-ing.²⁴

²⁴ Thus, what it takes to do well at A-ing *fully* rather than only partially determines what it takes to do well at responsibly controlling one's A-ing. Thereby, the crucial difference between this relation and the relation between A-ing and assessing A-ing is maintained. As I have discussed in § 2.2, what it takes to A well *partially* but *not* fully determines what it takes to do well at assessing A-ing and at teaching A-ing.

I shall now turn to my second argument. This relies on the following analogy between practicing an activity and teaching it:

- (4) (a) The competence to A is the capacity to guide one's current and future A-ings on the basis of one's assessments of one's past and current A-ings.
- (b) The competence to *teach* A-ing is the capacity to guide one's *students'* current and future A-ings on the basis of one's assessments of *their* past and current A-ings.

Obviously, this characterization of competence at teaching A-ing is entirely parallel to my characterization of competence at A-ing. The only difference between (4 a) and (4 b) is that the one assessed and guided in the relevant activity is *oneself* in the case of the competence to A, whereas it is one's *student* in the case of the competence to *teach* A-ing.

My argument, then, is this. If (4 b) is a plausible view of teaching A-ing, and if the notion of guidance by assessments is intelligible in *this* case, then the same is true for (4 a) as a view of the competence to A, and of the notion of guidance by own assessments which is involved *there*. If it is plausible to deny a difference in the activity of teaching A-ing and in the activity of guiding students on the basis of one's assessments of *their* performances and options, then the same holds for the false impression of a difference between the activity of A-ing and the activity of responsibly controlling one's *own* A-ings on the basis of one's assessments of self and situation.

I would also like to take this opportunity to mention that there is indeed an important social dimension of self-assessment and guidance which can be brought out with the aid of these considerations about teaching. In the passage I quoted on page 123, Ryle suggests an analogy between an external referee and an internal self-refereeing competence. This analogy is very apt. As already discussed in § 2.3, somebody who learns how to do something often does so on the basis of her teacher's testimony about how well she has been doing so far and what she can do in order to improve. But as I have already argued, such considerations about learning can be generalized because the process of improving one's competences is open-ended.

This shows that the assessments in the light of which one guides oneself do not necessarily *all* have to stem from the exercise of one's *own* assessment capacity. Instead, one may also rely on the *testimony* of, say, a sports coach or an academic instructor who also exercised her competence to assess these acts.²⁵ This social dimension leads to a substantial enrichment of the

²⁵ This has already played an important role in § 2.5, where I also referred to the important general debate about testimonial knowledge (cf. footnote 8 on page 67).

assessment capabilities of an individual, simply because others may have an even better understanding, or an understanding more to the point of the current problem, or simply more time to reflect or a better angle to observe.

One may worry that this social dimension of assessment constitutes a problem for my account since my coach's assessments of my performances are clearly not my *self*-assessments anymore. But this is beside the point. For I am still the one who is responsible for what to do about them. If I accept the assessment of my performances proposed by my coach, either because I understand her reasons and come to share this judgment on my own terms or even because I take her word on her expertise, then I still come to possess a genuine *de se* self-assessment. And as just discussed in § 4.4, only this form of assessment is what explains my guiding myself.²⁶

§ 4.6 Understanding without Regress

The account of guidance as responsible control just discussed in § 4.4 and defended in § 4.5 finally brings me back to a problem already highlighted at the end of § 4.1. In this section, I shall propose a way to solve it.

The problem is this. If a competence involves guidance by an understanding of what the activity in question demands, and if this understanding is the capacity to correctly assess performances of that activity, then this capacity is clearly a *further* competence. Assessing performances of an activity is itself an activity which is governed by norms (cf. § 1.1), and doing well in assessing requires an understanding and guidance by these norms (cf. § 1.5). Thus, one cannot block this regress problem by claiming that assessment capacities are mere dispositions or mere abilities.²⁷ Instead, a competence involves a further competence as a guiding understanding.

²⁶ What if somebody *merely* takes somebody else's word on authority and *never* exercises her own judgment save in the decision to blindly follow these instructions? If this happens in the process of acquiring the competence in question in the first place, I take it that this phenomenon is neither problematic nor uncommon. It simply illustrates that every instance of teaching and learning has to pass the vague borderline between not yet possessing a competence and already possessing it. Taking these suggestions to the extreme, one may also imagine a pair of people who deliberately 'team up' in such a way that one blindly follows the other's instructions. But it is hard to spell out such a case more clearly, and a lot will depend on the details. Intuitively, the 'practitioner' will fail to have a genuine competence to do something because she relies on the lucky coincidence that her 'instructor' gives the right commands. Thus, we may have an example of what I will analyze as 'practical luck' in § 5.4 and § 5.5. But, again, such a scenario is too far-fetched and too unclear to present a firm view here.

²⁷ Meanwhile, Stanley & Williamson (2016) have proposed a view of know-how and skill which relies on dispositions to gain relevant knowledge, insisting that these dispositions are not or not always themselves competences. I will discuss this view in § 8.6.

But then, how should we understand the understanding which guides the exercises of *this* competence? If it also consists in a competence to self-assess, as it *must* since the first capacity is a piece of know-how, then the account I have proposed seems to lead to an infinite regress of self-assessment competences. The competence to A would require the competence to assess one's A-ings, which would require the competence to assess one's assessments of one's A-ings, and so on. Clearly, such a view renders the idea of guidance by the exercise of any of those entirely absurd.

This problem of regress is different from the problem discussed in § 4.5. There, the regress seemed to arise from the distinction between the competence to A on the one hand and the ability to act in the light of one's assessments of self and situation with respect to one's A-ing on the other hand, where the latter also qualifies as a competence and thereby threatens to create an infinite regress. Now, the regress seems to arise from a further distinction between the competence to make these assessments of self and situation on the one hand and the understanding which guides the exercise of this competence on the other hand, where the latter also qualifies as a competence and thereby threatens to create an infinite regress.

The two kinds of regress problems just distinguished are structurally similar, but importantly different in their content and nature. In fact, my answers to these problems will also turn out to be structurally similar. In § 4.5, I denied that the two competences in question are indeed distinct and argued that they are identical. The same will also turn out to be true with regard to the second regress problem, albeit for different reasons.

Let me develop this core idea of my answer more slowly. I am committed to the view that what guides a self-assessment capacity, which may in turn guide the exercise of another competence, must also be *some* self-assessment capacity. But this does not entail that it must be an *additional* one. I would like to suggest that these are the very *same* capacities. The alleged infinite regress of assessment competences stops at the level of conceptions of activities. In a word, assessment competences are *self-reflexive*.

In order to support this view, I shall first formulate a short and intuitive idea and later spell out a total of three more detailed arguments.

To begin with, consider an example. Somebody who knows how to fish and exercises her competence to fish is guided by her understanding of what it takes to fish well – that is, by her capacity to assess her own acts against the standards of fishing. In other words, she applies the concept of good or proper fishing to her own performances. In doing so, she is guided by her understanding of what it takes to do well in these assessments – that is, by her capacity to assess assessments of fishing which involves her capacity to

employ the concept of a good assessment of fishing. What I need to show now is that these are merely two sets of exercises of the very same capacity.

My first and intuitive argument in favor of this claim is this. If you are capable of making assessments of something, you are thereby also capable of making assessments of such assessments. In order to make such meta-assessments, you only have to make a first-level assessment and go on to *compare* it with other first-level assessments. On the basis of your first-level assessment, you can judge the accuracy of other such assessments. These meta-assessments will then depend on the accuracy of the first-level assessment which is employed as a standard for others.

In order to strengthen this argument, I shall now spell out a more detailed version of these considerations, relying on the form of the judgments which are reached in assessments and in -assessments discussed in § 4.3. I have explicated their general form as (B 1) on page 116,²⁸ which I repeated below alongside (B 1'), an iterated version of (B 1).

- (B 1) Token act x in token circumstances c is a performance of A-ing of quality Q.
- (B 1') *That token act x in token circumstances c is a performance of A-ing of quality Q* – i.e. (B 1) – is, in circumstances c, an accurate assessment of x as a performance of A-ing.

These verdicts are conceptually connected in the following way.

- (E) The self-reflexivity of assessments can be expressed as the conceptual truth (E 1), an instance of which is (E 2).
- (E 1) x has quality Q with respect to y just in case *that x has quality Q with respect to y* is an accurate assessment of x with respect to y.
- (E 2) Token act x in token circumstances c is a performance of A-ing of quality Q just in case *that token act x in token circumstances c is a performance of A-ing of quality Q* is, in circumstances c, an accurate assessment of x as a performance of A-ing.

I must confess that I fail to see how (E 1) could possibly be false. When my knife is good as an instrument for cutting, then *that my knife is good as an instrument for cutting* is an accurate assessment of my knife with respect to cutting. I contend that (E 1) is a conceptual truth about what it is for something to be a good assessment. And (E 2) immediately follows from this

²⁸ Of course, there is not only the assessment of *acts* in (B 1), but also the assessment of *situations* in (B 2). But as already discussed in § 4.3, these are merely two sides of the same normative coin. And in any case, one may also construe an argument for (B 2) which is parallel to the one presented here with respect to (B 1).

claim, ensuring a conceptual bridge between (B 1) and (B 1'). This shows that assessing performances of an activity and assessing such assessments of that activity are exercises of the same competence – the conceptual capacity to correctly employ the concept of that activity.

A natural source of skepticism about this inference is that, on the face of it, second-order assessments of the form (B 1') which refer to first-order assessments of the form (B 1) seem to involve a *further* concept over and above those employed in first-order assessments alone – namely, the concept of an *assessment* of A-ing. But I contend that, in fact, (B 1) *does* involve the concept of an assessment – not explicitly, of course, but *implicitly*.

I contend that every token act x in token circumstances c is a performance of A-ing of quality Q just in case, *assessed as* a performance of A-ing in circumstances c , x *has* quality Q . This shows that every assessment of the form (B 1) implicitly involves the concept of an assessment. For example, for an act to be good *as* a performance of a given activity just is for it to be good *as assessed* as a performance of that activity. The relation between a particular act and a particular activity in virtue of which the act meets the norms of that activity *already is* the assessment relation we were looking for. The competence to make assessments of the form (B 1) already entails having the concept of an assessment. Thus, the conceptual truth expressed in (E 2) shows that the competence to assess performances of an activity is one with the competence to assess these assessments.

This concludes the first of my three arguments. The second argument also proceeds by way of supporting (E 2). Here, I mainly rely on the conceptual nature of the assessment capacities, as discussed in § 4.2.

This argument draws on a principle which I shall formulate as follows:

(F) The self-reflexivity of judgment and concept use can be expressed as (F 1), an instance of which is (F 2), and a further specified version of this, in turn, is (F 3).

- (F 1) It is conceptually true that p just in case *that p* is true.
- (F 2) It is conceptually true that 'F' applies to x just in case *that 'F' applies to x* is accurate.
- (F 3) It is conceptually true that, in circumstances c , 'F' applies to x just in case *that, in circumstances c, 'F' applies to x* is, in circumstances c , an accurate assessment of x as a candidate for applying 'F'.

First, (F 3) clearly entails (E 2). Unlike (E 2), (F 3) is not concerned with the use of a specific concept of the form 'performance of A-ing of the quality Q ' and its application to a specific performance x , but instead with the use

of *any* concept ‘F’ and its application to *any* object x.²⁹ Second, (F3) can be seen as a further specified version of (F2), where (F2) abstracts away from two factors which are explicitly mentioned in (F3). On the one hand, it abstracts away from the particular circumstances of concept use and instead takes the perspective of the objective accuracy of an application of a concept. And on the other hand, it also abstracts away from the specific *acts* of judging and concept use and instead look at their *contents*. Finally, given that (F3) is supported by (F2), it is easy to see why (F2) is true in light of the fact that it follows from the uncontroversial platitude in (F1).

In short, my second argument in favor of (E2) can be summarized as follows. First, the principle concerning the objective accuracy of concept use in (F2) is true because it follows from the platitude about truth in (F1). Second, the principle concerning the accuracy of specific uses of concepts in specific circumstances in (F3) is true because it constitutes a fine-grained version of the more abstract principle (F2). And finally, the crucial principle (E2) follows from the principle just supported, (F3), as a special case for concepts of the form ‘performance of A-ing of the quality Q’.

In sum, somebody can have the competence to assess A-ing – and to reach the verdict that some performance x is a good performance of A-ing – only if she also possesses the concept of an assessment of A-ing. In virtue of this, she also has the competence to assess assessments of A-ing. For she will reach the judgment that x is a good performance of A-ing just in case she will also reach the judgment that this judgment is a good assessment of x as a performance of A-ing. The capacity to assess A-ing and the capacity to assess assessments of A-ing are merely different sets of exercises of one and the same competence – the conception of A-ing.

This concludes the second of my three arguments. The final argument I can offer relates back to the issues already discussed in in § 4.5.

There, I have argued that the capacity to A and the capacity to responsibly control one’s A-ing are different sets of exercises of one and the same competence – the competence to A. I have already commented on the difference between this problem and the current one at the beginning of this section. One of my arguments in § 4.5 was that A-ing and responsibly controlling one’s A-ing are one and the same activity because they are constituted by the same set of norms. And my current argument can also be connected with this former reasoning. The crucial premise that activities can be individuated in terms of the relevant sets of norms, defended in § 1.2 and § 3.4, also leads to a structurally similar conclusion in the current case.

²⁹ Of course, ‘F’ may also apply to an ordered set of objects x_1 – x_n . I only speak of its application to single objects x, but this is merely a shorthand.

In any assessment activity, the core norm is the norm to make *accurate* assessments. Thus, the core norm of the activity of assessing A-ing is to make an assessment of the form (B 1) on page 135 just in case this assessment is accurate. However, *that* such an assessment is indeed accurate is just what (B 1') on the same page asserts. And as I have discussed at length, the conceptual principle given in (E 2) ensures that to make a judgment of the form (B 1) is *ipso facto* to make the corresponding judgment of the form (B 1'). One meets the norms of accurately assessing a performance of A-ing just in case one also meets the norms of accurately assessing this very assessment of A-ing. Thus, assessing A-ing and assessing assessments of A-ing are exercises of one and the same conceptual assessment capacity.

§ 4.7 Self-Reflexivity without Circularity

I have argued that the competence to A is one with the competence to act in the light of one's assessments of self and situation with respect to one's A-ing (cf. § 4.5). And I have argued that the competence to make such assessments is one with the competence to assess such assessments (cf. § 4.6). Thus, the account of guidance as responsible control is indeed an account of intellectual guidance without regress.

However, the arguments brought forward so far are bound to raise further questions and create other problems. This final section of the current chapter is therefore devoted to two of the most pressing concerns. First, I shall discuss the question of the results established so far threaten the general account of guidance presented in § 4.4. And second, I will address the problem whether this makes my account of know-how viciously circular. This will also give me the opportunity to discuss and explain a phenomenon already highlighted in § 2.2 – the fact that the understanding of a particular class of activities, but only of this proper subset of activities, requires the competence to engage in them.

The first problem can be presented as follows. The view just defended has it that assessing A-ing and assessing assessments of A-ing are different sets of exercises of one and the same competence. In § 4.6, I have argued that somebody who possesses a conceptual capacity will reach the judgment that a given concept applies just in case she also reaches the judgment that this application of the concept in question is correct. But if this is true, how can the latter competence to assess assessments *guide* the first competence to produce these assessments in the first place? This would require a single competence to guide *itself*. How is this intelligible?

I think that this problem is less serious than it seems. As I have spelled out in § 4.4, an assessment competence can guide an act not only in that it causally antecedes it or takes place simultaneously. I may also exercise my competence to assess my employments of a concept later and go on to shape my habits of concept use accordingly. In such cases, the easy equivalence between reaching the judgment that a given concept applies and also reaching the judgment that this application of a concept is correct fails. For here, the circumstances in which the second-order judgment is made differ from the circumstances in which the first-order judgment was made. In reassessing my judgments later, I may have more time to think or access to new evidence or even the testimony of a peer or a teacher.³⁰

This answer brings out an important simplification with which I have been working so far. Until now, I have assumed that the token circumstances in which the second-order assessment takes place are the same token circumstances in which the first-order assessment takes place. This is why I used the same schematic letter ‘c’ on both sides in the conceptual principle (E 2). However, the circumstances in which a second-level assessment of a first-order assessment is undertaken can, but need not be the same circumstances in which that -level assessment is undertaken.

Both possibilities are crucial for my proposal. Since the circumstances are the same in *some* cases, I was able to argue that one cannot have the capacity to make first-order assessments without also having the concept of an assessment and thereby the capacity to make second-order assessments. But this capacity to make second-order assessments is not limited to those cases in which one also makes a first-order assessment. Since it is not always the case that both assessments are undertaken in the same circumstances, I can now account for the phenomenon that an assessment capacity is guided by an understanding of what it takes to make these assessments – by itself.

As I have elaborated in § 4.4, the core notion of guidance which grounds all other such notions in play here is that I am the one who guides myself in my A-ing. That I am also guided by my capacity to make such assessments is true, but only derivatively so. Thus, the fact that, in self-reflexive capacities, my A-ings and my assessments of my A-ings are different sets of exercises of one and the same competence does nothing to show that I cannot guide the former by acting in the light of the latter. However convincing my arguments in § 4.4 may have been, they apply equally to cases of guidance where first-order assessments are guided by second-order assessments of these, even if these are all exercises of one single competence.

³⁰ This seems to be particularly important in sports, where coaches and analysts play a crucial role in pointing athletes to aspects of their performances (cf. § 4.6).

In sum, assessment competences are self-reflexive conceptual competences. That is, they belong to a unique kind of know-how the exercise of which can be guided by other exercises of itself. Since these are the kind of competences which play the role of guiding the exercises of other competences, there is still no threat of an infinite regress of further competences in the explanation of any first-order skill. The alleged regress always stops at the level of conceptual capacities.

This result, however, immediately invites the second problem I promised to address at the beginning of this section. That is, the regress seems to be avoided only by running straightforwardly into a vicious form of circularity. How can the Rylean responsibilist account of know-how be defended if it bluntly *assumes* precisely what it sets out to explain?

In response to this worry, I would like to simply *restate* the account on offer, and do so in such a way that the explanandum play no role whatsoever in the explanans. Stated in this way, the explanation reads as follows:

- (G) A competence is a reliable ability the exercises of which are guided by the exercises of the reliable ability to correctly assess performances of the relevant activity and options to do so, which are in turn guided by other exercises of the very same assessment ability.

In this explanation, I have simply employed the concept of reliable ability in the explanans and avoided to employ the notion of a competence again. But why is it possible to do so? Establishing this was the whole point of § 4.6, i.e. establishing the view that an assessment competence is self-reflexive, a reliable ability the exercises of which are guided by other exercises of itself. On this basis, I contend that my explanation of the concept of know-how is plausible and defensible because it explains this concept completely in terms of the independent concepts of a reliable ability, of guidance in the sense of responsible control, and of the self-reflexivity of the guidance of a certain class of reliable abilities – namely, conceptual capacities.

This also allows me to come back to a phenomenon I have mentioned in § 2.2 and promised to explain later. The issue was this. Generally speaking, it is perfectly possible to have a *mere* understanding of an activity or *mere* knowledge about an activity without having the competence to engage in it oneself. However, there seem to be cases where this is false. A special proper subset of activities is such that understanding the activity in the relevant sense already requires the competence to engage in it. But why?

In § 2.2, I already discussed one example of such an activity when I quoted Ryle's view that the capacity to assess philosophical work requires the capacity to philosophize oneself (cf. Ryle 1949, 53–54). As I argued

there, Ryle falsely generalized from this example and held that every assessment capacity requires the corresponding skill. However, it has remained an open question how this local peculiarity can be explained and why it is that the capacity to assess philosophical work requires the knowledge how to philosophize oneself. Now, I can finally give a straightforward answer to this problem. This answer is simple. Philosophy is concerned with concepts, at least primarily and possibly even exclusively so. Of course, the philosophy of philosophy is itself a controversial field on the intricacies of which I cannot even begin to comment here. I shall nevertheless venture the assumption that the activity of philosophy is, at least largely, engaged with concepts and with their change, analysis, creation, clarification, deconstruction, improvement, and so forth.

As I have just argued, conceptual capacities are the unique kind of competence which is self-reflexive in that the competence to assess its performances is just a meta-level instance of the competence itself. Accordingly, my account directly predicts that every conceptual capacity is such that one cannot merely understand what it takes to exercise it well without also being able to do so oneself. And Philosophizing is a conceptual activity.³¹

A further such example has been proposed by David Carr:

Multiplication (Carr 1981a, 54)

To understand the rules of multiplication is *ipso facto* to be able to multiply and hence there is an air of paradox about the statement—‘he understands multiplication but cannot multiply’.

This is perfectly correct. Understanding the rules of multiplication in the relevant sense and the competence to multiply oneself cannot be separated. The account of know-how I have offered can explain this fact by saying that *performing* acts of multiplication and *assessing* acts of multiplication are two sets of exercises of one and the same conceptual competence – the use of the concept of multiplication.³²

To clarify, the problem brought up by Carr is *not* that one may understand multiplication and fail to know the meanings of expressions like ‘+’ or ‘plus’.³³ One may have no clue whatsoever about these terms, but

³¹ In § 7.1, I will discuss a further case in point by John Bengson and Marc Moffett. This example will be quoted in sentences (5 a–d) on page 212.

³² The view that competence at multiplication is a conceptual affair does not rule out the fact that simple machines like pocket calculators *can* multiply, too. However, their capacities do not constitute genuine competences, but remain on the level of mere dispositions or mere abilities. I will come back to these problems in § 6.7.

³³ A thought experiment along these lines underlies Saul Kripke’s famous interpretation of Ludwig Wittgenstein’s remarks on rule-following (cf. Kripke 1982). On the relationship between these issues and the question of know-how, see § 1.6.

nevertheless understand multiplication perfectly well, maybe on the basis of learning to calculate with non-standard symbols or simply in a different language. In such a case, the person in question would both understand multiplication and have the competence to multiply. She would simply be unable to exercise her competence with the aid of certain tools, which she fails to understand, but which are, in fact, entirely optional. A different sign with the same meaning would do the job as well.

This can be explained very well on the basis of the fact that an understanding of multiplication involves the concept of this activity. As I have argued in § 4.2, there is an important distinction between concepts and the words or symbols which express them. One may therefore possess an understanding of multiplication and thereby possess the concept of this activity, but nevertheless fail to be able to identify a foreign symbol which expresses it. On the other hand, understanding multiplication and thereby having the concept of this activity is already sufficient for knowing how to multiply.

This concludes my discussion of the self-reflexivity of conceptual assessment competences, the final element of my Rylean responsibilist account of know-how. I have argued that Ryle correctly explains know-how as an intelligent ability and I have explained this notion as a reliable ability the exercises of which are guided by the competence to assess performances of this ability and options to do so. This latter competence is itself a piece of know-how, but it can be understood in terms of its self-reflexivity, the fact that its exercises are guided by further exercises of itself.