

Part Two

Cases, Language, and Intellectualism

Man muß manchmal einen Ausdruck aus der Sprache herausziehen, ihn zum Reinigen geben, – und kann ihn dann wieder in den Verkehr einführen.

(Wittgenstein 1977a, 504)

Sometimes you have to take an expression out of the language, to send it for cleaning, – then you can put it back into circulation.

MS 117 156: 5.2.1940

(Wittgenstein 1977b, 44)

Chapter 5

The Practical Nature of Know-how

In Part One of this book, I have presented an account of know-how which does justice to the crucial phenomenology of intelligent practice introduced by Gilbert Ryle (cf. chapter 1). In doing so, I have sometimes discussed specific cases and examples, arguing why they should or should not be treated as a case of know-how. Beginning with *Ryle's Range of Cases* on page 14 as an intuitive starting point for the scope of the concept of know-how, I have later employed further specific examples. I have used *Clocks & Seals* on page 32 to distinguish genuine know-how from mere ability (cf. § 1.5) and I have used *Bela Karoli* on page 54 to distinguish know-how from a mere understanding of or knowledge about an activity (cf. § 2.1).

In this chapter, and in chapter 6, I will return to this project of accounting for the phenomena in general and for specific cases in particular. I shall present the further examples and puzzle cases which have been discussed in the debate about know-how and show how Rylean responsibilism is able to account for them. While chapter 6 will be concerned with the distinctively *cognitive* nature of know-how, this chapter deals with the *practical* nature of competences – that is, with the fact that they are a specific kind of *ability*. However, since some examples are mentioned elsewhere, it may at times be useful to consult the index of cases on page 317.

The first pair of sections in the present chapter concerns my commitment that genuine ability is necessary for know-how. § 5.1 will defend this idea against some initial objections and § 5.2 will continue to defend this view even in the face of particularly strong counterexamples which purport to establish cases of know-how without any possibility of their being exercised.

The second pair of sections discusses the claim that not all abilities are cases of know-how. § 5.3 will argue that the Rylean responsibilist account defended here can correctly explain that and why certain abilities are mere

abilities as opposed to genuine competences. In §5.4, I will then try to make the same case with respect to a particularly problematic family of mere abilities, those which involve what I shall call ‘practical luck’.

In the final pair of sections, I will connect these considerations about know-how with the debate about dispositions. §5.5 will strengthen some of the arguments from the preceding sections by discussing an analogy with structurally identical cases for dispositions. And §5.6 will discuss whether this means that know-how is itself a disposition.

§ 5.1 Ability as Necessary for Know-how

In §1.3, I have argued that having know-how requires having the ability to engage in the activity in question. This, I take it, is the common sense view of the relationship between know-how and ability. And it is what underwrites the identification of the concepts of know-how, skill and competence with which I have been working in Part One of this book. However, the view that ability is necessary for know-how has met with forceful criticism, especially by intellectualists. Already in 1974, David Brown stated:

There is an abundant supply of people who cannot do things but know how to do them. They can be found among aging athletes, neurotics, coaches, actors with stage fright, architects, and male experts on natural childbirth. (Brown 1974, 303)

However, it should be noted that part of the disagreement with my intellectualist rivals is merely verbal. As I shall argue at length in chapter 7, the English expression ‘knows how to’ can unproblematically be used even in cases without ability. As I shall show in §7.5, the overall most plausible view is that ‘knows how to’ is polysemous and can semantically express both genuine know-how and a mere understanding of or knowledge about an activity. But regardless of these linguistic questions, I am nevertheless committed to showing what exactly my account of know-how entails about the examples which have been cited as alleged cases of know-how without ability. That is the topic of this section and of §5.2.

I shall argue that all alleged counterexamples to the claim that the ability to do something is necessary for the know-how to do so can be rejected since they can be sorted into three groups. First, there are examples which are misdescribed as lack of ability and amount to genuine know-how after all. Second, there are cases which are misdescribed as know-how in the first place and only involve mere understanding. And third, there are puzzling cases where the exercise of know-how seems to be entirely impossible. In §5.2, I will argue that these are borderline cases between the first two groups.

The first group involves cases where somebody knows how to do something but to lack the ability to do so because they cannot do this *now*:

Distance (Snowdon 2003, 9)

A group of friends want to open my safe in England while I am away in New Zealand. They ring me up and ask, somewhat tentatively given my forgetfulness, whether I know how to open my safe. I answer that I do and tell them how to do it. My claim to know how to do it is obviously true, and it is clearly unaffected by my being so far away that I am quite unable to open it myself for at least thirty six hours, and what is worse, have consumed so much Speight's Ale that I have developed a tremor in my hands preventing me, for some time, from opening safes. Part of the point of this example is to generate a sense of how *totally* irrelevant, in such a case, my own capacities for performing the action in question are to the issue of whether I *know* how to do it.

As I have discussed in § 1.4, know-how involves an ability which is reliable, but reliable only in those situations which are sufficiently normal with respect to the activity in question. In other words, a *general* ability cannot be questioned by pointing to situations in which the person does not have the corresponding *specific* ability. Thus, *Distance* is unproblematic because the current lack of the specific ability to open one's safe does not entail that one does not have the general ability to open one's safe. Cases of this sort are simply misdescribed as cases of know-how without ability. These are genuine cases of know-how involving ability without a current opportunity.

The second group of examples involves cases where somebody is said to know how to do something because they can teach others how to do so but are unable to do it themselves. Paul Snowdon comments:

To construct such examples we need to describe cases in which the subject can show, teach, or tell (or otherwise convey to) us how to do something, and hence *must* be credited with knowing how to do it, but is for some reason or other unable to do it. There is no assumption here that the presence of knowledge entails that it can be passed on by the knower, but it makes a denial of the knowledge ascription very hard when the subject can, apparently, convey the relevant information to someone else. (Snowdon 2003, 9–10)

However, as discussed in § 2.1, there is an important distinction between the exercise of one's know-how and a different way of manifesting one's understanding of an activity. The knowledge how to engage in an activity and the knowledge how to teach others how to do so share the understanding of what it takes to perform well, but these competences are still distinct. Most importantly, one may know how to teach how to do something without having the competence to do so oneself, as exemplified by Fridland's paradigm case *Bela Karoli* on page 54.

Two much-discussed examples along these lines have been proposed by Stanley & Williamson, crediting Jeffrey King, and by Bengson & Moffett.

Ski Instructor (Stanley & Williamson 2001, 416)

[A] ski instructor may know how to perform a certain complex stunt without being able to perform it herself.

Quintuple Salchow (Bengson & Moffett 2007, 32, 32 fn. 2)

[A] figure skater might know how to perform an extremely difficult jump, such as a quintuple salchow, though she cannot actually do the jump herself. [...] The salchow (pronounced *sal-kow*), named after the Swedish skater Ulrich Salchow, is a figure skating jump with a takeoff from a back inside edge and landing on the back outside edge of the opposite foot after one or more rotations in the air. The quintuple salchow would then require five complete rotations in the air. To our knowledge, no skater has ever landed a quintuple salchow.

I hold that the coach in *Ski Instructor* merely knows *about* how to perform the stunt, that she possesses a mere understanding of how to do so, but does not know how to perform it *herself*.¹ Likewise, the figure skater in *Quintuple Salchow* merely knows *about* how to perform a quintuple salchow, she possesses a mere understanding of how to do so, but does not know how to perform this maneuver *herself*. Still, the intuition remains that it is in most contexts perfectly acceptable to say of these people that they ‘know how to’ do these things, simply because it would be sufficiently clear that what is meant is not genuine competence, but merely ‘know-how’ in the sense of knowledge about or understanding of an activity. Again, this intuition will be discussed in chapter 7.

Bengson & Moffett have also offered a more detailed version of *Ski Instructor* in an attempt to support the intuition that there are such cases of know-how without ability. Their case is the following:

Ski Experts Pat and Albert (Bengson & Moffett 2011c, 168–169)

Pat has been a ski instructor for twenty years, teaching people how to do complex ski stunts. He is in high demand as an instructor, since he is considered to be the best at what he does. Although an accomplished skier, he has never been able to do the stunts himself. Nonetheless, over the years he has taught many people how to do them well. In fact, a number of his students have won medals in international competitions and competed in the Olympic games. [...] [C]ontrast Pat with Albert, an unathletic (nonskiing) scientist who studies the mechanics of skiing, including but not limited to the mechanics of complicated ski stunts. As a

¹ In his more recent book *Know How* (2011b), Jason Stanley seems to change his view and to suggest that there is some sense of ability involved in this case, after all. After presenting the relevant theoretical background in chapter 8, I will comment on this issue in footnote 5 on page 272.

result of his theoretical studies, Albert knows how *one* does the stunts (namely, by contracting such-and-such muscles in such-and-such ways). Suppose that Pat, too, knows the mechanics of the ski stunts he teaches his students (he studies them in his spare time). Then Pat and Albert both know how one does the stunts; neither is able to do the stunts.

According to Bengson & Moffett's proposal, the difference between Pat and Albert is best explained in terms of know-how. They continue:

But plainly a significant difference remains: only Pat knows how to do the stunts. Indeed, even though Pat cannot do them, he grasps the stunts in a way that Albert, who only knows the theory, does not. [...] [T]his not-purely-theoretical grasp is arguably part of what *enables* Pat to teach the stunts to Olympic-caliber students. (Bengson & Moffett 2011c, 169, 169 fn. 21)

However, Bengson & Moffett mislocate the distinction between Pat and Albert. As they themselves demonstrate, this distinction can be explained entirely independently from their knowledge how to perform the stunts, but instead in terms of the *kind* of understanding of the ski stunts which Pat and Albert have and the other competences they have. Unlike Albert's, Pat's understanding of the ski stunts also constitutes a part of a *further* competence, the knowledge how to teach others how to perform the stunts.

A further case by Bengson & Moffett is also instructive. They write:

Ear-Wiggling Frustration (Bengson & Moffett 2007, 40 fn. 15)

Consider, for instance, an expert in the anatomical underpinnings of facial appendage-wiggling who knows, for any facial appendage, precisely which muscles he must flex in order to wiggle that appendage. Over the course of many years, he has used this knowledge to teach himself to wiggle his nose, which he now does with ease. An aspiring ear-wiggler, he practices wiggling his ears daily, though to his frustration he has never succeeded. Of course, he knows how to wiggle his ears: this is, in part, why his failure is so frustrating.

Again, I am not convinced. While Bengson & Moffett describe this person as knowing how to wiggle their ears, partly on the grounds that this is what explains their frustration, I contend that we can explain this frustration equally well with appeal to the fact that such an expert has a perfectly accurate understanding of ear-wiggling and nevertheless does not manage to wiggle their own ears. Independently, such examples are further complicated because ear-wiggling can plausibly be regarded as a basic action. As I have argued, basic actions do not allow for genuine know-how, but only for mere abilities. I have discussed these problems, as well as the vagueness and context-dependence involved here in § 1.5 and § 1.7.²

² In fact, ear-wiggling was one of my own examples with respect to these problems. I mentioned this in footnote 3 on page 60.

Thus, the cases in this second group are also misdescribed as cases of know-how without ability, but for other reasons than in the first group. They are not cases of know-how at all, precisely because they lack ability. Instead, they only involve an understanding of the activity in question.

§ 5.2 On Impossible Exercises of Know-how

I have argued that standard counterexamples against the view that know-how requires ability are misdescribed. Either they do entail ability, after all, and therefore also know-how. Or they involve a mere understanding of or knowledge about an activity rather than genuine know-how. In this section, I shall continue to discuss further alleged counterexamples against the view that ability is necessary for know-how against this background. These cases are particularly puzzling since they seem to involve know-how which it is impossible to exercise, at all. But without the possibility to do something at least in some scenarios, there is no sense whatsoever in which the person in question can have the relevant ability.

To begin with, Paul Snowdon offers the following case.

Impossible Pudding (Snowdon 2003, 8)

I know how to make Christmas pudding, and have done so frequently. Alas, a terrible explosion obliterates the world's supply of sugar, so that no one is able to make it. I still know how to but, like everyone else, cannot.

I think it is fairly clear that this case falls into the first group of cases discussed in § 5.1. In *Distance* on page 149, the corresponding ability could be exercised after traveling elsewhere or undergoing treatment for the tremor in one's hands. *Impossible Pudding* seems different because there seems to be no opportunity whatsoever to make Christmas pudding, given that there is no sugar in the entire world anymore. But this is only a difference in degree. The person in question still has the ability to make Christmas pudding in normal situations, which crucially includes the availability of sugar. If there is no sugar anymore, there are also no normal situations with regard to making christmas pudding anymore. But then again, it seems unlikely that it is actually impossible to produce this missing ingredient.

Thus, *Impossible Pudding* is indeed just like *Distance*. Even if the person in question is *currently* not in a normal situation with regard to the exercise of her know-how, there *could* be such a situation, and then she could do so.

Of course, this assessment of such cases is nothing new. For example, Ruth Millikan uses a case very much like *Impossible Pudding* in order to illustrate the whole distinction between general and specific abilities:

If there was a time when people knew how to make tasty dodo stew, they didn't suddenly stop knowing how on the expiration date of the last dodo. They still had the ability, but they were no longer able to apply it. True, in ordinary speech the word 'ability' may be a little fuzzy around this edge. But let us settle on using it in this unwavering manner. Abilities don't disappear just because the world is uncooperative in supplying the necessary conditions for their exercise. (Millikan 2000, 54)

But can we imagine a cognate case where normal situations have indeed become entirely impossible? I have said that I do not take it to be very plausible that nobody in the world can produce a missing ingredient such as sugar. But there are clearer examples. In fact, three families of cases have been proposed which fall under this description. The first of these families can be introduced with the following paradigm:³

Speech Impediment (Snowdon 2003, 8–9)

Susan, having spent a lifetime in the Royal household, knows how to address the queen. She can tell you that the rule is 'Ma'am to rhyme with spam, not Ma'am to rhyme with harm.' She is, however, unable herself to address the monarch correctly, (or, indeed, recite the rule in the monarch's presence) since being of a nervous disposition, she develops a speech impediment when in the royal presence, and cannot pronounce any word beginning with 'm'.

The phenomenon exhibited here can be put as follows. Somebody knows how to do something but is nevertheless unable to engage in this activity because in every situation in which she has this opportunity, some interfering factor makes it impossible to do so. Crucially, the interfering factor occurs *whenever* this opportunity arises and precisely as a *cause* of this. In *Speech Impediment*, it is a necessary enabling condition for exercising the competence to address the queen properly that the queen is present. But this is precisely what triggers Susan's speech impediment.

But does this mean that there are indeed *no* options for Susan to exercise her competence? I do not think so. Just like in *Impossible Pudding* where we only need to imagine the production of sugar in order to make it possible to make Christmas pudding again, we only need to imagine that Susan

³ A second case which can be treated completely analogously is this: "Ann is in a room at the top of a burning building. There is no escape through the door since the corridor is ablaze. The only way to escape is to climb out of the window and crawl along a narrow ledge on one side of which is a sheer drop. Ann realises that that is the only way to escape. Unfortunately, the sight of the drop has the effect of making it true that she simply cannot get onto the ledge. She is, as we say, paralysed by fear. It seems plausible to me to say in this case that Ann actually *knows* how to escape, since she certainly realises that the one and only way to escape is to crawl along the ledge, but she is unable to do that, and hence unable to escape." (Snowdon 2003, 9)

undergoes a suitable therapy in order to manage to actually go ahead and address the queen correctly. Thus, it does not seem as if this case adds anything to the discussion. The normal situations for Ann to exercise her competences appear to be somewhat more remote since undergoing therapy appears to require a more substantial change than artificially producing sugar, arguably because the former is an internal change of her character and not a merely external change in her environment. Still, there are normal situations for Ann to exercise her competences.

In sum, these are cases where there are *momentarily* no normal situations for the exercise of the respective competence because of some further element which blocks these exercises. But this impediment can be removed.⁴

This brings me to the second family of cases purporting to show that one may have know-how but no ability because there are no normal situations for its exercise. These are cases where the performances *themselves* are such that beings like us can impossibly succeed in them.

Two such examples have been discussed by Bengson & Moffett. The first, *Quintuple Salchow* on page 150, involves an ice-skating stunt which nobody ever successfully performed. In § 5.1, I have already argued that this case is best described as involving mere understanding rather than genuine know-how. Their second case is more problematic, however:

Pi (Bengson & Moffett 2011c, 170)

Louis, a competent mathematician, knows how to find the n^{th} numeral, for any numeral n , in the decimal expansion of π . He knows the algorithm and knows how to apply it in a given case. However, because of principled computational limitations, Louis (like all ordinary human beings) is unable to find the 10^{46} numeral in the decimal expansion of π .

This case is distinctive in that the relevant performances are far out of reach of any remotely normal human being. However, I think that this case is entirely compatible with the view that genuine know-how requires ability.

To see this, note that Louis is described as someone who “knows how to find the n^{th} numeral, for any numeral n , in the decimal expansion of π ” because he “knows the algorithm and knows how to apply it in a given case.” (Bengson & Moffett 2011c, 170) While Louis certainly knows the relevant algorithm, the crucial question is what it means to say that he ‘knows how to apply it in a given case’. For Bengson & Moffett’s argument to go though, this must mean that he knows how to apply it in *any* case, with respect to *any* numeral n . But is this really the true?

⁴ Readers who are not convinced by this will find further support in the analogy with examples from the debate about dispositions which I discuss in § 5.5.

Louis certainly knows how to apply the algorithm in a vast number of cases. But if the numbers are much too high for his, or anybody's, computational capacities, then it seems odd to say that he knows how to apply the algorithm *there*. Instead, it is much more natural to say that he knows what it would take to do so, but not how to do so *himself*. Thus, I propose to understand Louis as having genuine knowledge of the relevant algorithm, and as having genuine know-how, including ability, of a vast number of cases of applying the algorithm. For sufficiently high numerals, however, he has no genuine knowledge how to apply the algorithm anymore and only possesses a mere understanding of how to do so. The competence which is his know-how is excellent, but still limited in a way in which his general understanding of the relevant calculations is not.

The third and final family of cases of know-how which can impossibly exercised in practice share the feature that they involve tragic accidents:

Amputee Pianist (Stanley & Williamson 2001, 416)

[A] master pianist who loses both her arms in a tragic car accident still knows how to play the piano. But she has lost her ability to do so.

Amputee Cyclist (Hawley 2003, 23)

The cyclist knows how to ride a bike, and does so daily. After an accident, her leg is amputated, and she can no longer ride her bike. Yet it seems that she still knows how to ride her bike—the accident does not cause her to forget how to ride her bike, and she may be able to teach others.

Amputee Chef (Snowdon 2003, 8)

Raymond Blanc, the world's greatest chef, knows how to make an excellent omelette. He loses his arms in a car accident, and is no longer able to make omelettes. However, he retains his knowledge how to make omelettes, and if you wish to learn how to make an omelette you should consult Blanc. He has, that is, not lost his knowledge, merely his capacity.

One way to react to these cases is to sort them into the first group by viewing them just like *Impossible Pudding* on page 152. This is what Alva Noë proposes when he comments on *Amputee Pianist* as follows:

[O]ne might be unable to play because, even though one does know how, conditions whose satisfaction is necessary for one to exercise one's ability are not satisfied. For example, no matter how good a piano player I am, I won't be able to play piano if there is no piano ready to hand. Lacking access to a piano would mean I would be unable to play, even though I would not, for that reason, lack the relevant know-how. This explains, I think, our shared judgement about [*Amputee Pianist*]. We judge she knows how to play even though she is now unable to play, because we think of the loss of her arms as comparable (in the relevant sense) to the loss of her piano [...]. (Noë 2005, 283)

I think that this account is very plausible, but I also think that Noë misses part of the force of these cases. To clarify, this account is plausible because people can sometimes retain abilities without further exercise and practice. For example, I have retained my knowledge how to play soccer for many years without playing soccer even once. True, my skill has grown much worse without practice, but I have retained it nonetheless to some degree. The same may be true in the amputee cases.

However, Noë's view misses an important aspect of these examples. In *Impossible Pudding*, we can easily imagine a way in which the crucial necessary enabling condition might come to be fulfilled again in the future – namely, by the production of artificial sugar. And, as I have argued, the same is true for *Speech Impediment*. But in the amputee cases, it is much harder to imagine such a possible future development. The possibility of sophisticated prosthetics notwithstanding, it is quite a stretch to maintain that somebody without the necessary limbs has the ability to do things like cycling and playing the piano. What gives these cases their intuitive force is that the possible situations in which the relevant necessary enabling conditions are met are *not*, or at least not plausibly, *real* possibilities in the future, but they are *merely counterfactual* possibilities.

Katherine Hawley is very explicit about this consequence. First, she assesses her own example *Amputee Cyclist* in a way which is entirely analogous to Noë's verdict of *Amputee Pianist*. She writes:

The apparent counterexample dissolves when we distinguish different tasks more carefully. There are two tasks here—the ordinary task of riding a bike with two legs, and the much harder task of riding a bike with one leg. Both before and after the accident, the cyclist knows how to perform the ordinary task, and does not know how to perform the more difficult task. But both before and after the accident, the cyclist has counterfactual success in the ordinary task, and not in the more difficult task. Even afterward, if she were to try to ride a bike, under the circumstances of having two legs, she would succeed in riding a bike. [...] The same strategy will handle apparent counterexamples regarding coaches (who know how to perform under the circumstances of being younger, fitter or more agile), prisoners (who know how to perform if free) and those who lack material resources (who know how to perform given resources). (Hawley 2003, 23)

Hawley also sides with Noë in assimilating the *Amputee* cases to *Impossible Pudding* and thereby sorts them into my first group of examples. And she explicitly puts alleged counterexamples involving coaches into the same category. But as I have argued, such examples like *Ski Instructor* on page 150 form a distinct group of cases. These are not examples of know-how involving ability which merely seem to lack opportunity, as in the first group.

Instead, the second group consists of cases of mere understanding the relevant activity which merely *seems* to involve know-how. Thus, the line of argument proposed by Noë and Hawley shows too much. It threatens to blur the distinction between the two groups of cases I carefully distinguished. In the end, this would lead to the loss of the crucial distinction between full-blown know-how and mere understanding of an activity I defended in § 2.1.

But how, then, are we to assess these puzzling examples? At this point, I would like to rely on my discussion of the vagueness and context-dependence of know-how in § 1.7. As I have argued there, it may be vague and context-dependent whether or not a person knows how to do something because it may be vague and context-dependent how exactly the relevant activity and therefore the relevant norms are individuated. Now, I shall argue that the *Amputee* cases exhibit a similar form of context-dependency.

First, to the extent that we can imagine a real possibility in the actual future where the amputated limbs are suitably replaced, we should take the inclination to ascribe know-how to these people at face value. Then, they do know how to do these things because they have the relevant ability after all. The fact that there are no current opportunities to exercise these skills is then just as unproblematic as in cases like *Impossible Pudding* on page 152. Missing limbs are not the same thing as missing ingredients, but both are unfulfilled necessary enabling conditions which could become fulfilled in the future. On this reading, the *Amputee* cases fall into the first group.

Second, however, to the extent that we are reluctant to ascribe the relevant ability precisely because of the amputation, I hold that this stems from the fact that we are equally inclined to say that the possibility of having the relevant limbs after all is not a real possibility in the actual future. Then, these people do not know how to do the things in question because they do not have the relevant ability and fall into the second group of cases distinguished above, the paradigm of which was *Ski Instructor* on page 150. They involve no full-blown know-how, but merely an understanding of the activity in question, however sophisticated this understanding may be and however fruitfully it may play a role in the distinct competence to teach.

In short, my proposal is to understand the *Amputee* cases as borderline cases between the two groups of cases distinguished in § 5.1. Either they are cases of know-how including ability, however remote the possibility of their exercise may have become. Or they are cases which merely involve an understanding of the activity in question, but no knowledge how to engage in them oneself. Depending on reasonable variation in the assessment of these cases, either of these options is plausible and can even be made more plausible if the relevant case is suitably spelled out in more detail.

§ 5.3 Ability as Insufficient for Know-how

I have argued that one can only know how to do something if one possesses the ability to do so. Beginning in § 1.5, however, I have made a distinction between abilities which do and abilities which do not amount to know-how – between competences and *mere* abilities. As I have argued in § 1.7, this is a vague and context-dependent distinction, but this is entirely adequate, for example with respect to the capacity to walk. Given this background, I shall now take up the contrast between genuine competence and mere ability again and discuss some further and more difficult examples which threaten this distinction. While this section will discuss a number of different cases, a separate section, § 5.4, will be devoted to the most puzzling kind of case, those which involve what I will call ‘practical luck’.⁵

Paul Snowdon has mentioned some instructive examples of ability without know-how. I shall start with the following:

Press-ups (Snowdon 2003, 11)

Martin is someone who can do fifty consecutive press-ups. Let us suppose that none of us here can do that. It would be, I suggest, quite counterintuitive to say that Martin knows how to do something we do not know how to do. Rather, he is, simply, stronger than we are. He is stronger, but not more knowledgeable.

On the face of it, this example involves an ability, namely, the ability to do fifty consecutive press-ups, and suggests that having or lacking this ability is not a matter of know-how, at all. If so, this is a case of mere ability. And there is a good *prima facie* reason to think so. Doing a press-up looks like a paradigm case of a basic action, something one does *simpliciter* rather than *by* doing something else. As established in § 1.5, this characterization already entails that basic actions are not the kind of thing which one may know how to do, but which one may only have the mere ability to do. However, § 1.7 has shown that the distinction between basic and non-basic actions is context-dependent. And in a context where we assess our comparative physical fitness, it is reasonable to say that doing fifty consecutive press-ups is something one may *learn* and thereby come to *know how* to do after all. At least, I shall assume this for the sake of Snowdon’s argument.

Further, *Press-ups* is also open to the two readings just discussed for *Exit* on the facing page. On the first reading, the activity in question is individuated narrowly as ‘doing fifty consecutive press-ups’ while, on the

⁵ Examples concerning the opacity of know-how ascriptions can also be construed as cases of ability without know-how. However, I will discuss these examples only later, in § 6.1 and § 6.2, where I focus more generally on the cognitive nature of know-how.

second reading, it is individuated more generally as ‘doing consecutive press-ups’. Again, I take it that the general reading is more natural but I agree that both readings are available. Thus, Martin either knows how to do something which we do not know how to do – doing fifty consecutive push-ups – and Snowdon’s question is how this can be a difference of know-how. Or Martin is the only one to possess a level of proficiency in an ability we share – the ability to do consecutive push-ups – and Snowdon’s worry is how this gradual difference can be due to what Martin *knows* how to do.

In response, recall that I have accounted for know-how as an ability to do well in an activity in virtue of being guided by an understanding of what it takes to do so. However, this does not entail that every difference in know-how must lie in the relevant understanding. Instead, it may also be a difference in the way in which this understanding succeeds in guiding the performances in question. And this, in turn, may depend crucially on a person’s bodily features which are largely independent from her understanding.

Thus, just like the coach in *Bela Karoli* on page 54, who understands performing a standing layout on beam just as well or even better than his students, but fails to have the skill to perform one himself, the people in *Press-ups* may understand doing fifty or any other number of consecutive press-ups just as well or even better than Martin, but others may still fail to have this specific skill or the same proficiency at this general skill. The difference in know-how is largely due to a difference in bodily features.

In sum, a difference in know-how can be due to a difference in the understanding of the activity in question or due to a difference in the guidance by this understanding which, in turn, can be largely due to a difference in the bodily enabling conditions for such guidance to be successful.

Apart from *Press-Ups*, Snowdon has also offered the following cases:

Exit (Snowdon 2003, 11)

A man is in a room, which, because he has not explored it in the least, he does, as yet, not know how to get out of. In fact there is an obvious exit which he can easily open. He is perfectly able to get out, he can get out, but does not know how to (as yet).

Rock (Snowdon 2003, 11)

There is an irregular and rather narrow opening in a rock. S, who is fairly agile and thin, can certainly get through it. If, however, he has no knowledge of the rock or the task it would be odd to say that he knows how to get through it.

According to Snowdon, the people in *Exit* and *Rock* lack the relevant know-how because he lacks some crucial propositional knowledge about the situation he finds himself in. This, I take it, is indeed a possible verdict of these

cases. Likewise, the man in *Exit* does not have the knowledge how to exit the room because he fails to know what it would take to do so – i.e. he lacks an understanding of this activity which could guide him. And the person in *Rock* does not know how to get through the opening because she fails to have an understanding of doing this.

However, I think that a second verdict is even more plausible. In most contexts, it is more natural to individuate somebody's know-how with regard to more *general* activities such as arguing certain *kinds* of cases in court or escaping from rooms of a certain *kind*, as opposed to overly *specific* activities such as arguing *this* case in court or exiting *this* room. It would be strange to say that a lawyer does not know how to argue a certain case in court, a case of a kind for which she is an absolute expert, merely because she has not yet opened the file on her desk. Instead, it seems more natural to say that she has the competence to argue this case, but merely needs to acquire propositional knowledge about the specifics of how to exercise this know-how with regard to that case. Likewise, it seems natural to say that the man in *Exit* does have the know-how to exit the room after all and that the person in *Rock* does know how to squeeze through the opening, simply because they are generally competent in such movements,⁶ even if they will only be able to exercise this know-how after learning something about the specifics of the situation. But acquiring such situation-specific knowledge before acting on it is at the heart of the exercise of competence for independent reasons (cf. § 4.4).

I contend that a further example deserves a very similar assessment:

Sight-Reading (Snowdon 2003, 11)

During an evening of music I sight-read the accompaniment to a song by Wolf that I had not seen or heard before. It is clear that I was able to do that, and, indeed was able to do it well before I actually did it. But it would seem quite incorrect to say that I *knew how* to sight-read that piece.

This case parallels *Exit* in that the person in question merely lacks certain knowledge and can therefore plausibly be seen as lacking the relevant know-how. But again, it is more natural to understand the person in question as possessing the relevant know-how after all, while merely failing to be acquainted with a novel situation in which this know-how can also be exercised. Thus, the person in *Sight-Reading* does know how to sight-read the song because he knows how to sight-read such songs in general. The fact that he did not know this particular song beforehand makes no difference.

⁶ Part of the peculiarity of *Exit* stems from the fact that moving out of rooms may plausibly be at the borderline between basic and non-basic action (cf. § 1.5 and § 1.7).

A more elaborate case which concerns essentially the same point has been proposed by Bengson & Moffett:

Kytoon (Bengson & Moffett 2011c, 172–173)

Chris forms the desire to build a kytoon—a lighter-than-air kite that may, like a balloon, be filled with gas (e.g. hydrogen, hot air, or helium). She has never built a kite before, let alone a kytoon. But she is very good with her hands and thus is confident in her ability to make one. Seeking information about how to build a kytoon, information she currently lacks, Chris goes online and performs a Google search for “building a kytoon.” She finds a Web site with instructions. The instructions are long, but she is able to understand and follow each step with a modest amount of effort. Over the course of the next few days, she succeeds in executing the steps. The result of her efforts is her own personal kytoon, which she then proceeds to learn to fly.

Crucially, Chris does not learn how to build a kytoon first and then proceed to exercise this newly acquired competence later. Instead, she follows the steps she finds in the instructions, and has already begun to build a kytoon when learning more about how this is to be done. Bengson & Moffett conclude that “at the time of her initial decision, Chris is reliably able to intentionally φ (build a kytoon), but at the time of her initial decision, she does not know how to φ (build a kytoon).” (Bengson & Moffett 2011c, 173)

Given Chris’s manual dexterity, her quick uptake and the reliability of the sources she proceeds to consult, I think it is indeed possible to describe her as having the reliable ability to build a kytoon even at the time of her initial decision. Lacking a grasp of that activity altogether, however, she does not possess the competence to do so at that time. But I contend that this verdict is not mandatory. In the light of my discussion of cases along the lines of *Exit* on page 159, it is easy to imagine a continuum of cases where different amounts of information have to be acquired on the fly. Where virtually every piece of information is lacking, the verdict that a person is reliably able to do something she nevertheless does not know how to do will be clearest. By contrast, where only a minimal amount of information is lacking, it is much more natural to say that the person in question does know how to do the job after all.

At the one extreme, suppose that Chris has never built anything on her own and that she never even read a manual for building something, but that she still happens to be extremely reliable in learning such things. I take it that this variation on *Kytoon* supports the verdict proposed by Bengson & Moffett even more clearly. Such a person may happen to have the reliable ability to build a kytoon which she has never exercised so far, but it does not seem correct to say that she knows how to build a kytoon.

At the other extreme, suppose that Chris is already very good at building kites and handles manuals with ease. She has already heard a lot about differences between kites and kytoons, but still has to look up a handful of facts while she works. This second variation on *Kytoon* illustrates that having to look up only a small amount of information is entirely compatible with having full-blown know-how. Maybe Chris has to look up only one piece of information, and maybe she considers her own notes instead of a Google search. In this case, it would be far too strict to maintain that she did not know how to build a kytoon already.

Luckily, I do not have to take a firm stance on these questions here. I intended to show that cases like *Kytoon* are a further group of borderline cases between mere ability and full-blown competence (cf. § 1.7). There is no clear threshold for how rich the understanding of an activity has to be and how much refinement of that understanding in the course of action is compatible with saying that somebody does know how to do something.

In sum, I have argued that my proposal for explicating know-how as a reliable ability guided by understanding does well at drawing the line between mere ability and full-blown know-how, including in borderline areas.

§ 5.4 Practical Luck

While I have already defended the distinction between mere ability and genuine competence for a range of cases in § 5.3, there is a final family of problematic cases for this distinction which deserves a separate treatment. The paradigmatic example of such cases stems from Bengson & Moffett:

Lucky Salchow (Bengson & Moffett 2007, 46)

Suppose that Irina is seriously mistaken about how to perform a salchow. She believes incorrectly that the way to perform a salchow is to take off from the front outside edge of her skate, jump in the air, spin, and land on the front inside edge of her skate. (The correct sequence is to take off from the *back inside edge* and land on the *back outside edge* of the opposite foot after one or more rotations in the air.) However, Irina has a severe neurological abnormality that makes her act in ways that differ dramatically from how she actually thinks she is acting. Whenever she actually attempts to do a salchow (in accordance with her misconceptions) this abnormality causes her to reliably perform the correct sequence of moves. So, although she is seriously mistaken about how to perform a salchow, whenever she actually attempts to do a salchow (in accordance with her misconceptions) the abnormality causes Irina to perform the correct sequence of moves, and so she ends up successfully performing a salchow. Despite the fact that what she is doing and what she thinks she is doing come apart, she fails to notice the mismatch. In

this case, it is clear that Irina is (reliably) able to do a salchow. However, due to her mistaken belief about how to perform the move, she cannot be said to know how to do a salchow.

I think that this verdict is entirely accurate. The fact that Irina's understanding of performing a salchow is completely misguided shows that she does not have the knowledge how to perform a salchow. Luckily for Irina, her misconception does not undermine her ability to perform the salchow since this shortcoming is cancelled out by the further shortcoming of her neurological condition which happens to cause her to perform the salchow whenever she intends to perform what she mistakes for the salchow. I propose to say that, in such a case, a person has *practical luck*.

This is also one of two plausible assessments of the following example:

Annoying Smoker (Hawley 2003, 27)

[C]onsider Susie, who likes to annoy Joe, and believes that she does so by smoking. In fact Joe is annoyed by Susie's tapping on her cigarette box, which she does whenever she smokes. Susie would succeed in annoying Joe if she tried, but it seems that she does not know how to annoy Joe, perhaps because she misconstrues the situation.

I take it to be reasonable to diagnose this case as an example of practical luck. But maybe it does not involve a sufficiently reliable ability in the first place. John Williams has argued that Susie in *Annoying Smoker* is not relevantly more reliable than Sally in *Avalanche* on page 26 where the relevant ability is clearly not sufficiently reliable to count as know-how (cf. J. Williams 2008, 122). Likewise, Susie does not seem to have a sufficiently reliable ability to annoy Joe. She may at some point smoke without tapping her cigarette box, or tap a different cigarette box which is not annoying at all. In such cases, she would not succeed in annoying Joe by smoking anymore. Thus, there seems to be a continuum between cases of practical luck and examples without sufficient reliability.

These difficulties notwithstanding, there are clear enough cases which exhibit the phenomenon of practical luck. But what to make of them? After proposing *Lucky Salchow*, Bengson & Moffett comment:

We take it that examples such as this one show that errors in understanding how to ψ are sufficient to undermine know-how attributions even when the corresponding abilities are in place. Abilities, it seems, are at most reliable dispositions to intentional behavior, whereas know-how involves some degree of understanding. We believe that this insight is the key to a general philosophical theory of know-how. (Bengson & Moffett 2007, 46)

Again, I am entirely in agreement. As a matter of fact, I have already argued for this key insight into the nature know-how and I even made the case that Ryle already saw this point very clearly (cf. § 1.5). However, it remains an open question why it is *exactly* that cases of practical luck fail to amount to know-how. For unlike *mere* abilities like my ability to digest, examples of practical luck *do* involve an understanding of the activity in question, which is even *responsible* for the ability's reliability.

As I shall show, cases of practical luck involve two elements both of which are incompatible with know-how. As Bengson & Moffett already pointed out, they involve a substantial *misunderstanding* of the activity in question. Further, they also rely on a substantially *deviant* form of guidance by this inaccurate understanding. The fact that these shortcomings cancel each other out in the end does not show that cases of practical luck are examples of know-how after all. Let me comment on these two elements in turn.

Intuitively, one may understand an activity but still be wrong about specific aspects of how the activity works exactly. One does not need to have a perfect grasp of all the intricacies of it, and one does not need to be able to describe one's understanding perfectly (cf. § 2.2). As I spelled out in § 4.1, in order for one's conception of an activity to count as an actual understanding, it must be sufficiently accurate overall. But examples of practical luck clearly violate this constraint. In *Lucky Salchow*, Irina's understanding of the salchow is substantially misguided because she explicitly believes that the salchow must be performed in a way in which it cannot be performed at all. And in the borderline case between failure of reliable ability and practical luck *Annoying Smoker*, there is sufficient reason to take the understanding involved to be substantially misguided, too. Susie severely misconceives her activity of annoying Joe because she mistakes her smoking to be the relevant factor rather than her tapping her cigarette box.

This brings me to the second element of practical luck, a substantially deviant form of guidance by an inaccurate understanding. This is what makes the difference between cases which lack reliable ability in the first place and examples of practical luck which do involve reliable ability, but for the wrong reasons. In the paradigm example *Lucky Salchow*, Irina's misunderstanding of doing the salchow only leads to her reliably performing the salchow because she *fails* to be guided by her actual conception of what she should do. If, counterfactually, she were to manage to get rid of her neurological problems in some way, and then began to be *properly* guided by her *actual* conception of the salchow, she would clearly fail to perform a salchow since she drastically misconceives what it takes to do so and would correct her performances in the wrong direction.

Something similar can also be argued for in the case at the borderline between practical luck and lack of reliable ability, *Annoying Smoker*. Here, Susie is only reliable at annoying Joe because her being guided by her actual conception of what she should do is accidentally coupled with a further fact such that it accidentally leads to her annoying Joe. If, counterfactually, she were *properly* guided by her conception and her conception *alone*, she might at some point also smoke a cigarette or a pipe *without* tapping a cigarette box, maybe because she has none at hand. Then, she would fail to annoy Joe, but also fail to correct her performances in the right direction, namely, towards tapping a cigarette box.

Following Bengson & Moffett, I have argued that cases of practical luck undermine know-how despite the fact that they involve reliable ability explained by an understanding of the activity in question. This is because the understanding involved is not sufficiently accurate, and because this understanding explains the person's performances in the wrong way, leading to the accidental and lucky reliability of the person's performances. Thus, even if I disagree with Bengson & Moffett's positive intellectualist account of know-how, I follow their insistence on the crucial role of cases like these.

Couched not in terms of guidance, but in terms of intentional action, Jason Stanley's response to *Lucky Salchow* is very similar. He writes:

Irina has a false belief about how to do the Salchow, and she is lucky enough that whenever she intends to do the Salchow, she succeeds. Though she intelligently and successfully performs the Salchow, she does not *intentionally do the Salchow* when she succeeds, anymore than it follows that I *intentionally win the lottery* when I win the lottery after buying a lottery ticket intending to win. Of course, when Irina performs the Salchow, she does it with the intention of performing the Salchow, and there is a causal connection between her intention to perform the Salchow and performing the Salchow. But as we have learned from Davidson, F-ing with the intention of F-ing does not entail intentionally F-ing, even when there is a causal connection between one's intention to F and one's F-ing. In order to intentionally F, there must be the *right kind* of causal relations between one's intention to F, and one's F-ing, and those are lacking in Irina's case. (Stanley 2011c, 218)

It seems as if this answer can easily be translated into the account I have proposed, simply by adding the view that intentional action and guidance are conceptually connected. However, this would not be completely sufficient because it does not cover cases of entirely automatic exercises of know-how – an issue I have discussed in detail in chapter 3. Thus, a person who benefits from practical luck may also exercise her alleged know-how entirely automatically rather than as an intentional action. In contrast to

Stanley's proposal, my account can handle even the automatic exercises of an alleged competence which is in fact undermined by practical luck. For the two criteria I mentioned are independent from the question if the exercises of the ability are intentional or automatic. A substantial misconception of the activity and a substantially deviant form of guidance to cancel it out can also play a role when an ability is exercised automatically.⁷

Thus, a treatment of practical luck in terms of understanding and guidance should be preferred over Stanley's view. His proposal may be plausible locally for the limited range of cases which involve intentional action, but the heart of the explanation does not lie in intentionality, but instead in the idea of a guiding understanding (cf. Bengson & Moffett 2011c, 172–174).

§ 5.5 Masks, Finks, and Mimics

In this section, and in § 5.6, I will discuss an analogy between know-how and ability on the one hand and dispositions on the other hand. As I will argue, this analogy will add credibility to some of my assessments of the cases over the course of this chapter. For my take on these cases of know-how is entirely analogous to the widespread view about dispositions which are masked, finkish or mimicked. In this section, I shall present these notions from the literature on dispositions and spell out the analogy. § 5.6 is then devoted to the question what this analogy shows and what it presupposes.

Let me introduce this analogy with reference to an example discussed in § 5.2. I have argued that *Speech Impediment* on page 153 is a case in which a necessary enabling condition for exercising the relevant competence contingently also causes a further element to interfere with the exercise of the competence.⁸ Cases of this sort have gained some prominence in the debate about dispositions. Mark Johnston writes in a seminal paper:

Consider a fragile glass cup with internal packing to stabilize it against hard knocks. Packing companies know that the breaking of fragile glass cups involves three stages: first a few bonds break, then the cup deforms and then many bonds

⁷ The paradigmatic example for automatic exercises of competences in § 3.1 was *Unwelcome Sign* on page 78. On this model, one may imagine a case in which I possess the ability to read a certain script automatically, while misunderstanding it entirely, but nevertheless associating every symbol with the right meaning in the end because I also automatically confuse everything. In such a case, I would still have practical luck and no full-blown know-how, because my understanding of the script I read is mistaken, and because, if, counterfactually, I were to guide myself properly by the way I think the script is to be read, I would 'correct' my performances in the wrong direction.

⁸ Again, everything I add about *Speech Impediment* here also applies to the parallel case involving acrophobia which I mentioned in footnote 3 on page 153.

break, thereby shattering the cup. They find a support which when placed inside the glass cup prevents deformation so that the glass would not break when struck. Even though the cup would not break if struck the cup is still fragile. The cup's fragility is masked by the packing which is a) something extrinsic to the glass cup and b) causes the glass cup when struck to withstand deformation without breaking. (Johnston 1992, 233)

This makes for an obvious analogy between dispositions and know-how and ability. Dispositions on the one hand and skills and other abilities on the other hand are similar in that both involve the possibility of manifestation. Just like the fragility of a thing can manifest itself in that it breaks, the competence to play chess can manifest itself in that one plays chess. However, both the manifestation of know-how and the manifestation of a disposition can go awry. And in both cases, something can go wrong *systematically*. One way in which this can happen is by what Johnston calls a 'mask' and what others, synonymously, call 'antidote' (cf. Bird 1998). A *masked* disposition, ability or competence is one where the enabling conditions for its manifestation are such that they *always* prevent its manifestation.

Crucially, however, a masked disposition is nevertheless a disposition. A cup does not cease to be fragile because of the protective packaging around it. Exploiting this analogy, I shall say that Susan in *Speech Impediment* has the *masked competence* to address the queen correctly. Her know-how is not undermined by the fact that her speech impediment systematically prevents her from exercising it.

True, Susan's know-how is of no help to her in actually confronting the queen. But this does not show that she does not have the relevant skill. Instead, it would take *additional*, indeed quite remarkable competences for her to succeed after all. Addressing the queen correctly in spite of a serious speech impediment is a much greater accomplishment than doing it without one. This would require more than the ordinary skill to address the queen correctly, which Susan has. In the same way, some cups may be so fragile that even normal protective packaging would still not prevent their breaking. Evidently, this does not mean that an ordinary fragile cup does not remain fragile even if the same protective packaging prevents its breaking.

There are also further features of dispositions that have interesting counterparts in the debate about know-how. I have said that a *masked* disposition, ability or competence is one where the enabling conditions for its manifestation are such that a further element prevents its manifestation, while the respective disposition, ability or skill nevertheless remains in place. A closely connected phenomenon is the so-called 'finkishness' of dispositions (cf. Martin 1994). David Lewis explicates this term as follows:

[A] stimulus s itself might chance to be the very thing that would cause the disposition to give response r to stimulus s to go away. If it went away quickly enough, it would not be manifested. In this way it could be false that if x were to undergo s , x would give response r . And yet, so long as s does not come along, x retains its disposition. Such a disposition, which would straight away vanish if put to the test, is called *finkish*. A finkishly fragile thing is fragile, sure enough, so long as it is not struck. But if it were struck, it would straight away cease to be fragile, and it would not break. (Lewis 1997, 144)

Generalizing from this explication of finkish dispositions in the same way as I generalized from the above explication of masked dispositions, we find a further way in which manifestations of a competence can fail systematically. A finkish disposition, ability or competence is one where the enabling conditions for its manifestation are such that they cause the respective disposition, ability or competence to vanish while these conditions obtain.

Katherine Hawley has explicitly discussed this with respect to know-how:

Panic (Hawley 2003, 25)

It looks as if knowledge-how may be finkish: Sylvia knows how to get home from the city center, but if she were in the city center she would have a panic attack and forget how to get home. If she were to try to get home from the city center she would not succeed.

This example of a finkish competence is importantly different from the masked competence of Susan in *Speech Impediment* on page 153. Sylvia in *Panic* forgets what to do altogether, unlike Susan, who still knows what to do, but is unable to do so. But what are we to make of this?

To begin with, it seems as if finkish know-how is simply a peculiarity, but not a problem for any account of know-how, at all. According to the definition of finkishness, a finkish competence is a skill which is sometimes present and sometimes lost. And in general, it is just a fact of life that people gain and lose skills. But what makes finkish competences special is that these competences are lost and regained *systematically*. Just like it is a fact of life that the exercise of a competence can be blocked because of some defeating condition while masked competences are nevertheless peculiar because a further element arises as a defeating condition *whenever* the enabling conditions are satisfied, so it is also a fact of life that competences can be lost and regained but finkish competences are nevertheless peculiar because a further element removes the competence altogether *whenever* the enabling conditions are satisfied. This is also the reason why I have not included finkish competences in my earlier discussion in § 5.1 and § 5.2. They do not purport to show that somebody can have know-how without ability. It is part of their definition that, at the relevant time, they have neither.

However, one may wonder if this is precise enough in order to distinguish finkish skills from masked skills. Katherine Hawley responds to her own case *Panic* in the same way in which I responded to *Speech Impediment* in § 5.2:

It is uncontentious that ordinary dispositions may be finkish [...]. But we need not accept that knowledge-how may be finkish. [...] Sylvia's knowledge-how matches her counterfactual success. Sylvia does not know how to get home from the city center under circumstances of being prone to panic attacks. She does know how to get home from the city center under circumstances which are normal for most people [...]. (Hawley 2003, 25)

Along these lines, Sylvia retains her knowledge how to get home after all. Her panic attack does not show that this know-how is finkish, it simply masks it. Sylvia's know-how is an ordinary piece of know-how for which a panic attack is certainly a defeating condition, not a remarkable piece of know-how which can also be exercised in such a drastic situation.

However, I think that Hawley's assimilation of the finkishness to the masking of competences blurs a nevertheless important distinction. I shall instead propose a more nuanced account for this difference between finks and masks with respect to know-how.

The crucial point is this. When a competence is *masked*, then the understanding of the activity in question remains in place, but it becomes impossible to engage in this activity for further reasons. Despite her nervous condition, Susan in *Speech Impediment* understands perfectly well what it would take to address the queen correctly and where she falls short of doing so. In contrast, when a competence is *finkish*, then this understanding of the activity vanishes altogether, at least as long as the enabling conditions for engaging in that activity obtain. Sylvia in *Panic*, while suffering from a panic attack, does not have any understanding of how to get home from the city center. It is not that she systematically fails to perform well *despite* her understanding of what to do as in a case of masking, but precisely because she *lacks* this understanding in the first place.

This shows that a finkish competence is substantially different from a masked competence. But I think that Hawley's assimilation of finkishness to masking is nevertheless plausible in a different way. In order to see this, consider the question how clear the notion of an understanding's vanishing and reappearing actually is. Does it really make sense to describe Sylvia in *Panic* as losing her understanding of how to get home altogether and regaining it a short time later? Instead, I contend that it makes perfect sense to describe a finkish skill as a case where it is indeed true that *something* is masked, but it is the *understanding* which is *involved* in this skill rather

than the *whole* skill. In these terms, Sylvia in *Panic* still possesses an understanding of how to get home even when she is suffering from a panic attack, but she cannot remember and employ this understanding anymore, precisely because it is masked by her panic attack.

The upshot of these considerations is that *know-how* is *finkish* just in case the *understanding* involved in know-how is *masked*. As I have argued in § 4.1 and over the course of chapter 4, these states of understanding are *themselves* competences. Thus, the possibility of masking such a state of understanding is built right into the account I have offered.

This concludes my discussion of finkishness and masking. At this point, I would like to note that I have given this account without commenting on the otherwise controversial question if what makes it the case that an object's disposition is finkish or masked must be an extrinsic rather than intrinsic property. On an intuitive construal of this otherwise complicated distinction, all the examples I discussed involve intrinsic finks or masks. However, I must confess that I find this phenomenon entirely unproblematic, both when it comes to know-how and when it comes to mere dispositions. A convincing defense of this position has been proposed by Lauren Ashwell (2010), who also discusses a number of plausible examples. Without further discussion, I shall simply quote one of these examples as an argument in favor of the view that intrinsic masks and finks are sufficiently unproblematic for my present purposes. Ashwell writes:

A berry could be poisonous, although due to its indigestible skin it is prevented from having a poisonous effect on those who ingest it. An intrinsic property of the berry—its having an indigestible skin—prevents the berry from having its noxious effect. When it is skinned, it is not that the berry *becomes* poisonous—its disposition to harm those that ingest it is simply unmasked. (Ashwell 2010, 636)

Next to the possibility of masks and finks, there is also a third feature of dispositions which can fruitfully be exploited in my analogy. The classic reference for this phenomenon is Arthur David Smith, who writes:

Imagine a sturdy block of wood which is tapped sharply at *t*. At *t*, also, a special ray (let us call it a Z-ray) is beamed at the block causing it to splinter in just the way in which one would have expected the block to have splintered had it been fragile. Perhaps it will be objected that this is not a fair counterexample because [...] [there is no] connection between the event mentioned in the antecedent and that mentioned in the consequent. This point may easily be satisfied, however, [...] by having the beaming of the Z-ray on to the block of wood caused by the tapping of the block. Here the condition [...] is fulfilled, with a real connection between the two events [...], and yet, *ex hypothesi*, the block of wood is not fragile. (Smith 1977, 440)

Such examples have come to be called ‘mimics’ (cf. Lewis 1997). The block of wood does not have the disposition of fragility, but the Z-ray mimicks that disposition for the block of wood. Whenever something which would cause a fragile thing to break occurs, this causes the Z-ray to be beamed at the block and the block thereby to splinter. Despite the fact that something happens which could be a manifestation of the block’s fragility whenever a fragile thing would manifest the same phenomenon, the block itself is still not fragile.

Just like in the cases of masked dispositions and finkish dispositions, this explication of mimicking dispositions can also be generalized. *Mutatis mutandis*, this yields the definition of what I have called ‘practical luck’ in § 5.4. When a disposition, ability or skill is *mimicked*, the object or person in question does not possess the disposition, ability or skill at all, but whenever the enabling conditions for their manifestation are fulfilled, a *further* factor causes them to behave in just the same way it would behave if the disposition, ability or skill were to manifest itself. To illustrate this with my paradigm example *Lucky Salchow* on page 162, Irina does not have the competence to perform a salchow, but in the right circumstances, the combination of her misunderstanding and her neurological problems causes her to behave just as if she were to exercise such a competence. And just like it is clear that the mimicking in Smith’s example does make it the case that the block of wood has the disposition of fragility, it is, as I have argued in § 5.4, equally clear that Irina’s misunderstanding and her neurological disorder do not make it the case that she knows how to perform a salchow.

This concludes my discussion of the scope of the analogy between dispositions on the one hand and know-how on the other hand. I have argued that the phenomena of masking, finkishness and mimicking, as they are called in the debate about dispositions, can be found analogously in the debate about know-how and competence, as well.

§ 5.6 Is Know-how a Disposition?

The analogy I laid out in § 5.5 brings back a question which has already been present in this book from the very start: Is know-how *itself* a disposition? Gilbert Ryle explicitly stated that “[k]nowing *how*, then, is a disposition” (Ryle 1949, 46), as I quoted on pages 48 and 51. The most straightforward explanation of the analogy just discussed is that this is indeed the case. In this section, I shall discuss the question what this view would entail and whether the account presented here is really committed to it.

To begin with, there is a lively debate about the concept of a disposition (cf. Choi & Fara 2012), and an equally lively debate about the concept of an ability (cf. Maier 2010, Clarke 2015). One of the questions under discussion is whether or not *abilities* are a kind of disposition (cf. e.g. Clarke 2009; Whittle 2010; Vetter 2016). If abilities are indeed a special kind of disposition, then my analogy between know-how and dispositions would follow straightforwardly from my declared view that *know-how*, or competence, is a special kind of ability. Thus, the question if *know-how* is a disposition can be answered by discussing the question if *ability* is a disposition.

However, nothing I have said presupposes that this is the case. Alternatively, the structural features of dispositions with regard to masking, finkishness and mimicking, and the analogous structural features of abilities and competences, could also be explained in another way. For example, all of these concepts are *modal powers or potentialities*, they involve the possibility of something's happening under certain conditions. Arguably, these structural features are only due to this common modality and entirely independent from the question if, given that competence is a species of ability, ability is also a species of disposition.

One way to see this is by appreciating that these common modal features are often understood in terms of *conditionals*, and that conditionals exhibit such features quite independently from their use in an account of dispositions. In the literature, such observations have indeed first been made independently with respect to conditionals (cf. Shope 1978) before they gained prominence in the debate about dispositions (cf. Martin 1994). Thus, in my account of the specific kind of ability which is know-how, I can indeed bracket the nature of dispositions and their relation to abilities.

But this creates a complication for my view. Masks, finks and mimicks are typically employed to show that a satisfactory account of dispositions cannot be one in terms of conditionals. David Lewis sketches such a proposal – the ‘simple conditional account’ – as follows:

A fragile thing is one that would break if struck; an irascible man is one who would become angry if provoked; and so on. In general, we can state the *simple conditional analysis* thus: Something *x* is disposed at time *t* to give response *r* to stimulus *s* iff, if *x* were to undergo stimulus *s* at time *t*, *x* would give response *r*. (Lewis 1997, 143)

The failure of this simple conditional account of dispositions in the face of masks, finks and mimicks has become common knowledge in the debate about dispositions (cf. Choi & Fara 2012). Examples of these phenomena such as those presented in §5.5 are direct counterexamples to this view.

However, Gilbert Ryle is canonically understood as one of the paradigm proponents of something like what Lewis calls the ‘simple conditional analysis’ of dispositions, for example in the survey article by Choi & Fara (2012), quoting from chapter V of *The Concept of Mind* (Ryle 1949). If so, it seems as if my proposal is inconsistent, at least as a proposal for how Ryle can plausibly be interpreted.⁹

However, the Rylean account of know-how which I have proposed is in fact entirely independent from this question. Rather than in terms of dispositions or conditionals, I have *straightforwardly* spoken of ability and possibility, without the additional premise that abilities are a kind of disposition and without any further commitment as to what dispositions are.¹⁰

As it turns out, this is entirely faithful to Ryle’s declared view. True, he says that dispositions are understood in terms of conditionals – for example, in that both are ‘inference tickets’ (cf. e.g. Ryle 1949, 119; Ryle 1950, 329–330). But this does not *contrast* with straightforward talk of modality. For Ryle forcefully denies any substantial difference between conditionals and modal statements:

[T]he differences between modal and hypothetical statements is in fact purely stylistic. There is only one colloquial way of correctly negating the superstitious hypothetical statement “If a person walks under a ladder, he comes to grief before the day is out,” namely by saying “No, a person may (might or could) walk under a ladder and not come to grief.” (Ryle 1950, 335)

An ‘if-then’ sentence can nearly always be paraphrased by a sentence containing a modal expression, and vice versa. Modal and hypothetical sentences have the same force. [...] There is only a stylistic difference between the ‘if-then’ idiom and the modal idioms. (Ryle 1949, 122–123)

Thus, I can maintain my view, and maintain it even as a plausible interpretation of Ryle. The features of know-how I have discussed are independent from the question whether the modal nature of this concept is analyzed in terms of dispositions.

I have argued that know-how is a kind of ability, but that I can remain neutral on the question whether ability is a kind of disposition, despite the fact that know-how and abilities share a certain modal profile with dispositions, a profile in virtue of which they also exhibit phenomena like masking, finkishness and mimicking. At this point, one may complain that I

⁹ In this case, there would even be further reason for doubt since it is even less clear if, simple cases notwithstanding, the relevant *multi-track* dispositions can be understood in terms of conditionals at all (cf. Manley & Wasserman 2008; Vetter 2013).

¹⁰ In this respect, I follow Katherine Hawley, who also brackets the issue of dispositions, even if she does appeal to conditionals (cf. Hawley 2003, 25).

was only able to do this because I have employed a largely *unanalyzed* notion of ability. While I have relied on this concept from the very beginning, in chapter 1, I have only presented certain general characteristics of abilities, such as the distinction between general ability and specific opportunity (cf. § 1.4), but remained silent on what it actually *is* to possess an ability. By contrast, if I were to finally *give* an account of ability, how may I possibly circumvent the appeal to dispositions?

In reply, I would first like to point out that I regard it as a virtue of the Rylean responsibilist account defended here that it only undertakes those commitments concerning the nature of ability which are strictly necessary. I take it that the discussion in Part One of this book was perfectly intelligible even without a *full* account of ability.

Second, however, and more substantially, whatever I say about ability is so far constrained only in one crucial respect – abilities cannot be understood along the lines of the simple conditional analysis of dispositions. Thus, I may opt for a view of abilities which does not resort to any such a view. Or I may even maintain that abilities *are* dispositions, but offer an alternative view of dispositions, one *not* in terms of conditionals. Such a metaphysics of dispositions and potentialities has recently been presented by Barbara Vetter (cf. Vetter 2014; Vetter 2015). Thus, this is one plausible option for me to take.

However, Vetter has also argued *against* the view that *abilities* are dispositions (cf. Vetter 2016). In the remainder of this section, I would like to discuss these considerations, as they are particularly instructive with respect to my project of accounting for know-how.

Vetter discusses two versions of the view that abilities are dispositions. The first version relies on some form of the conditional analysis of dispositions, where the relevant stimuli for the disposition are intentions, desires, wishes or something along these lines (cf. Fara 2008; Vivhelin 2004). However, Vetter argues that such accounts fail, among other things because some exercises of abilities do not rely on wishes, intentions or desires, at all (cf. Vetter 2016, sec. 2). This is evidently much in the spirit of my discussion of the automatic exercise of know-how (cf. § 3.1) and I agree that such accounts of ability in terms of dispositions are implausible.

But the second version of this view is much closer to my own positive account of know-how, in particular to my appeal to normativity and normative guidance (cf. chapter 1). In the words of Ernest Sosa, this is the idea that “[c]ompetences (and abilities) are dispositions to succeed” (Sosa 2015, 95). The claim Vetter extracts from virtue reliabilists like Ernest Sosa and John Greco (cf. e.g. Greco 2007; Sosa 2010) reads as follows:

x has the ability to A iff x is disposed to A successfully when A'ing at all, i.e. iff, if x were to A at all, then (interferences aside) x would (probably) A successfully. (Vetter 2016, sec. 3)

While this approach is perfectly able to handle exercises of abilities without wishes, intentions or desires, Vetter argues that it fails for other reasons. It is unable to account for abilities to engage in activities without any clear standard of success. Her examples are aimless ambling around or doodling. She writes:

These are performances which, by their very nature, are without any aim, and presumably without any evaluative standard. (There are, it may be said, no better or worse ways of ambling or doodling; that's the very point of them.) (Vetter 2016, sec. 3)

Vetter considers the reply that activities like these *do* have standards of success, after all, but that there is *no gap* between engaging in them and meeting those standards. Success at doodling just is to doodle in some way or other, and success at ambling aimlessly just is to amble around anywhere and in any way. Then, however, the analysis of abilities as dispositions to succeed will be undermined because the conditional that “if x were to A at all, then (interferences aside) x would (probably) A successfully” will be *tivially* true, as soon as such an activity is substituted for ‘A’. This entails the absurd consequence that such abilities are possessed by everybody (cf. Vetter 2016, sec. 3). Moreover:

The same kind of consideration would seem to apply quite generally to simple motor abilities: the ability to move my eyes, bend my leg, wiggle my foot, and so on. All of these performances are such that to perform them at all is already to perform them successfully. (Vetter 2016, sec. 3)

Does this mean that the second version of the view that abilities are dispositions – the idea that abilities are dispositions to succeed – is also bound to fail? Probably not. As Vetter also points out, there are plausible versions of this idea which can, or at least purport to, solve these problems (cf. Maier 2013; Jaster 2016). But even if this objection should prove to be fatal, I would like to point out that it is only fatal to the account of *abilities* in terms of dispositions to succeed, but not to the view that *know-how* or *competence* is a disposition to succeed, which I would still be able to maintain.

To see this, it is crucial to note that all of Vetter's examples are such that we can easily talk of ability, but where it sounds odd to attribute know-how. I *can* amble aimlessly, I am *able* to do so, but I do not *know* how to do it – i.e. I am not guided by an understanding of what it takes to do well

in ambling aimlessly when I do amble aimlessly. This point is even clearer with respect to the simple motor abilities Vetter mentions. Moving my eyes, bending my leg and things like these are clear cases of *basic actions*. As I have argued in § 1.5 and § 1.7, abilities to perform basic actions are *mere* abilities – abilities which do *not* amount to competences. In contrast, cases of genuine know-how all admit of a more substantial characterization of the normative standards of the relevant activity (cf. chapter 1). Thus, the view that *know-how* is a disposition to succeed is safe from the objection against the view that *all* abilities are such dispositions.

One consequence of this discussion is that virtue reliabilists like Sosa and Greco may defend themselves against arguments like Vetter's by carefully framing their view in terms of genuine *competence* as opposed to merely ability. In a different context, Vetter notes that Sosa in fact mainly uses 'competence' rather than ability (cf. Vetter 2016, sec. 3). Thus, the Rylean responsibilist account of know-how as competence presented here may prove congenial to the aims of virtue reliabilism. Conversely, I have myself already relied on a virtue reliabilist argument in the development of my account – in order to show that the individual assessments of acts of an activity which one reaches by exercising one's assessment competence amount to genuine propositional knowledge (cf. § 4.3). However, I cannot explore these questions here in more detail.

Chapter 6

The Cognitive Nature of Know-how

Rylean responsibilism sheds light on the examples and puzzle cases which have been discussed in the debate about know-how, and offers plausible explanations as to where and why something is a genuine case of know-how. In chapter 5, I have spelled this out in detail with respect to the *practical* nature of know-how. In this chapter, I shall continue to do so, but focus on the cognitive, semantic and epistemic aspects of competences.

The first pair of sections will address the semantic fact that ascriptions of know-how are opaque, quite unlike ascriptions of ability. In § 6.1, I present these phenomena and offer an account of them. This will lead, in § 6.2, to the question whether there is such a thing as merely *de re* know-how.

The second pair of sections will turn to specifically epistemic aspects of know-how, focusing on the question of epistemic luck. § 6.3 will show that know-how cannot be undermined by certain forms of epistemic luck, which poses a severe problem for the view that know-how simply consists in propositional knowledge. § 6.4 will offer an explanation of this phenomenon and argue that it is compatible with my Rylean responsibilist account of know-how, even if this explicitly includes propositional knowledge.

In the end, I turn to a number of themes from cognitive science. § 6.5 discusses the notion of procedural knowledge. § 6.6 offers an account of two famous clinical cases which have been cited in the debate about know-how. Finally, § 6.7, discusses the question of know-how outside of adult humans.

§ 6.1 The Opacity of Know-how Ascriptions

A number of philosophers have recognized that ascriptions of know-how may be opaque. In this section, as well as in § 6.2, I will discuss this phenomenon and show how Rylean responsibilism can account for it.

To begin with, opacity is well-known and well-discussed when it comes to propositional attitude reports (cf. McKay & Nelson 2010). Let me illustrate it with an example which is arguably less worn-out than many others.¹

Klement

Let us assume that international investigators searching for undercover Nazis in South America in the late 1950s had a comprehensive overview of recently emigrated Europeans and their whereabouts. Among other things, let us assume that they knew that a certain Riccardo Klement works as an electrician for Daimler-Benz in González Catán in Argentina. Also, they were looking for the former SS-Obersturmbannführer Adolf Eichmann. As it happens, Riccardo Klement *is* Adolf Eichmann. Eichmann used this name on a forged Swiss passport in order to escape to South America. However, the investigators did *not* know that *Eichmann* works as an electrician for Daimler-Benz. This is true despite the fact that Klement *is* Eichmann and they knew that *Klement* worked there. They did not know that Klement is Eichmann.

As this example makes clear, an ascription of a propositional attitude, such as propositional knowledge, may be correct despite the fact that an ascription of a slightly altered propositional attitude involving a different, but co-extensional expression is *incorrect*. In *Klement*, the only substitution was that of one proper name for a co-extensional proper name.

As many philosophers have remarked, something similar is the case in ascriptions of know-how. John Williams offers two paradigmatic examples. Here, I shall focus on one of these cases, as everything I say until then is equally applicable to both. But I will present the other example, *Triangles*, on page 186. The crucial case, then, is this:

Superman (J. Williams 2008, 110)

Lois may know how to contact Clark Kent (she knows that she has his telephone number) yet not know how to contact Superman, despite the fact that contacting Clark is necessarily contacting Superman (because Clark is necessarily Superman). Since Lois has the ability to contact Clark, she has the ability to contact Superman, although she might not know that she has the latter ability.

In this scenario, Lois has the ability to contact Clark Kent, and she has the ability to contact Superman since to contact Kent would *be* to contact Superman. However, while she knows how to contact Kent, she does *not* know how to contact Superman, *despite* the fact that to contact Kent would *be* to contact Superman. Like the ascriptions of propositional knowledge in *Klement*, the ascriptions of know-how in *Superman* must be assessed differently, despite the fact that the knowledge concerns the same thing.

¹ Wolfgang Barz has used this case in one of his classes, and it has stuck with me.

Before Williams, Katherine Hawley has independently proposed that the following case exemplifies the opacity of know-how:

Light Switch (Hawley 2003, 26)

[T]o flip the switch is to alert the prowler, so if Sarah is able to do one of these, she is able to do the other (Davidson 1963). Yet it seems that Sarah could know how to turn on the light without knowing how to alert the prowler, if it doesn't occur to her that turning on the light will alert the prowler.

A structurally similar case, also due to Hawley, is this:

Impressive Skating (Hawley 2010, 401)

[S]uppose I devise—then master—a sequence of moves on my skateboard, and I wonder whether showing this off will impress the kids at the local skate park. As it happens, they will be very impressed if they see me perform this sequence. Do I know how to impress the kids at the local skate park? [...] I entertain the proposition that performing my routine is a way for me to impress the kids, and I entertain this proposition under a practical mode of presentation (that is, differently from the way in which a non-skating spectator would entertain this thought), but I do not know whether the proposition is true, and so I don't know how to impress the kids at the skate park.

Bracketing the notion of a 'practical modes of presentation' until §9.2, both of Hawley's examples can be explained in the same way. And here, I side with John Williams' remark that "opacity must involve substitutions of necessary equivalence. So this example does not establish the opacity of know-how since it only involves substitution of contingent effect for actual cause." (J. Williams 2008, 110 fn. 3) I think that these further criteria are indeed important in order to bring out the full force of the phenomenon of opacity and to distinguish it clearly from the problem of practical luck which discussed in §5.4. Thus, the core problem with *Light Switch* and *Impressive Skating* is that they are cases of practical luck. Still, such examples show how opacity may also play a role in cases of practical luck.

Williams draws the sweeping conclusion that "unlike ability, know-how is opaque. Since what is true of actual ability is not true of know-how, know-how is not actual ability." (J. Williams 2008, 110) But to deny that know-how is actual ability is only *one* way to explain why know-how is opaque, while ability is not. Instead, I shall explain this fact by explaining why the special *kind* of ability which is know-how is opaque.

As I have argued, a *mere* ability is an ability to do something well which is *not* guided by an understanding of what it takes to do so, while know-how is an *intelligent* ability in the sense that it *is* so guided. On this basis, I have already accounted for a number of puzzle cases, including what I

called ‘practical luck’ in § 5.3 and § 5.4. The opacity of know-how can also be explained along these lines. It follows straightforwardly from the opacity of the guiding understanding of an activity which is crucial for every competence. Now, I take it to be independently and intuitively plausible that states of understanding admit of opacity. But the account of understanding presented in chapter 4, especially with respect to the conceptual nature of understanding as discussed in § 4.2, makes this case even stronger.

Given these resources, we can account for *Superman* as follows. Lois has the ability to contact Clark Kent, and she has an understanding of what it takes to do so. When she exercises this ability, she is guided by this very understanding. Depending on the situation, she will choose a suitable method to reach Kent (say, by phone). She will be able to learn more about which of these methods is indeed helpful in which circumstances and to do better next time. True, Lois also has the ability to contact Superman. She even has an understanding of what it takes to do so. Say, she knows that she only needs to call his number or send him a message. Crucially, however, she does *not* possess the ability to contact Superman *in virtue of* her understanding of what it takes to do so, but *independently*. And when Lois exercises her ability to contact Superman – say, by calling the number written in her notebook next to the name ‘Clark Kent’ –, she is *not guided* by her understanding of what it takes to contact *Superman*. She understands what it would *take* to contact him all right, but she cannot *act* on this understanding in order to do so since she simply does not know which number to call. Thus, this is not a case of know-how but of mere ability.

In this way, I can also account for a further case which illustrates the opacity of know-how. This was discussed much earlier by Thomas Steel:

Philadelphia (Steel 1974, 43)

Among the many people who know how to get to Philadelphia there is at least someone, we may imagine, who does not know that Philadelphia is the City of Brotherly Love, and, consequently, does not know how to get to the City of Brotherly Love, even though he does know how to get to Philadelphia, and Philadelphia is that city.

Unlike in *Superman*, the reason why the people in this example are not guided by an understanding of the activity in question is not that they have such an understanding, but fail to be guided by it, but that they do not even have such an understanding in the first place. The people described in *Philadelphia* do not know the City of Brotherly Love and therefore do not understand how to get there. But this does not make a difference for the present problem. Both cases involve mere ability as opposed to know-how

because the ability in question is not an ability to do something well in virtue of being guided by an understanding of what it takes to do so.

This, however, does not conclude the problem of the opacity of know-how. Thomas Steel had already noticed a further phenomenon which John Williams fails to mention. He sees that there is at least *some* sense in which the problematic ascription of know-how is accurate, after all. He continues:

But the action of getting to Philadelphia is identical with the action of getting to the City of Brotherly Love. Thus, among the many people who know how to get to Philadelphia there is at least someone who both knows how to get to Philadelphia and also does not know how to get to Philadelphia. (Steel 1974, 43)

If this is correct, then the opacity of know-how ascriptions must pertain to only one of *two senses* which these ascriptions have. Steel writes:

As in regard to statements about belief, there is an ambiguity in statements of the form ‘S knows how to F’. Following the procedure which is sometimes made use of in connection with the belief statements, we might say that such statements may be taken either *de dicto* or *de re*. More specifically, I suggest that ‘S knows how to F’ is ambiguous as between: (1) ‘S has knowledge how with respect to the doing of that action which is known to him as F-ing’ – what we might call its *de dicto* sense; and (2) ‘The action of F-ing is such that the doing of it is something S has knowledge how with respect to’ – what we might call its *de re* sense. (Steel 1974, 47)

In order to fully appreciate this proposal, let me start with an explication of what the distinction between *de re* and *de dicto* comes down to with regard to belief. In his seminal paper on *de re* belief, Tyler Burge writes:

From a semantical viewpoint, a *de dicto* belief is a belief in which the believer is related only to a completely expressed proposition (*dictum*). *The epistemic analogue is a belief that is fully conceptualized.* That is, a correct ascription of the *de dicto* belief identifies it purely by reference to a “content” all of whose semantically relevant components characterize elements in the believer’s conceptual repertoire. What is the appropriate epistemic characterization of *de re* belief? I think one should explicate the notion simply in terms of *the negation of our epistemic characterization of de dicto belief.* (Burge 1977, 345–346)

Accordingly, the investigators in *Klement* on page 178 believed, but believed only *de re*, that Adolf Eichmann works as an electrician for Daimler-Benz in González Catán in Argentina. Since they did not conceive of Riccardo Klement *as* Adolf Eichmann, their *de dicto* belief that Klement works as an electrician for Daimler-Benz in González Catán in Argentina cannot be described as a *de dicto* belief that Eichmann works there, but only as a *de*

re belief to that effect. Belief *de dicto* requires that the person in question conceives of the elements of her belief *in the very way* in which they are characterized in the relevant ascription, while *de re* belief only requires that she conceives of the elements of her belief in *some way* or other.

This proposal also makes perfect sense in the case of know-how, and my account of know-how as an intelligent ability can explain this fact. As I have argued, the exercises of know-how are guided by an understanding of the activity in question (cf. § 1.5), an understanding which crucially involves concepts (cf. § 4.2). Accordingly, I can explain Steel's distinction between *de re* and *de dicto* know-how in the very same way in which Burge explains the corresponding distinction for belief. That is, know-how *de dicto* requires that the person conceives of the activity in question in the very way in which it is characterized in the ascription used. And know-how *de re* only requires that she conceives of this activity in some way or other.

This explains why Steel is inclined to say that there is a sense in which the people described in *Philadelphia* know how to get to the City of Brotherly Love – the *de re* sense. They have the ability to get to the City of Brotherly Love and an understanding of that activity, but an understanding in which the City of Brotherly Love is *not* conceived of in *these* terms, but rather under the name 'Philadelphia'. Likewise, Lois in *Superman* knows how to contact Superman in the *de re*, but not in the *de dicto* sense. She has the ability to contact Superman and an understanding of this activity, but she is *not* guided by an understanding of contacting Superman conceived of in *this* way, but rather conceived of as contacting *Clark Kent*.

Thus, the fact that ascriptions of know-how are opaque is perfectly intelligible. Correspondingly, it is no surprise that there are cases where somebody ascribes mere *de re* know-how to somebody else. This latter phenomenon has not been noted at all in the current debate about know-how. I shall therefore discuss it more thoroughly in the next section.

§ 6.2 Is there *de re* Know-how?

What is the status of *de re* know-how? In the case of belief, there is widespread agreement that mere *de re* belief is nonetheless genuine belief. And as I have already argued in § 2.5, § 4.2 and § 4.3, know-how can involve mere *de re*, merely demonstrative conceptions of the kinds of performances which are assessed with respect to their quality as acts of exercising the relevant activity. Does this mean that mere *de re* know-how is just as unproblematically a case of genuine know-how as mere *de re* belief is a case

of genuine belief? In this section, I shall argue that this is not the case. *Genuine* know-how must be *de dicto*.

The key to see this is that my previous arguments only concern the question how to conceptualize the individual performances which are assessed with respect to the whole of the activity of A-ing. These may well be *de re*. However, the examples just discussed in § 6.1 are cases where the activity of A-ing *as a whole* is conceptualized *de re*. But this is impossible without the loss of genuine know-how. Competence at A-ing requires conceiving of A-ing *as* A-ing, i.e. *de dicto*.

To show this, I shall go back to § 3.3, where I argued that a necessary condition for know-how is that it is possible to *intentionally* engage in the activity in question. But it is hard to imagine the people in the cases under consideration as intentionally doing what they have merely *de re* know-how to do. In *Superman* on page 178, Lois cannot intentionally contact Superman because, if she were to intend to contact Superman, she would have no idea what to do. Likewise, the people described in *Philadelphia* on page 180 cannot intentionally get to the City of Brotherly Love because, if they were to intend to get to that city, they would not know where to go.

Of course, this is not quite correct. As Donald Davidson has famously argued, these examples are nevertheless intentional actions. He writes:

[A] man is the agent of an act if what he does can be described under an aspect that makes it intentional. What makes this answer possible is the semantic opacity, or intentionality, of attributions of intention. Hamlet intentionally kills the man behind the arras, but he does not intentionally kill Polonius. Yet Polonius is the man behind the arras, and so Hamlet's killing of the man behind the arras is identical with his killing of Polonius. (Davidson 1971, 46)

However, the fact that such acts are intentional because they are intentional under *some* description, even if they are not intentional under the description currently used, does not undermine my argument to the effect that *de re* know-how falls short of genuine know-how. To show this, I shall reconsider the specific argument I gave in § 3.3 in support of the view that it is necessary for genuine know-how that it is possible to intentionally engage in the activity in question. As I shall go on to show now, these earlier considerations *also* require that every competence to A can be exercised intentionally under the description of this activity *as* A.

In § 3.3, I have argued that every competence can be exercised intentionally since every competence is *acquired* and since, therefore, competences whose exercises are not intentional, but entirely automatic, essentially depend on the intentional practice of these activities. However, one cannot

intentionally practice something in order to improve one's skill unless one also conceives of what one is doing *as* exercising that competence under the very description in question. For example, when I practice my skill at playing squash, I must conceive of my performances *as* performances of playing squash in order to have the chance to assess how well I live up to the standards of squash and where exactly it is that I need to improve. If I do not conceive of my doings as playing squash, but rather as a random exercise with a peculiar combination of ball and racket, I may certainly improve my muscles and dexterity and thereby also increase my competence at playing squash, but it would be false to say that I was *practicing* to play squash. Thus it is essential for competence that it can be exercised intentionally *and* intentionally under the right description.

This ensures that what I called mere *de re* know-how cannot count as *genuine* know-how. In *Superman*, Lois can intentionally contact Superman because she can do so intentionally under the alternative description of contacting Clark Kent. But she cannot *practice* how to contact Superman because this would require her to intend to contact Superman under that very description and this is something she does not do because she believes that she cannot contact Superman. The same holds for the people described in *Philadelphia*. They can intentionally get to the City of Brotherly Love because they can do so intentionally under the alternative description of getting to Philadelphia. But they cannot *practice* how to get to the City of Brotherly Love since this would require them to intend to get to the City of Brotherly Love under that description, which they evidently would not do while trying to get to what they conceive of as Philadelphia.

I have argued that only *de dicto* know-how is genuine know-how because the understanding which guides the exercises of such a competence must conceive of the activity in the same way in which the know-how is characterized. With the aid of this result, we can also understand a related phenomenon which has been discussed in a further treatment of opacity by David Carr. First, he presents the following further paradigmatic case:

Gray's Elegy (Carr 1979, 407–408)

Suppose a famous dancer was to perform before an audience, an item from his repertoire to which he has himself given the following title: [.] [a] performance of Improvisation No. 15 [...] [T]he movements of the dancer turn out to resemble an accurate (movement perfect) semaphore version of Gray's 'Elegy', though the dancer is quite unaware of this fact. Now clearly we can use the identity thesis as a warrant for regarding ['a performance of Improvisation No. 15'] and ['a semaphore recital of Gray's 'Elegy'] as alternative descriptions of the same action (or set of actions) and [...] we can also offer either ['a performance of Improvisation No. 15']

or [‘a semaphore recital of Gray’s ‘Elegy’] as a description of what the dancer was able to do. But equally clearly, although we can describe the dancer as knowing how to bring about [a performance of Improvisation No. 15] we cannot reasonably suppose that he also knows how to bring about [a semaphore recital of Gray’s ‘Elegy’].

Evidently, this case is completely analogous to *Superman* and *Philadelphia*, which I have already discussed explicitly. However, Carr goes on to contrast *Gray’s ‘Elegy’* with the following example:

Conjurer (Carr 1979, 408)

Oddly enough, some substitutions in knowing how contexts seem to be safe enough. For example the same magician’s performance might be described as [.] [a] display of conjuring tricks or alternatively as [.] [a] performance of prestidigitation. Now it seems clear enough that if the magician knows how to bring about [a display of conjuring tricks], he also knows how to bring about [a performance of prestidigitation].

While I must confess that my intuitions are not completely settled, I nevertheless tend to agree with Carr that this is indeed an example in which the phenomenon of opacity does not arise. But why? Carr writes:

The reason why we can safely switch [a display of conjuring tricks] and [a performance of prestidigitation] in sentences about knowing how is that both descriptions refer to the same activity, but ‘same’ not in the sense of being one self-identical pattern of physical movements; rather ‘same’, in the sense of having one point or purpose. (Carr 1979, 409)

Carr’s idea that two descriptions pick out the same activity in terms of the same “point or purpose” resonates with my view that an activity can be identified by the set of norms which govern it (cf. § 1.2). It is impossible to have an accurate understanding of conjuring and to have an accurate understanding of prestidigitation *without* thereby understanding that this is the same activity. Understanding what it takes to do well in one of these *just is* understanding what it takes to do well in the other one.

This is where *Conjurer* on the one hand differs from *Superman*, *Philadelphia* and *Gray’s ‘Elegy’* on the other hand. In the latter cases, the relevant activities are described with different proper names for the same individual – ‘Clark Kent’ vs. ‘Superman’, ‘Philadelphia’ vs. ‘The City of Brotherly Love’, and ‘Improvisation No. 15’ vs. ‘Gray’s ‘Elegy’’. Like the pair of synonymous general terms ‘conjuring’ and ‘prestidigitation’, these three pairs of proper names are necessarily co-extensional – or this is how Carr construes *Gray’s ‘Elegy’*. But here, it is perfectly possible to have accurate conceptions of

the individuals referred to by such proper names *without* realizing that they are co-extensional. To know an activity as conjuring just *is* to know it as prestidigitation. By contrast, to know an activity with regard to an individual under *one* proper names differs significantly from knowing this activity under a *different* proper name because this understanding *depends* on these names as linguistic devices.²

However, this point does not essentially depend on the use of proper names in the descriptions of the relevant activities. To see this, consider a further case from John Williams:

Triangles (J. Williams 2008, 110)

[C]onsider Stan, whose job involves selecting three equal lengths of wood and then gluing them together to make equilateral triangles. Stan, who is not very bright, has the concept of equal length but has no clue what an angle is. He may know how to make equilateral triangles (by following the method above) yet not know how to make equiangular triangles (he has no idea what these are), despite the fact that making equilateral triangles is necessarily making equiangular triangles.

Again, this case is completely analogous to the earlier ones, *Superman*, *Philadelphia* and *Gray's 'Elegy'*. The only difference is that it does not involve the use of different, but co-extensional proper names, but instead the use of different, but co-extensional general terms – ‘equilateral triangle’ and ‘equiangular triangle’. However, it still differs from *Conjurer* because these general terms are not synonymous, despite the fact that they are necessarily co-extensional. These general terms express complex concepts formed on the basis of the distinct concepts of equiangularity and equilaterality, and the ability to employ these concepts is sufficiently independent from each other.

Thus, mere *de re* know-how is not genuine know-how because know-how requires a *de dicto* conception of the relevant activity. However, cases like *Conjurer* suggest that what it means to have the *right* kind of conception is best understood in terms of the *norms* governing the activity. The crucial thing is to understand of what it takes to do well in a given activity.

To conclude, I would like to consider a plausible objection to these considerations. Arguably, *de re* know-how is indeed genuine know-how, simply because this is the most extreme example of a case where somebody merely needs further information in order to acquire a genuine competence. I have discussed such examples already in § 2.5 and § 5.3, but cases of *de re* know-how are extreme in that it is only *one* single proposition knowledge of which the person in question needs to acquire, and in that this proposition is even

² Of course, I cannot discuss the semantics of proper names more thoroughly in this book (cf. e.g. Cummings 2013). But I take it to be safe enough to assume that proper names differ in this crucial respect from descriptions of activities.

necessarily true. For example, Lois in *Superman* only needs to learn that Clark Kent *is* Superman in order to acquire the *de dicto* knowledge how to contact Superman. Coming to know a single proposition which is necessarily true is, arguably, the smallest step imaginable which may separate a person from having a competence or not. Talking loosely, it is therefore perfectly fine to describe mere *de re* know-how as actual competence. Properly speaking, however, I maintain that such cases are no genuine competences.

§ 6.3 The Problem of Epistemic Luck

In the debate about know-how, many authors have looked at specific characteristics which have been established for *propositional* knowledge in order to discuss the question whether or not these are *also* present in the case of know-how. Three such features have already been discussed in the literature – the possibility of transmission by testimony, the requirement of true belief, and the question of susceptibility to epistemic luck. In this section, and in § 6.4, I shall concentrate only on the last of these problems.

Part of the reason why I bracket the other issues is that I have already dealt with them earlier, in chapters 2 and 4. With respect to testimony, it may be true that know-how differs from propositional knowledge in that it cannot be transmitted by testimony *alone* (cf. Poston 2016). But testimony certainly has an important role to play in the learning and teaching of competences (cf. § 2.5).³ With regard to true belief, I have already pointed out that my account rests on the idea of competences to assess performances rather than on beliefs about what to do when (cf. § 4.1). But as Brownstein & Michaelson (2016) have discussed in detail, one may describe various kinds of cases where somebody allegedly knows how to do something, but fails to have any true beliefs about how to do it. Here, I am happy to commit myself to the claim that they only *truly* have know-how if they *at least* have the capacity to *acquire* such beliefs as soon as they exercise this assessment capacity. And *qua* exercise of this competence, these beliefs will amount to genuine propositional knowledge, even if other beliefs about the relevant activity are entirely absent or even false (cf. § 4.3).⁴

³ For further discussion of these problems, see Hawley (2010) and Small (2014).

⁴ Yuri Cath describes a case, *The Non-Dogmatic Hallucinator*, where somebody explicitly rejects the relevant true belief about how to juggle, even though she knows how to juggle, because of credible, but accidentally misleading evidence (Cath 2011, 116–117). Examples like these are difficult for many reasons (cf. Brownstein & Michaelson 2016) and much will depend on how the story is continued from there. I tend to think that know-how is retained, but that it is *finkish* because the guiding understanding of how to juggle is *masked* by the belief in the misleading evidence (cf. § 5.5).

With respect to epistemic luck, then, there is a widespread consensus in contemporary epistemology, following one of the main lessons of the seminal paper on by Edmund Gettier (1963), that epistemically lucky true beliefs do not amount to full-blown propositional knowledge (cf. e.g. Pritchard 2005; Orozco 2011).⁵ In this section, I shall discuss cases of know-how which involve epistemic luck. This is a severe problem for propositionalist intellectualism, the view that know-how is a species of propositional knowledge. In §6.4, I will go on to show that my own Rylean responsibilist proposal makes perfect sense of the possibility of epistemic luck, and this despite the fact that I also hold that know-how involves propositional knowledge.

Several candidate Gettier cases involving epistemic luck have been discussed in the literature on know-how. The first of these stems from Jason Stanley and Timothy Williamson. They write:

Flight Simulator (Stanley & Williamson 2001, 435)

Bob wants to learn how to fly in a flight simulator. He is instructed by Henry. Unknown to Bob, Henry is a malicious imposter who has inserted a randomizing device in the simulator's controls and intends to give all kinds of incorrect advice. Fortunately, by sheer chance the randomizing device causes exactly the same results in the simulator as would have occurred without it, and by incompetence Henry gives exactly the same advice as a proper instructor would have done. Bob passes the course with flying colors. He has still not flown a real plane. Bob has a justified true belief about how to fly. But there is a good sense in which he does not *know* how to fly.

This verdict strikes me as false. As Ted Poston and Yuri Cath have pointed out, it is intuitively much more plausible to say that Bob may have been lucky in how he acquired his know-how, but that it is nevertheless genuine know-how that he has acquired (cf. Poston 2009, 744; Cath 2011, 125). Poston and Cath have also proposed further cases in point:

Nolan's Curve (Poston 2009, 745)

Nolan wants to be a good pitcher and needs to learn how to throw a fastball, breaking ball, slider, and curve. He has good reason to believe that Mike is a great pitching coach. Mike doesn't know much about baseball and even less about how to pitch. Nevertheless, Mike tells Nolan that to throw a fastball he

⁵ I remain neutral on the question whether the phenomenon of epistemic luck leads to genuine 'Gettier cases', that is, to cases of justified true beliefs which are nevertheless not cases of propositional knowledge. While this is the most widespread view in the current debate, one may reasonably disagree about this point and hold instead that epistemic luck certainly undermines propositional knowledge, but also undermines justification (cf. e.g. Hetherington 2011a). Along these lines, the influence of epistemic luck also leads to states which fall short of being knowledge, but not to Gettier cases of justified true beliefs, and instead simply to true beliefs without justification.

should V, to throw a breaking ball he should X, to throw a slider do Y, and for a curve do Z. Mike's wrong about everything except a curve. Nolan internalizes the instructions so that he can competently carry them out.

Olavi's Finnish Tango (Poston 2009, 746)

Olavi wants to learn the Finnish tango, an established variation on the Argentine tango. He finds a website that aims to specialize in the Finnish tango. Olavi downloads the instructions and learns those instructions. Olavi, though, is very lucky to have what are in fact the correct instructions. The website is devoted to causing mass confusion about the Finnish tango by uploading different instructions each second.

I agree with Poston about both of these cases (cf. Poston 2009, 745–746). While Nolan certainly does not know how to throw a fastball, a breaking ball or a slider, he nevertheless knows how to throw a curve, despite the unreliability of his coach. And Olavi knows how to dance the Finnish tango, despite the unreliability of the website he consulted.

Yuri Cath proposes the following third example:

The Lucky Light Bulb (Cath 2011, 115)

Charlie wants to learn how to change a light bulb, but he knows almost nothing about light fixtures or bulbs (as he has only ever seen light bulbs already installed and so he has never seen the end of a light bulb, nor the inside of a light fixture). To remedy this situation Charlie consults *The Idiot's Guide to Everyday Jobs*. Inside, he finds an accurate set of instructions describing the shape of a light fixture and bulb, and the way to change a bulb. Charlie grasps these instructions perfectly. And so there is a way, call it ' w_1 ,' such that Charlie now believes that w_1 is a way for him to change a light bulb, namely, the way described in the book. However, unbeknownst to Charlie, he is extremely lucky to have read these instructions, for the disgruntled author of *The Idiot's Guide* filled her book with misleading instructions. Under every entry she intentionally misdescribed the objects involved in that job, and described a series of actions that would not constitute a way to do the job at all. However, at the printers, a computer error caused the text under the entry for 'Changing a Light Bulb,' in just one copy of the book, to be randomly replaced by new text. By incredible coincidence, this new text provided the clear and accurate set of instructions that Charlie would later consult.

Here, too, Charlie has learned how to change a light bulb, at least given the implicit assumption that Charlie does not only believe that w_1 is a way for him to change a light bulb, but that he can furthermore act in way w_1 and thereby change a light bulb. Cath continues: "The fact that Charlie is extremely lucky to read accurate (as opposed to misleading) instructions just seems irrelevant to whether or not he comes to know how to change a light bulb on the basis of reading those instructions." (Cath 2011, 117)

In sum, these cases all involve the acquisition of know-how from highly unreliable sources and thereby an instance of epistemic luck. Poston's general conclusion is therefore entirely correct: "Knowledge-how isn't constrained by the same anti-luck intuitions as propositional knowledge." (Poston 2009, 746) More specifically, Adam Carter and Duncan Pritchard have recently clarified this line of thought by pointing out that know-how is compatible with the kind of epistemic luck called 'environmental luck', a kind of luck which would undermine propositional knowledge (cf. Carter & Pritchard 2013). I shall come back to these classificatory questions in §6.4.

This provides a very straightforward argument against propositionalist intellectualism. If know-how is propositional knowledge, and if all propositional knowledge can be undermined by epistemic luck, then this must be the case for know-how, as well. Since know-how *cannot* be undermined by epistemic luck, it must be false that know-how is propositional knowledge. Since I am committed to the view that know-how involves propositional knowledge, even if it does not consist in it, I will discuss Jason Stanley's response against this charge. §6.4 will then present my own position.

In defense of propositionalist intellectualism, Stanley has suggested that Bob in *Flight Simulator* does not know how to fly, despite appearances, because he is unable to fly *intentionally* (cf. Stanley 2011c, 235 fn. 8). But this, too, strikes me as false. Bob has an entirely accurate conception of flying and is guided by this conception when he flies (cf. Cath 2011, 129). Therefore, nothing prevents him from flying intentionally, and intentionally under the right kind of description (cf. §3.2 and §6.2). Of course, Stanley is correct to remind us of the Davidsonian lesson that A-ing with the intention of A-ing is not the same thing as *intentionally* A-ing because the latter requires that intention and action are connected in the right way (cf. Stanley 2011b, 178; Stanley 2011c, 218). But the problem is that epistemic luck in the acquisition of know-how does not interfere with the aptness of this connection at all. Stanley's reply threatens to conflate epistemic luck with the importantly distinct phenomenon of *practical* luck discussed in §5.4.

In a second reply to this problem, Stanley compares ascriptions of know-how with other ascriptions of what is often abbreviated as knowledge-*wh*, which I shall discuss at length in chapter 8, e.g. knowledge where Russia is or knowledge who came to the party. Stanley argues that "in the case of knowledge-*wh*, Gettier intuitions are less robust than in the case of ascriptions of explicit knowledge-that, if present at all." (Stanley 2011c, 219)

As I will discuss in §8.1, Stanley argues extensively that know-how should also be treated as an example of knowledge-*wh*, and that all knowledge-*wh* is propositional knowledge. On this basis, he suggests that this

whole category of propositional knowledge is *generally* less susceptible to epistemic luck than other categories, despite the fact that it remains genuine propositional knowledge. And he supports this with appeal to pragmatic considerations concerning what is normally at stake in the relevant context (cf. Stanley 2011b, 178–181; Stanley 2011c, 219–220).

This reply is entirely in line with Stanley's work on propositional knowledge. He prominently defends a view known as 'pragmatic encroachment', the idea that practical stakes and interests legitimately influence the question which true beliefs amount to knowledge (cf. Stanley 2005b). But it remains unclear why *exactly* these features should be more relevant when it comes to propositional knowledge ascribed with the aid of *wh*-complements. Maybe such knowledge can also resist epistemic luck, at least in certain cases (cf. Carter & Pritchard 2013, 9–10). Thus, without a convincing account of why this should be the case, this reply seems simply *ad hoc*.

Furthermore, *ceteris paribus*, it constitutes a substantial cost for the position of propositionalist intellectualism if, in the light of the problem of epistemic luck, it can only be maintained if a controversial further commitment is made – an endorsement of pragmatic encroachment. For on the face of it, there are strong objections to this view (cf. e.g. Blome-Tillmann 2009; Roeber 2016) and the general debate about these and many related questions is far from settled (cf. Rysiew 2011; Ichikawa & Steup 2012).

Thus, the problem of epistemic luck continues to be a substantial problem for propositionalist intellectualism. On these grounds, several philosophers, most prominently Yuri Cath, have recently endorsed *revisionary* versions of intellectualism. These views have it that know-how is a state of believing or otherwise relying on a true proposition which is nevertheless compatible with epistemic luck (cf. Cath 2011; Cath 2014).⁶ I cannot discuss these proposals in detail in this book. But in any case, I contend that they suffer from the same problems as propositionalist intellectualism with respect to the explanation of intelligent practice. That is, the arguments presented in chapter 9 also apply to revisionary versions of intellectualism.

§ 6.4 Accounting for Epistemic Luck

I have argued that know-how has an epistemic feature which distinguishes it from propositional knowledge: it cannot be undermined by environmental epistemic luck. In this section, I shall argue that Rylean responsibilism

⁶ Similar revisionary versions of propositionalist intellectualism can be found in Brogaard (2011), Capone (2011) and Zardini (2013).

directly predicts this result. Then, however, I shall turn to a tension which results from this view. If know-how crucially entails propositional knowledge (cf. § 2.5 and § 4.3), then it seems to be incoherent to say that the former, but not the latter, is immune to epistemic luck. However, I shall propose a way to solve this problem.

To begin with, Rylean responsibilism is very well positioned to explain why it is that know-how cannot be undermined by environmental epistemic luck. As I have already foreshadowed in footnote 1 on page 108, the notion of understanding has the same feature. One may possess genuine understanding of something despite the fact that this understanding has been acquired luckily. Given that know-how builds on such an understanding of the activity in question, this feature of understanding directly translates into a feature of know-how.⁷

To clarify this argument, consider the following much-discussed case:

Comanche Dominance (Kvanvig 2003, 197–198)

Consider, say, someone's historical understanding of the Comanche dominance of the southern plains of North America from the late seventeenth until the late nineteenth centuries. Suppose that if you asked this person any question about this matter, she would answer correctly. Assume further that the person is answering from stored information; she is not guessing or making up answers, but is honestly averring what she confidently believes the truth to be. Such an ability is surely constitutive of understanding [...]. But does she have knowledge? Ordinarily, yes; but it is not required. For, on the usual theories of knowledge, all those answers could be given from information possessed and still fail to be known to be true, because the answers might only be accidentally true. For example, most history books might have been mistaken, with only the correct ones being the sources of the understanding in question and with no basis in the subject for preferring the sources consulted over those ignored. Such a case fits the model of a standard type of case found in the Gettier literature (in particular, the fake barn case), where such accidentally true beliefs are not justified in the way needed for the beliefs to count as knowledge.

What Kvanvig here references as the 'fake barn case' is indeed a classic example, one which involves what Duncan Pritchard has categorized as 'environmental' luck – that is, as being luckily right in an epistemically unfriendly environment. Pritchard presents this case as follows.⁸

⁷ Clearly, the very same argument is also available to the objectualist intellectualism defended by John Bengson and Mark Moffett. This position is also immune to the problem of epistemic luck since it identifies know-how with states of objectual knowledge or objectual understanding (cf. Bengson & Moffett 2007; 2011c).

⁸ This case was first put forward by Alvin Goldman (1976), who credits Carl Ginet.

Barney (Pritchard 2012, 251)

Using his reliable perceptual faculties, Barney noninferentially forms a true belief that the object in front of him is a barn. Barney is indeed looking at a barn. Unbeknownst to Barney, however, he is in an epistemically unfriendly environment when it comes to making observations of this sort, since most objects that look like barns in these parts are in fact barn façades.

Kvanvig elaborates on *Comanche Dominance* as follows:

The basic idea here is that although knowledge is incompatible with a certain kind of epistemic luck, understanding is not. Upon learning of the disturbed etiology of beliefs about the Comanches, as in the case imagined here, we might say that a person has true beliefs or even justified true beliefs, but no knowledge, if we have heeded our lessons from Gettier. We would not, at least we should not, say that because of these factors, she is lucky to have the knowledge that she has, for knowledge rules out this kind of luck. But we needn't say the same about the claim of understanding. If the etiology were as imagined, one would be lucky to have any understanding at all of the Comanche dominance of the southern plains. So such understanding would count as understanding not undermined by the kind of luck in question. (Kvanvig 2003, 199)

I agree with Kvanvig's assessment⁹ and I think that this case about the understanding of the Comanche's dominance at a certain point in time can easily be turned into a case where the object understood is a normative activity. For example, the person reading about the Comanche could perfectly well also have learned a lot about the Comanche's leathercraft. And if she is dexterous and quick-minded enough, this may already be sufficient for her to acquire the competence to make leather goods in such ways herself. The fact that the source from which she acquired this understanding was terribly unreliable would not count against her epistemic state's being a genuine understanding of Comanche leathercraft or against her ability's being a genuine competence.

Cath explicitly acknowledges that *Comanche Dominance* was his model for creating *The Lucky Light Bulb* on page 189 (cf. Cath 2011, 115 fn. 3). Therefore, it should not be surprising to see that the view that Charlie knows how to change a light bulb goes hand in hand with the view that Charlie understands the activity of fixing light bulbs, that he understands what it takes to fix a light bulb. Given Charlie's reliable ability, the persistence of

⁹ There are many subtleties involving *Comanche Dominance* on which I cannot comment here (cf. e.g. Brogaard 2006; Grimm 2006; Pritchard 2010). But I think that it emerges very clearly that if Kvanvig is correct, as I hold, then he is correct only about the objectual understanding of the Comanche's dominance and not, or at least not necessarily, about some cases of *propositional* understanding *why* certain things have happened. Fortunately, it is only objectual understanding that is relevant here.

his guiding understanding explains the persistence of his know-how. The same is also true of the other cases, *Flight Simulator*, *Nolan's Curve* and *Olavi's Finnish Tango*.

In sum, the account presented in Part One of this book can explain the phenomenon that know-how is immune to being undermined by luck in its acquisition.

At this point, one may object that, while compatibility with epistemic luck is indeed an important feature of know-how, the specifics of my Rylean responsibilist account of know-how create some further trouble. After all, I have argued that possessing an understanding of an activity goes hand in hand with possessing propositional knowledge (cf. § 2.5 and § 4.3). If propositional knowledge is *incompatible* with epistemic luck, while understanding *is* compatible with it, how is my account coherent? I hold that this problem can be solved, but that these questions are also interesting from a more neutral perspective. And as far as I can see, I will be able to point out an example of a neglected sub-category of epistemically unproblematic luck.

To begin with, I have pointed out a way in which a competence is generally closely associated with propositional knowledge, and a way in which propositional knowledge is strictly necessary for know-how. On the one hand, § 2.5 has shown that true beliefs play a crucial role in the acquisition of know-how, and that they are often, indeed paradigmatically, cases of full-blown knowledge, namely knowledge by testimony from experts. Clearly, one can easily construe cases of epistemic luck in which no knowledge, but only true beliefs are stepping-stones in acquiring the competence in question. As far as I can see, nothing essential depends on this difference in the context of acquiring know-how.

On the other hand, § 4.3 has shown that one may *only* possess an understanding of an activity if one *also* possesses propositional knowledge in the form of correct assessments of certain individual acts and of circumstances. These constitute propositional knowledge as opposed to true beliefs since they stem from the understanding of that activity – they are the results of the non-accidentally successful exercise of this assessment competence.

But is it possible that somebody possesses a genuine understanding of an activity which has been acquired despite the presence of epistemic luck, but that she *then* acquires genuine propositional knowledge on the basis of this luckily acquired understanding? How can it be unproblematic to have epistemic luck in the background of propositional knowledge merely because there is a middle man involved which is immune to epistemic luck?

To answer this question, I shall rely on the classification of kinds of luck provided in the seminal work of Duncan Pritchard (2005). For on the

face of it, these cases fit the description of one of the forms of luck which he identifies as entirely unproblematic and compatible with propositional knowledge, namely “capacity epistemic luck” which he defines as a case where “[i]t is lucky that the agent is capable of knowledge.” (Pritchard 2005, 134) He goes on to add that this includes “the luck that might be involved in the agent having the cognitive capacities that enable her to know (the luck regarding her ‘abilities’)” (Pritchard 2005, 134). Clearly, the case just spelled out fits this description perfectly. The competence in question was acquired luckily – under the influence environmental epistemic luck – which, however, does not undermine this as a genuine competence. Then, the exercise of this luckily acquired competence makes the acquisition of propositional knowledge possible. This latter acquisition of propositional knowledge is lucky in the *mediated* sense that it results from the *non-lucky* exercise of a luckily acquired capacity.

To be fair however, I should note that all of Pritchard’s examples in which an agent luckily has certain abilities and nevertheless gains propositional knowledge by exercising them are cases in which she luckily *retains* these abilities in the face of a very high probability of having lost them a moment earlier (cf. Pritchard 2005, 135–136). By contrast, the problem at hand is one where somebody luckily *acquires* these competences in the first place. They therefore constitute a specific sub-category of Pritchard’s ‘capacity epistemic luck’. This sub-category can be illustrated with what I take to be an independently plausible case:

Falcon

Tara learns about falcons from her teacher. She thereby acquires an understanding of what they look like, how they behave, and so on. This is how Tara comes to possess the cognitive capacity to identify falcons on sight. On her way home from class, she indeed comes across a falcon which she correctly recognizes as such. Intuitively, Tara gains the propositional knowledge that this is a falcon. But suppose that Tara was very lucky to have acquired her capacity to identify falcons. Maybe the teacher would normally have lied or misled her but abstained from doing so only on this single day. Maybe her school could easily have been destroyed by a natural disaster and was only saved by incredible luck, thus allowing Tara to go to class in the first place. In such a scenario, it may be lucky for Tara to have acquired her capacity to identify falcons, but the *exercise* of this capacity nevertheless allows her to gain genuine propositional knowledge.

Evidently, if this is the right way to think about these examples, the same will be true in the case of conceptual capacities other than bird classification, including the capacities to assess performances of an activity which are identical with states of understanding of these activities.

I conclude that I can consistently maintain that competences are immune to certain specific forms of epistemic luck in their acquisition and that the propositional knowledge which is gained on the grounds of the understanding involved in these competences is not threatened by this fact.¹⁰

§ 6.5 Procedural Knowledge

The concept of know-how, of competence or skill, is not only an important philosophical concept. It also plays a prominent role in cognitive science, broadly construed to include psychology, ethology, and relevant parts of linguistics and neuroscience. In this and the following sections, I will offer a brief discussion of these complex problems. I shall begin by discussing how the account of know-how I defend bears on cognate concepts in cognitive science. Then, in § 6.6 and § 6.7, I will turn to the question of the scientific evidence concerning specific cases of abilities, broadening my examples beyond the realm of normal human adults.

Cognitive scientists employ a distinction which is akin to the distinction between know-how and propositional knowledge, namely between *procedural* and *declarative* knowledge or memory. This distinction originally stems from computer science where it was used as a distinction between two ways of implementing one and the same state rather than between two kinds of state.¹¹ As Zoltan Dienes and Josef Perner write in a seminal article:

The procedural-declarative knowledge distinction was introduced in artificial intelligence [...] and later taken over in psychological modelling [...]. It concerned how best to implement knowledge: Should one represent the knowledge that all men are mortal as a general declaration ‘for every individual it is true that if that individual is human it is also mortal’? [...] The alternative would be to have a specialised inference procedure: ‘Whenever an individual is introduced that is human, represent that that individual is mortal.’ (Dienes & Perner 1999, 743)

These concepts have now acquired a psychological sense in which they do purport to distinguish kinds of states of knowledge or memory. A standard definition of these concepts, as they are used in cognitive science, is this:

¹⁰ To clarify, this should not be taken to entail that know-how can be acquired in *any* lucky way. In particular, it does not follow from my account that creatures like Millikan’s and Davidson’s famous ‘Swampman’ possess know-how. For if “by some cosmic accident a collection of molecules formerly in random motion were to coalesce to form your exact physical double”, then “that being would have no ideas, no beliefs, no intentions [...] because the evolutionary history of the being would be wrong.” (Millikan 1984, 93). Since such a creature does not possess any mental states (cf. Davidson 1987, 443–444), it cannot understand or learn anything either. At most, it would have mere dispositions. For discussion of such cases, see also Fridland (2010, 139–140).

¹¹ For discussion of this point with respect to know-how, see Stanley (2011b, 150–157).

Declarative memory encompasses the acquisition, retention, and retrieval of knowledge that can be consciously and intentionally recollected[.] [...] In contrast, non-declarative or procedural kinds of memory encompass the acquisition, retention, and retrieval of knowledge expressed through experience-induced changes in performance. (Gabrieli 1998, 89–90)

Often, it is explicitly asserted that this distinction coincides with the distinction between know-how and propositional knowledge, for example in the title of a well-known study of amnesia, on which I shall comment later (cf. Cohen & Squire 1980). An authoritative encyclopedia states that, in skill learning, “one is said to have acquired ‘procedural knowledge’ (knowing how to carry out some procedure), as opposed to ‘declarative knowledge’ (knowing that some proposition is correct).” (Reisberg 1999, 461)

By contrast, the authors of a prominent introduction to cognitive science are more cautious:

Traditional epistemology distinguishes between knowing how and knowing that. Though this distinction is not the same as the one psychologists draw between procedural and declarative knowledge, the two are closely related. [...] However, these distinctions do not coincide exactly. (Stillings *et al.* 1995, 369)

While I cannot provide a detailed metatheoretic analysis of the notions of procedural and declarative knowledge in the present context, I would nevertheless like to mention where this distinction fails to coincide with the distinction between know-how and propositional knowledge. This can be seen most clearly in comparison with a prominent view by Michael Devitt. He presents a *tour de force* through the scientific literature and concludes:

[T]he psychological distinction between declarative and procedural knowledge, and the related distinctions about memory and learning, are well established in empirical science. And all of this is evidence for the nature of, near enough, the folk distinction between knowledge-that and knowledge-how. I say “near enough” because I think that the psychologists would have done better to identify their procedural knowledge with one common kind of the folk’s knowledge-how, *mere* knowledge-how. [...] [I]t is that Rylean kind of knowledge-how that is thought not to involve knowledge-that. Still, we need not fuss about this: mere knowledge-how is still knowledge-how. So, if the psychologists are right and procedural knowledge does not involve declarative knowledge, then declarative knowledge is not essential to knowledge-how. (Devitt 2011, 212)

I shall assume for the sake of argument that the psychological notion of declarative knowledge is the same as the philosophical notion of propositional knowledge, even if there are a number of difficulties with this assumption. Most importantly, declarative knowledge is typically characterized as

fully verbalizable and as explicit rather than implicit (cf. e.g. Devitt 2011, 209).¹² Clearly, though, this is not true of propositional knowledge as it is discussed in philosophy. It is probably true that all declarative knowledge is propositional knowledge. But it is most certainly false that all propositional knowledge is declarative knowledge, i.e. fully verbalizable and fully explicit.

But bracket this. On the assumption that declarative knowledge is propositional knowledge, Devitt's argument seems to be this. Know-how and procedural knowledge are the same thing because both are concerned with the agent's performances and procedures, and just like no declarative knowledge is necessary for procedural knowledge, no propositional knowledge is necessary for know-how, or at least for the know-how Devitt calls 'Rylean'.

Unsurprisingly, Devitt's *modus ponens* is my *modus tollens*. True, know-how concerns performances and procedures. But I have already explained why know-how requires propositional knowledge in §2.5 and §4.3. This entails that know-how cannot be identical with any concept which is independent from propositional knowledge, including procedural knowledge. In fact, as Devitt's symptomatic distinction between 'knowledge-how' and 'mere knowledge-how' already foreshadows, I contend that the notion of procedural knowledge turns out to be much closer to and maybe even identical with the concept of *ability*.

This point can also be appreciated from another angle. In §1.5, I distinguished between mere abilities, i.e. abilities which do not involve an understanding and propositional knowledge about the activity in question, and abilities which *do* amount to know how, partly because they *do* involve these elements. Analogously, cognitive scientists agree that at least some kinds of procedural knowledge require declarative knowledge.

In sum, my hypothesis is that procedural knowledge should not be equated with know-how, but instead, roughly but near enough, with ability.

While I have not been able to make a more comprehensive case for this view, I would nevertheless like to mention that it makes perfect sense of two crucial facts – the fact that procedural knowledge should indeed play a role in the philosophical debate about know-how, and the fact that one cannot simply read off a philosophical account of know-how from the psychological

¹² Further, the distinction between explicit and implicit knowledge can also be applied within the realm of declarative knowledge (cf. e.g. Dienes & Perner 1999). And the question of verbalizability is complicated by the fact that declarative knowledge is further subdivided into episodic and semantic knowledge (cf. e.g. Tulving 1999). While semantic memory is typically assumed to be verbalizable descriptively, there is no such requirement on episodic memory. And as Jason Stanley has convincingly shown, if demonstrative verbalization is allowed, then episodic memory can be verbalized just as well as semantic memory (cf. Stanley 2011b, 157–163).

notion of procedural knowledge. On the one hand, since know-how is a form of ability, and since abilities are studied under the heading of ‘procedural knowledge’, there is much to be gained from bringing these debates together, and I have already referred to the scientific literature myself at various points. On the other hand, if the only notion on offer is that of ability or procedural knowledge, then the distinction between mere ability and genuine know-how is simply out of the picture. Then, there is no way to account for the crucial phenomenon of normative guidance as opposed to mere conformity to norms (cf. § 1.5), and corollary phenomena such as the problem of practical luck (cf. § 5.4), the problem of epistemic luck (cf. § 6.3 and § 6.4), and the opacity of know-how ascriptions (cf. § 6.1 and § 6.2).

Thus, where philosophers merely rely on the distinction between procedural knowledge and declarative knowledge, many crucial questions with respect to know-how are left open, and the distinction between mere ability and full-blown competence is in danger. In what follows, I shall try to substantiate this conclusion with pertinent case studies.

§ 6.6 The Cases of D.F. and H.M.

In the debate about know-how, several examples from psychological studies have been put forward and discussed as cases of know-how. The most prominent of these are a woman abbreviated D.F. and a man abbreviated H.M. In this section, I shall discuss these cases.

The first example, D.F., is a tragic victim of carbon-monoxide poisoning, a woman with visual apperceptive agnosia who has participated in a number of much-discussed studies. The core result was this:

D.F. (Goodale & Milner 1992, 22)

Despite her profound inability to recognize the size, shape and orientation of visual objects, D.F. showed strikingly accurate guidance of hand and finger movements directed at the very same objects. [...] Thus, when presented with a large slot that could be placed in one of a number of different orientations, she showed great difficulty in indicating the orientation either verbally or manually (i.e. by rotating her hand or a hand-held card). Nevertheless, she was as good as normal subjects at reaching out and placing her hand in the slot, turning her hand appropriately from the very onset of the movement.

While this case has been prominently discussed by a number of philosophers (cf. e.g. Kelly 2002), it has also been used in order to object to the view that know-how is propositional knowledge, i.e. as a counterexample to propositionalist intellectualism. According to these critics, D.F. does not

have the relevant propositional knowledge because, unable to recognize the orientation of the slot, she lacks the required concept of the way in which to post the card which is necessary for the relevant propositional knowledge (cf. Toribio 2008, 43–44; Jung & Newen 2010, 127–128).

While I am not committed to propositionalist intellectualism myself, I *am* committed to the view that know-how requires propositional knowledge. Thus, it seems, D.F. also constitutes a challenge to the Rylean responsibilist account defended here. But, as I shall show now, it is not clear that the case of D.F. really poses a problem for my view.

On the one hand, it may well turn out that D.F. only possesses a mere ability to post a card through a slot rather than a full-blown competence. Since such performances are anyway good candidates for basic actions, which only admit of mere ability as opposed to genuine competence anyway (cf. § 1.5), this is a plausible answer to the case of D.F. But on the other hand, judging from what I gather from the details of her case, it seems more likely that she possesses know-how after all. D.F. seems to have an understanding of the activity of posting cards through slots, and she also seems to be guided by this understanding. If so, she knows how to do so.

True, there is a difference between D.F. and other people. Normally, people can see the orientation of the slot beforehand and consciously guide themselves in posting the card. In contrast, D.F. can only move her hand appropriately without consciously pre-planned guidance. Further, while D.F. cannot explicitly describe the orientation of the slot, this is normally possible. But as I have discussed at length in Part One of this book, especially in chapter 3, this is compatible with the exercise of genuine know-how.

Given the account of guidance as responsible control I presented in § 4.4, D.F.'s ability can be shown to amount to a genuine competence as long as she can assess her performances on the fly and judge when she does well and when she does badly at posting a card through a slot, influencing herself to succeed in the end. And in fact, it looks like D.F. does possess this capacity. At least, this is strongly suggested by Sean Kelly's observation that "if you ask her to report the orientation of the slot, she'll begin to move her hand toward the slot as if she were going to push it through, and then at the last moment she'll stop, saying '*This* is the orientation it's in'" (Kelly 2002, 388). As Jason Stanley remarks, this suggests that D.F. possesses a demonstrative conception of what to do in order to post a card through a slot (cf. Stanley 2011b, 172 fn. 22).

In sum, the tragic and intriguing case of D.F. is most plausibly an example of genuine know-how, but an example where the relevant understanding and propositional knowledge remains very large at a demonstrative level.

And at the same time, D.F. may be a borderline case, given the fact that the relevant performances could also be understood merely as basic actions.

This brings me to the second clinical case which has gained some prominence in the debate about know-how. This is the case of the amnesiac H.M. After his neurosurgeon William B. Scoville resected H.M.'s medial temporal lobe (MTL) in 1953, he was found to have acquired anterograde amnesia, sparking a series of studies led by Brenda Milner (cf. Scoville & Milner 1957). Until his death in 2008, H.M. has been the subject of many important studies. As Suzanne Corkin summarizes:

H.M. (Corkin 2002, 153)

H.M.'s anterograde amnesia manifests as deficient acquisition of episodic knowledge (memory for events that have a specific spatial and temporal context) and of semantic knowledge (general knowledge about the world, including new word meanings). The evidence strongly supports the conclusion that the MTL structures that were removed in H.M. are crucial for long-term declarative memory (conscious recollection of facts and events), including the acquisition of new semantic knowledge. Other studies indicate that H.M.'s short-term memory is intact and, therefore, not dependent on these MTL structures. So, he can encode new information, but rarely uses this information to establish a long-term trace. (Corkin 2002, 153)

What is important for the question of know-how, however, is that H.M. showed significant improvements in his practical capacities on the basis of training, despite the fact that he was unable to remember any of this training in the later trials in which he did significantly better. Beginning with Brenda Milner's groundbreaking mirror-drawing experiment, this result has now been firmly established for a number of tasks (cf. Corkin 2002, 154–156). This includes a famous experiment in which amnesiac patients were shown to perform just as well as normal people in mirror-reading of new words, while words already read in the mirror in previous trials were read faster by normal people than by amnesiac patients (cf. Cohen & Squire 1980).

In the debate about know-how, many philosophers have interpreted these results as establishing the independence of know-how from propositional knowledge (cf. Bzdak 2008; Wallis 2008; Adams 2009). Evidently, this argument requires the premise that H.M., and other patients like him, have acquired and retained some relevant know-how but no relevant propositional knowledge. And on the face of it, this is indeed a plausible description of such cases.

To show why my account is nevertheless compatible with this evidence, I shall rely on two distinctions. First, there is the distinction between mere ability and genuine know-how. This allows me to propose that when an

amnesiac patient is shown to retain procedural knowledge, this may nevertheless remain a mere ability (cf. § 6.5). And second, there is a further important distinction between what a patient *acquires* in each trial and what they *retain* over the course of these trials. As I shall suggest, what patients like H.M. *acquire* in each trial is genuine know-how including propositional knowledge. But since they lose the understanding which underlies this competence afterwards, what they *retain* from trial to trial is indeed a mere ability.

To support this view, let me take the second point first. Here, I can rely on a methodological observation by Jason Stanley and the neuroscientist John Krakauer. They write:

Here is a fact about HM. Each time HM performed the task he received *explicit verbal instruction*, and was able to use that knowledge each time. HM of course forgot that he had used explicit knowledge. But that of course does not entail he did not require the knowledge at the time. To understand what the original results do or do not mean, it is useful to consider more recent experiments conducted in patients with similar medial temporal lobe lesions to HM since the 1960s. The general approach in follow-up studies in patients with medial temporal lobe lesions, as in the original Milner experiment, is to demonstrate dissociation between improvement in motor performance variables, usually time to completion and error/accuracy measures, and ability to explicitly recall aspects of the task. What becomes apparent when considering this literature is that the amnesic patients could not perform any of the tasks *unless instruction was provided on each day*. (Stanley & Krakauer 2013, 8)

On this basis, it seems entirely plausible to say that H.M. acquired a genuine competence each time he was explicitly verbally instructed and performed the task in question. Given the testimony of the researchers, he acquired an understanding of what he was supposed to do and guided himself on the basis of this understanding, including the relevant propositional knowledge (cf. § 2.5). H.M., as described on the previous page, *does* possess a functioning short-term memory, and this underwrites the possibility that, for a short period of time, he retains a newly acquired competence.

This brings me back to my first point. As far as I can see, H.M. should not be described as *retaining* these competences, but rather as losing them in between the test trials and as re-acquiring them during the subsequent trials. After all, his explicit memory of everything concerning these tasks is completely lost. Nevertheless, there is *something* which he retains and which underwrites his increase in success over time. But since H.M. relies on explicit verbal instruction in order to perform the relevant tasks on subsequent trials, I do not find it plausible to say that he retains an un-

derstanding of this task, and hence that he retains a piece of know-how. Instead, he seems to retain ‘procedural knowledge’ in the form of mere automatic abilities. And on the basis of these, re-acquiring the relevant competences can then proceed faster, which leaves time to learn more *now* than was learned earlier.

This conclusion also sheds light on a further argument by Michael Adams. Concerning the study by Cohen & Squire (1980), he points out:

[W]hen developing a skill such as riding a bike, one improves this skill by repeated, intentional engagement in the activity, but one does not, however, improve one’s digestion by digesting repeatedly. The skills involved in the mirror-reading study mentioned above, however, were developed and improved only because the subjects repeatedly directed their attention toward the task at hand. (Adams 2009, 111)

However, Adams is too quick to draw the conclusion is that these subjects must have retained genuine know-how as opposed to mere ability (cf. Adams 2009, 110–111). Given the fact that they lose all explicit memory of previous trials, it is much more plausible to describe them as losing their know-how, since they lose the crucial element of conceptual understanding, and as retaining only the improved mere automatic abilities. Again, it would be odd to describe these subjects as continuing to practice the respective competences over the whole course of trials precisely because of their loss of memory.

However, part of Adams’ general outlook sits very well with the position I have established earlier. With reference to a further study by Paul Reber, Barbara Knowlton, and Larry Squire (cf. Reber *et al.* 1996), he writes:

In this study, the amnesic subjects acquired the new skill at a rate equal to that of the control group. When asked to abstract principles away from the skill and apply it to a new task, however, the amnesic group was severely limited. This is because they were unable to have procedural and declarative knowledge combine together to apply the skill knowledge they had acquired to new and different contexts. This study supports the traditional characterization of knowledge-how as rigid in application; knowledge-that, however, is characterized by its plasticity. For new skills to have such plasticity, there must be an interplay between knowledge-that and knowledge-how. (Adams 2009, 111)

In this passage, Adams equates procedural knowledge with know-how and declarative knowledge with propositional knowledge-that, arguing that these must combine in the right way in order to form a competence. Given that I have argued that ‘procedural knowledge’ is best understood as, roughly, ability (cf. § 6.5), and given that I have been using ‘know-how’ synonymously

with ‘competence’ or ‘skill’, I can entirely agree with Adams’ point in this passage. Terminological questions aside, what Adams calls the “interplay” between propositional knowledge and ability is precisely what characterizes genuine competence or know-how.

§ 6.7 Machines, Animals, and Infants

In this chapter, I have discussed various aspects of the cognitive nature of know-how with respect to specific examples and puzzle cases. In this final section, I shall finally turn to three large groups of important examples – machines, non-human animals, and human infants – all of which have already been present in the debate since Ryle, for instance in the examples which I labelled *Clocks & Seals* on page 32 and *Animals & Infants* on page 43. I will only be able to offer a brief and cursory discussion of these cases, mainly pointing out which commitments follow from the Rylean responsibilist account defended in Part One. However, I shall close with some general methodological considerations, and with some thoughts on the scope and the ambitions of a philosophical account of know-how.

The first group starts with an extreme, but very clear kind of example – simple machines. What Ryle already made clear with respect to clocks is nowadays often discussed with respect to pocket calculators. Given the standard definition of this term (cf. § 6.5), such a machine has procedural knowledge of calculating operations. One may probably also say that it is able to calculate. But, *pace* Lihoreau (2008), it certainly does not *know how* to do so (cf. Stanley 2011b, 153 fn. 2; Bengson & Moffett 2011b, 53 fn. 68). Clearly, a pocket calculator fails to have an understanding of calculating, and it fails to be guided by such an understanding in calculating. Crucially, this is true about a calculating machine even if it implements, say, a mechanism of double-checking its results. At the very least, such a mechanism does not result in the kind of learning discussed in chapter 2, and it does not exert the kind of responsible control which I have identified as crucial for know-how in chapter 4. Such a machine has remarkable procedural knowledge – that is, remarkable abilities – but *mere* abilities.

Naturally, these questions lead to the problem of artificial intelligence. Suppose that a machine – a robot, say – can indeed be said to take up information from experience and alter its behavior on this basis, very much in the spirit of the account of guidance as responsible control I offered in § 4.4. As far as I can see, it will still remain an open question if we are able to understand this uptake of information as genuine conceptual understanding

and knowledge. If this is indeed the case, and only if it is, we will have genuine know-how. But whether or not this is the case is a complicated question which is simply beyond the scope of this book.¹³

The second group of cases consists of various examples involving non-human animals such as Ryle's circus seals and dogs in the examples referenced above. In these passages, Ryle seems to be very dismissive of non-human animals, treating them practically on a par with clocks. In contrast, many other writers have called upon the results of ethology, where the capacities of animals are uniformly stated in terms of procedural knowledge, and argued that non-human animals possess know-how – for example, that some dogs know how to catch a frisbee, that bees know how to communicate the distance, direction and quality of food sources by dancing, or that certain chimpanzees know how to extract termites from their nests using simple tools. And it has also been suggested that these non-human animals possess know-how *without* possessing concepts or propositional knowledge (cf. e.g. Noë 2005; Johnson 2006; Jung & Newen 2010; Devitt 2011).

What to make of this? To begin with, there is a danger that these considerations are motivated largely by the notion of procedural knowledge which, as discussed in § 6.5, corresponds more closely to the concept of ability than to genuine know-how. But we can bracket this here. For I am happy to embrace the consequence of the Rylean responsibilist account of know-how and reject the descriptions of these creatures as having know-how without propositional knowledge. Possessing know-how, possessing concepts and possessing propositional knowledge is only possible in concert.

This entails that all candidate creatures must consistently be understood *either* as possessing mere abilities and information states which fail to be conceptual, *or* as possessing genuine know-how, as well as concepts and propositional knowledge. As far as my intuitions go, there are clear cases in each of these categories – the former including animals like earthworms, and the latter including humans, chimpanzees (cf. e.g. Matsuzawa 2002), maybe ravens (cf. e.g. Heinrich 2002) and dolphins (cf. e.g. Hermann 2002). But these are mere intuitions. As with the problem of artificial intelligence, drawing this boundary is beyond the scope of this book, and doing so would require a deeper engagement with the debate about non-human animals.¹⁴

The third group of cases comes into view by turning from *phylogeny* to *ontogeny*. This is exemplified in the passage I quoted from Ryle as *Animals & Infants* on page 43. Just like there is an *evolutionary* continuity

¹³ For this debate, see Boden (1990) and Frankish & Ramsey (2014).

¹⁴ For this debate, see Bekoff *et al.* (2002), Allen & Trestman (2014), Andrews (2012), and Beck (2012).

between non-human and human animals, there is also a *developmental* continuity between the human newborn and the human adult. And just like I am committed to the view that, at some point in the evolutionary history of humans, there was a joint transition from mere ability to full-blown know-how *as well as* concepts and propositional knowledge, I am equally committed to an analogous transition in the cognitive development of every individual human organism. These transitions may well happen gradually, and there may well be points in time which can reasonably be described as borderline. But know-how and conceptuality go hand in hand.

This is not to deny that very sophisticated abilities have to be in place *before* the advent of propositional thought. The only thing I am committed to is that these remarkable abilities remain *mere* abilities up until the point where concepts, understanding, and normative guidance come into play. Thus, we do not have to rewrite the widespread narrative that know-how comes first and propositional knowledge is only developed later. It only needs to be relabelled. What comes first is ability.

Again, I have to close this topic with the remark that anything approaching a substantial discussion of the development of human infants is beyond the scope of this book – both with respect to scientific issues and with respect to the philosophical questions surrounding them.¹⁵ My only aim was to point out what the Rylean responsibilist account of know-how entails and what it does not entail.

This concludes my considerations about machines, non-human animals and human infants. Clearly, what I have been able to say is not much. In essence, I simply applied the commitment of the interdependence of know-how, propositional knowledge and conceptual understanding to these groups of examples, and pointed out some intuitive ideas and open questions, all with the background knowledge that this is one way in which the concept of know-how is vague and context-dependent (cf. § 1.7). However, this is in keeping with the mainstream philosophical approach in epistemology. The core of the discussion, I take it, should focus on paradigm cases of human adults. A canonical statement of this methodology is due to Keith Lehrer:

[W]e shall not be concerned with the sort of knowledge attributed to animals, small children, and simple machines that store information, such as telephones that store telephone numbers. Such animals, children, or machines may possess information and even communicate it to others, but they do not know that the information they possess is correct. They lack any conception of the distinction between veracity and correct information, on the one hand, and deception and misinformation, on the other. Any child, animal, or machine that not only possesses information but

¹⁵ For this debate, see Matthews & Mullin (2015).

knows whether the information is correct is, of course, a candidate for being a knowing subject. In those cases in which such knowledge is lacking, however, we shall assume ignorance in the information sense of knowledge under investigation here. (Lehrer 1990, 7–8)

Thus, my point in this section can also be summarized as the insistence that the debate about know-how should be treated just like the debate about propositional knowledge. Just like the debate about propositional knowledge should start with cases of human adults and only consider other beings later, the debate about know-how should also not be started with reference to anything but the paradigm cases of intelligent practice in human adults – examples of which have been present in the debate from the very beginning, in *Ryle's Range of Cases* as presented on page 14. At the very least, this is the methodology exemplified in this book.

One may also conclude that there is a hierarchy of *grades* of knowledge, with full-blown knowledge at the top and lower grades of knowledge beneath. With respect to propositional knowledge, this has been proposed by Ernest Sosa, who distinguishes between “animal knowledge” and “reflective knowledge”, where reflective knowledge is more demanding and, very roughly, not only includes reliably true belief, but a reflective understanding of one’s epistemic credentials (cf. e.g. Sosa 2007a). In this vein, Jeremy Fantl has proposed to introduce an analogous hierarchy of grades of know-how, where the lowest grades start with mere possibility, leading to what I called mere ability, and eventually reaching full-blown know-how as I have discussed it here (cf. Fantl 2008, 465–466).

Such proposals are perfectly compatible with the position I have defended here. And they would even open up the possibility to be more concessive towards certain views I rejected. For example, one may conclude that anti-intellectualism is entirely correct, but only with respect to ‘animal know-how’ rather than with respect to ‘reflective know-how’. Likewise, the appeal to the concept of procedural knowledge as a case of know-how may also be entirely correct, but correct only when it comes to ‘animal know-how’, not in the case of ‘reflective know-how’. In the end, however, this is mainly a question of terminology. While I have sympathies for the idea of grades of know-how, I have opted for a clear change from one term to another to draw the distinctions in question. It does not matter *how* these are expressed, as long as they *are* expressed.

Chapter 7

On ‘Knows how to’

In Part One of this book, I have argued for a philosophical explication of the concept of know-how. I claimed that know-how is a skill or a competence, a reliable ability to do well in an activity which is guided by an understanding of what it takes to do so. To state the obvious, I have employed the English language in order to formulate this view. Where I used the expression ‘knows how to’, I used it to express the concept of know-how so explicated. But *does* this locution express that concept? In this chapter, and in chapter 8, I finally return to this problem.

I will begin, in §7.1, by making an intuitive case for the view that ‘knows how to’ can indeed express the concept of competence, while noting the *prima facie* problem that the use of ‘knows how to’ is often entirely acceptable even in cases where no corresponding ability is present. However, I will propose two independently sufficient strategies for explaining this fact. First, §7.2 appeals to the question which alternatives to ‘knows how to’ are pragmatically salient in the relevant context. Second, §7.3 presents an argument for the view that ‘knows how to’ can express both the concept of genuine know-how and the concept of a mere understanding of an activity.

This view will then be defended and elaborated in the final pair of sections. §7.4 argues that ‘knows how to’ may very well be an ambiguous expression, or that there is at least no conclusive evidence that it is not. Finally, §7.5 will suggest other ways of understanding the bifurcation hypothesis and argue that the most plausible proposal consists in treating it as polysemous, especially given the many interconnections between this view and the earlier considerations over the course of this chapter.

§ 7.1 ‘Knows how to’ and Ability

As I have already made clear, what I have laid out so far commits me to the claim that the English expression ‘knows how to’, as I have used it, actually expresses the concept of know-how or competence as I have explicated it. In principle, this is compatible with the view that this expression actually has a very different meaning, and that I used it in a merely technical way. But other things being equal, it would be preferable to have a philosophical account of know-how which is not *merely* compatible with the syntax and semantics of ‘knows how to’, but which fits with its meaning as well as possible. And this is what I shall go on to explain and justify.

To begin with, I contend that it is the default and common-sense position that ‘knows how to’ at least sometimes expresses the concept of competence, know-how or skill. Therefore, the burden of proof should clearly lie with those philosophers who argue that this is false. And indeed, I will mainly discuss and object to those who have presented arguments for this revisionary conclusion, particularly to the arguments by Jason Stanley, John Bengtson, and Marc Moffett.

However, there are also positive reasons in favor of the common sense position I defend. The strongest argument I can offer will be presented in § 7.3. But one may even make a simple pre-theoretic case for this view.

This consideration starts with a look into a typical authoritative dictionary and thesaurus of the English language like Merriam-Webster Online. One among several explicitly defined meanings of the verb ‘know’ listed there is “to have a practical understanding of” (2014a; 2014b), which is illustrated with the example “knows how to write” (2014a). In the same way, one among several explicitly defined meanings of the noun ‘knowledge’ is stated as “acquaintance with or understanding of a science, art, or technique” (2014a, cf. 2014b). Furthermore, the definition of the noun ‘know-how’ is formulated as “knowledge gained by actually doing or living through something” (2014a) or, alternatively, as “knowledge of how to do something well” and “knowledge of how to do something smoothly and efficiently” (2014b), which is illustrated with the example “you’ll gain some practical know-how in this auto mechanics class” (2014a). Finally, and most clearly, the entry for the noun ‘know-how’ explicitly lists *synonyms*, including the terms ‘expertise’, ‘proficiency’ and ‘skill’ (2014a; 2014b).

On this basis, it seems to be very clear that the Rylean concept of know-how as competence can indeed be expressed in English with the expression ‘knows how to’. In fact, it seems that the existence of the distinct noun ‘know-how’ can be explained very well given the hypothesis that the verb

'know' and the noun 'knowledge' express a *broader* category of concepts while *one* of these can be expressed more precisely with the specification of the verb 'knows' as 'knows how to' or with the further and more specific noun 'know-how' instead of 'knowledge'. The fact that there even *is* a further noun 'know-how' over and above the noun 'knowledge' suggests that 'know-how' specifies a *kind* of knowledge otherwise left unspecified. Thus, 'knows' and 'knowledge' can express a number of concepts, including know-how and propositional knowledge.¹

Thus, there is good *prima facie* evidence for the common-sense position that 'knows how to' can indeed express the concept of competence.

However, as I have argued in § 1.3 and defended at length in § 5.1 and § 5.2, the ability to do something is necessary for the knowledge how to do it. Competences are a kind of ability. But this creates a problem. After all, it is at least sometimes entirely acceptable to apply the expression 'knows how to' even in cases where clearly no ability is present. In the remainder of this section, I shall elaborate on this problem.

For example, I maintain that the coach in *Bela Karoli* on page 54 fails to know how to perform a standing layout on beam. However, it would often be entirely unproblematic to describe him as 'knowing how to' do so. After all, what is actually *meant* by this attribution would be very clear, namely that he understands this activity well enough in order to be able to teach it. The same is true of the paradigm cases cited by intellectualists like Stanley & Williamson and Bengson & Moffett as allegedly establishing the independence of know-how from ability. In *Ski Instructor* on page 150, we may unproblematically use the expression 'knows how to', but only because it is clear that this is intended to express an understanding of the activity of skiing, independently of genuine competence to ski. Likewise, the figure skater in *Quintuple Salchow* on page 150 only has an understanding of what it takes to perform a quintuple salchow, rather than the knowledge how to perform one herself. Since this is all a speaker may mean to imply with the expression 'knows how to', it is again unproblematic to do so.

Bengson & Moffett have also tried to establish that what 'knows how to' expresses does not involve ability at all. In § 7.2, I will present these arguments and offer an alternative explanation of their linguistic data.

However, Bengson & Moffett have also argued that there is a certain non-standard class of examples of activities for which 'knows how to' and 'is able to' actually interact in just the way I defended. In these cases, an

¹ It should be noted that this argument only concerns 'knowledge' and 'knows' and remains neutral on the further question whether 'knows how to' expresses a uniform concept. This will be the topic of § 7.3.

attribution of the state ascribed with ‘knows how to’ semantically *entails* an attribution of the state ascribed with ‘is able to’. In the remainder of this section, I discuss this issue and return to the initial problem in § 7.2.

Bengson & Moffett’s paradigm example for the semantic entailment in question is this (Bengson & Moffett 2007, 36):

- (5) (a) Irina knows how to add.
- (b) Irina is able to add.
- (c) * Irina knows how to add, but she is unable to do so.
- (d) * Not only does Irina know how to add, she can actually do so.

Bengson & Moffett claim: “The ability attribution [(5 b)] is neither cancellable [(5 c)] nor reinforceable [(5 d)]” and therefore, (5 c–d) are ungrammatical (and marked as such by an asterisk) (Bengson & Moffett 2007, 36). The explanation they propose has it that these cases involve what they call ‘ability-based concepts’, i.e. concepts which one cannot have without having the relevant ability (cf. Bengson & Moffett 2007, 44–45). That is, knowing how to add requires a concept of addition, which in turn is ability-based, i.e. requires the ability to add.

This raises an interesting question. Why is it that one cannot have these concepts without having the relevant ability? I will come back to the full account of know-how Bengson & Moffett propose later on, in chapters 8 and 9. For now, I would like to note that the Rylean responsibilist proposal developed in Part One of this book provides a very straightforward explanation of this fact, even if it simultaneously denies that cases like these contrast with other cases which fail to require abilities.

As defended in § 4.2, the understanding of *all* activities requires reasonable conceptual mastery. And as defended in § 1.5, such an understanding is crucial for every competence. What makes activities like adding special is not that knowing how to engage in them always requires the actual ability to do so. Instead, these cases are special in that they do not allow for *merely* understanding them or merely knowing *about* them, without any competence to engage in them oneself. In § 2.2, I have already discussed a parallel example by Gilbert Ryle, the activity of philosophizing, and in § 4.7, I have discussed a further parallel case by David Carr, *Multiplication* on page 141, on which Bengson & Moffett also comment in their discussion of their own case. As § 4.7 shows, this peculiarity stems from the self-reflexivity of conceptual capacities. Philosophizing, adding, and multiplying are essentially conceptual competences. Understanding these activities – i.e. having the conceptual capacity to assess performances of them (cf. § 4.1 and § 4.2) – already requires having the competence to engage in them.

Thus, Bengson & Moffett are correct to point out a crucial difference between standard cases and cases like adding. However, this is not about the difference between know-how which requires ability and know-how which fails to do so. Instead, it is the difference between competences which are not essentially conceptual and competences which are.

§ 7.2 Pragmatically Salient Alternatives

I have maintained that ‘knows how to’ can be used to express the concept of competence as defended in this book, i.e. a concept which requires ability. However, Bengson & Moffett have tried to establish that what ‘knows how to’ expresses does not involve ability at all. In this section, I will discuss two lines of argument which these philosophers have used and argue that both fall prey to a common problem. Both are inconclusive because they fail to take into account the question which alternative expressions other than ‘knows how to’ are pragmatically salient in the relevant context.

The upshot of this discussion will be that ‘knows how to’ expresses the concept of competence, including ability, at least sometimes. Maybe ‘knows how to’ always semantically expresses competence, and the phenomena discussed here are indeed *merely* pragmatic. But maybe ‘knows how to’ can semantically express *multiple* concepts. Indeed, this latter view will be discussed and defended in § 7.3.

In their first line of argument, Bengson & Moffett take another look at their example *Quintuple Salchow*, cited on page 150, and at the following sentences which may be used to describe it:

- (6) (a) Irina knows how to do a quintuple salchow.
- (b) Irina is able to do a quintuple salchow.
- (c) Irina knows how to do a quintuple salchow, but she is unable to do one.
- (d) Not only does Irina know how to do a quintuple salchow, she can actually do one.

To determine whether know-how attributions entail ability attributions, we can check for the presence of such characteristics as cancellability or reinforceability, which serve as the standard tests for entailment. Consider, for instance, the know-how attribution [(6 a)]. The ability attribution [(6 b)] is both cancellable [(6 c)] and reinforceable [(6 d)]. [...] If knowing how to do a quintuple salchow [(6 a)] entailed being able to do one [(6 b)], then [(6 c)] would be internally inconsistent and [(6 d)] would be redundant. We, however, find [(6 c)] and [(6 d)] unproblematic. (Bengson & Moffett 2007, 35)

This presentation of the issue makes it easy for me to identify the burden of proof which results from the account I have proposed. According to Bengson & Moffett, the fact that (6c) and (6d) are perfectly adequate English sentences poses a *prima facie* problem for me. Unfortunately, though, these examples fail to show what they purport to show.

First, I doubt that the acceptability of (6c) really entails the cancellability of an attribution of ability *vis à vis* an attribution of know-how. To begin with, the phrase ‘is unable to’ in (6c) is entirely open as to whether it refers to a general ability or a specific ability (cf. § 1.4) – to the general ability to do something under suitable circumstances or to the opportunity to exercise this ability here and now.

In order to disambiguate these two readings of (6c), one could distinguish between ‘is unable to’ and ‘does not have the ability to’, but even this might fail to make the difference clear enough in some cases. Instead, the following statements should be unproblematically clear. The words I have added to the original wording of (6c) are emphasized in order to mark the contrast I introduce myself, and the indices indicate the intended meanings of ‘s’ for specific versus ‘g’ for general ability.

- (6c_s) Irina knows how to do a quintuple salchow, but she is unable to do one *here and now*.
 (6c_g) Irina knows how to do a quintuple salchow, but she is unable to do one *under the right circumstances for performing this stunt*.

I take it that (6c_s) is acceptable and informative. But this fails to show what Bengson & Moffett want (6c) to show, namely that one may know how to do something and fail to be able to do so whatsoever. Instead, (6c_g) supports Bengson & Moffett’s view. But it is clearly much less acceptable than (6c_s). It might indeed be ‘internally inconsistent’, as they put it.

But this does not mean that (6c_g) is not acceptable *at all*. Even semantically inconsistent statements can be *pragmatically* acceptable. Such a statement will often receive a charitable reinterpretation in the relevant context, depending on the pragmatically salient alternatives. Along these lines, an utterance of (6c_g) may plausibly be interpreted as intended to convey that Irina possesses an understanding of how to do a quintuple salchow, albeit a *mere* understanding since she does not have the ability to perform one herself. Arguably, this would be a straightforward *conversational implicature* of (6c_g) – a part of what is *meant*, even if not *said*.²

² Implicature is an important notion, but there is considerable debate about how to characterize it and put it to use (cf. e.g. Blome-Tillmann 2013; Davis 2014). For a further discussion of implicatures in ascriptions of know-how, see Capone (2011).

This reinterpretation is easily at hand because of a fact which many others have also cited (cf. Fridland 2012, 22 fn. 16). There is a common-sense distinction between the phrase ‘to know how to A’ and phrases such as ‘to know *about* how to A’ or even ‘to know how *one* As’. For clearer reference, here is a list of the forms of these alternatives:

- (7) (a) S knows how to A.
- (b) S knows about how to A.
- (c) S knows how one As.

Since these distinctions between (7 a–c) are easily missed in everyday speech, an inadequate use of (7 a) is easily understood as intended in the sense of (7 b–c), which, arguably, expresses an *understanding* of A-ing rather than the competence to A *oneself*. Of course, the *result* of this reinterpretation is very close to what Bengson & Moffett claim to be the literal meaning of such an utterance. The *pragmatic* result is very close to their *semantic* assumption. Still, this shows that their conclusion is not mandatory.

Distinguishing between the two distinct readings of (6 c) given in (6 c_s) and (6 c_g), I have argued that there is a good explanation of the acceptability of (6 c), even if it is correct to hold that ‘knows how to’ expresses the concept of know-how or competence which includes ability. This brings me to the second sentence quoted by Bengson & Moffett, (6 d) on page 213. I shall suggest that this sentence can be treated analogously.

Again, I shall disambiguate the relevant phrase ‘can actually do one’ and present a clear distinction between general ability and specific opportunity:

- (6 d_s) Not only does Irina know how to do a quintuple salchow, she can actually do one *here and now*.
- (6 d_g) Not only does Irina know how to do a quintuple salchow, she can actually do one *in the right circumstances for performing this stunt*.

Again, (6 d_s), the first reading involving specific opportunity, is entirely acceptable. The fact that somebody knows how to do something does not entail that this competence can be exercised here and now. But this was not Bengson & Moffett’s point. Instead, they would need to claim that even (6 d_g) is unproblematic. However, this claim is clearly much less acceptable than (6 d_s). It may even indeed be ‘redundant’, as they put it.

Again, one may also ask if there is a sense in which (6 d_g) is acceptable at all. But I take it to be obvious that semantically redundant statements may nevertheless be unproblematic because they, too, will typically receive a charitable reinterpretation in the relevant context. Along these lines, the first conjunct in (6 d_g) would not be understood literally, along the lines of

(7a) on the preceding page, but instead along the lines of (7b–c), i.e. as meant to express a mere understanding. Even if this is not what is *said*, it may nevertheless be precisely what is *meant*. Again, it is natural to assume that this would be a conversational implicature.

If this is what is meant, then the whole of (6d_g) will turn out to be informative, after all. For it will then be interpreted as stating that Irina does not merely know *about* how to perform this stunt, but that she is actually able to perform it herself – presumably, by being guided by her knowledge about the stunt. Again, the result of this reinterpretation is what Bengson & Moffett take to be the literal meaning of (6d_g) – the pragmatic implicature leads to what they posited as the semantic meaning. But the point remains that this shows that their conclusion is not the only explanation of why (6d_g) is acceptable.

In sum, I have argued that the intuitive acceptability of the sentences proposed by Bengson & Moffett fails to show that one may ascribe know-how without ascribing ability. Instead, these facts can also be explained in terms of the easy availability of charitable reinterpretations and pragmatic implicatures, while maintaining the view that ability is necessary for know-how. However, this purely pragmatic story is not the only explanation available to those who wish to reject the argument from Bengson & Moffett. I will add further semantic considerations in § 7.3, § 7.4 and § 7.5.

This brings me to the second line of argument by Bengson & Moffett, developed in collaboration with Jennifer Wright (2009). They take issue with part of Alva Noë’s reply to an example by Stanley & Williamson which I have discussed in § 5.1, *Ski Instructor* on page 150. Noë writes:

Is it Stanley and Williamson’s view that, if polled, most English speakers would share their intuition that the instructor is unable to do the jumps even though she knows how to do the jumps? I would predict that this is not true, or rather, that the outcome of such a poll would depend on how we tell the back-story. (Noë 2005, 283)

Bengson, Moffett and Wright comment:

In order to settle this issue, we tested Noë’s prediction by giving 194 participants the following vignette:

Pat has been a ski instructor for 20 years, teaching people how to do complex ski stunts. He is in high demand as an instructor, since he is considered to be the best at what he does. Although an accomplished skier, he has never been able to do the stunts himself. Nonetheless, over the years he has taught many people how to do them well. In fact, a number of his students have won medals in international competitions and competed in the Olympic games.

Participants were asked both whether Pat knows how to perform the complex stunts and whether Pat is able to perform the complex stunts. [...] [T]he vast majority (81%) judged both that Pat knows how to perform the stunts and that he is unable to do them. (Bengson *et al.* 2009, 391–392)

In order to broaden their evidential base, they proceed to also test *Quintuple Salchow* on page 150, or a mildly expanded version of it:

[W]e gave 190 participants in the same study an additional vignette:

Jane is an Olympic-caliber figure skater practicing a complex jump called the Salchow. When one performs a Salchow, one takes off from the back inside edge of one foot and lands on the back outside edge of the opposite foot after one or more rotations in the air. A single Salchow requires one complete rotation. A double requires two. A triple requires three. A quadruple requires four. And a quintuple requires five. Like virtually all Olympic skaters, Jane is consistently able to perform a triple Salchow. Although Jane can land a quadruple Salchow one out of every three attempts, she is unable to do a quintuple Salchow. In fact, at the present time, nobody is able to perform one. Nevertheless, Jane wants to be the first skater to ever land a quintuple Salchow and so she occasionally practices them in her free time. She knows that in order to do a quintuple Salchow, she must take off from the back inside edge of one foot and land on the back outside edge of the opposite foot after five complete rotations in the air. Whenever she attempts this, however, she cannot make it around the full number of rotations without falling.

Participants were asked both whether Jane knows how to do the quintuple Salchow and whether Jane is able to do the quintuple Salchow. [...] [T]he vast majority (76%) judged both that Jane knows how to do the quintuple Salchow and that she is unable to do it. (Bengson *et al.* 2009, 392–393)

They conclude that “the folk are perfectly comfortable attributing know-how in the absence of ability” and that “ordinary judgments of know-how appear to be insensitive to the absence of ability.” (Bengson *et al.* 2009, 393) This suggest that know-how is independent form ability.

Arguments along those lines are typical for *experimental* philosophy, the attempt to use empirically elicited statistics of intuitive judgments as evidence in philosophical discussions. While I cannot discuss this approach in general,³ I will still comment on the present case.

³ Of course, there are important arguments to be made both against and in favor of experimental philosophy in general. Compare Alexander & Weinberg (2007) for an overview, and Sytsma & Buckwalter (2016) for a collection of most recent contributions to the debate. Vocal criticisms of experimental philosophy with which I am sympathetic include Sosa (2007b) and Kauppinen (2007). Plausible defenses of such approaches include Knobe (2007) and Nagel (2012).

Here, I am very sympathetic with Alva Noë’s own reply (2011). He argues that all such empirical results are able to contribute is raw data of intuitive judgments rather than philosophically reflected insight. The core point is that “whatever people say could at most be the beginning of our conversation, not its end” (2011, 198–199). In other words, it is still an open question how this data is supposed to be explained.

I think I can lend some further weight to these considerations. In fact, I contend that the empirical results presented by Bengson, Moffett & Wright do not constitute an *additional* argument over and above the first line of argument by Bengson & Moffett on which I have already commented. These results merely confirm a fact on which I agree – that statements of the form ‘S knows how to A, but is unable to A’ are sometimes acceptable. But it remains an entirely open question how this should be explained.

Above, I have argued that what Bengson & Moffett take to be the literal meaning of such statements can equally well be revealed as a pragmatically plausible reinterpretation of an otherwise inconsistent claim. As I shall argue, an analogous line of argument can also be employed in order to explain the majority judgments in the studies just quoted.⁴

The key to this idea is a fact about the English language which I have already exploited above, in (7 a–c) on page 215 – the subtle, but important distinction between the phrase ‘to know how to A’ and phrases like ‘to know *about* how to A’ or ‘to know how *one* As’, which, as I have suggested, could very well track the distinction between full-blown know-how and mere understanding I have been defending

Crucially, the options with which the participants in the studies by Bengson, Moffett and Wright were confronted did *not* include options which track this distinction. They were only asked to assess sentences of the forms ‘S is able to A’ and ‘S knows how to A’, i.e. (7 a), with no third option along the lines of (7 b) or (7 c).

If it is true that there is a subtle semantic distinction between (7 a) and (7 b–c) which can easily be missed, then the *absence* of this contrast can, in the light of the presence of only the *single* contrast between ‘S is able to A’ and ‘S knows how to A’, create the impression that the latter, i.e. (7 a), should be read along the lines of (7 b–c). In particular, this single

⁴ I should also note that the empirical studies by Bengson, Moffett and Wright also contain a second pair of vignettes (cf. Bengson *et al.* 2009, 395–397). These are modelled on *Avalanche* which I have quoted on page 26 and on *Lucky Salchow* which I will discuss on page 162. In these cases, the intuitive verdicts of the participants of their studies happen to confirm the view I have defended in § 1.3 and § 1.5 and which I will continue to defend in § 5.3 and § 5.4. That is, most participants hold that these are cases of mere abilities and not of genuine know-how.

contrast will typically create the pragmatic assumption that the contrast between these options is both *informative* and *relevant* with regard to the case under consideration. A natural way to save this pragmatic assumption is to interpret (7a) along the lines of (7b–c).

In sum, both lines of argument in favor of the view that it is admissible to ascribe know-how without ascribing ability fail to constitute an argument against the picture of know-how I have been defending in this book. If my proposal is correct, then these phenomena can simply be explained pragmatically, in spite of apparent semantic inconsistencies.⁵ However, it seems plausible that these phenomena are not *merely* pragmatic. In § 7.3 and § 7.4, I shall therefore also consider semantic explanations of these phenomena. § 7.5 will then suggest a way in which these pragmatic considerations can lead to the development of a specific semantic phenomenon.

§ 7.3 The Argument from Translation

I have argued that it is at least pragmatically acceptable to employ the expression ‘knows how to’ in cases without ability. This is true even if this expression uniformly expresses the concept of know-how, competence or skill. In this section, I shall discuss an alternative explanation of the acceptability of such ascriptions which is not only pragmatic, but instead semantic. Following David Wiggins, I shall call this hypothesis “the bifurcational view” (Wiggins 2012, 123). It states that ‘knows how to’ can express at least two non-equivalent concepts – both genuine competence and something which falls short of genuine competence, i.e. an understanding of an activity or propositional knowledge about the activity.

Of course, this idea is not new. Ian Rumfitt and Greg Sax, for example, have suggested that what Ryle meant to express with the phrase ‘knows how to’ was what one may also reformulate with the existing noun ‘know-how’ in order to disambiguate the otherwise unclear meaning of the verb phrase ‘knows how to’ (cf. Rumfitt 2003, 166; Sax 2010, 514). Similar hypotheses have already been defended by many other philosophers.⁶

⁵ Obviously, my pragmatic explanation builds on conversational maxims along the lines of those discussed in the seminal work of Herbert Paul Grice (1989). This is nothing new in the debate about experimental philosophy. For example, Antti Kauppinen (2007) has proposed a general argument along those lines with which I am sympathetic. But as I have indicated on page 217, I cannot discuss this here in general.

⁶ The defenders of this view include Hintikka (1975), Carr (1979; 1981a), Katzoff (1984), Rumfitt (2003), Rosefeldt (2004), Lihoreau (2008), Damschen (2009), Wiggins (2009; 2012), Sax (2010), Brogaard (2011), Devitt (2011), Ren (2012), Abbott (2013), and Kremer (2016), among many others.

In this section, I will present an argument in favor of the bifurcation hypothesis. But this argument will leave open how exactly this conclusion should be understood, for example as positing a form of semantic ambiguity or in other ways. I will come back to these questions later, in § 7.4 and § 7.5.

My argument for the bifurcation hypothesis follows a pattern which is well-established among linguists and semanticists. It builds on the question how an expression would be translated into other languages. My argument can thus be dubbed ‘The Argument from Translation’. The idea is that if an expression in one language can be translated into multiple non-equivalent expressions in another language, then this is a *prima facie* reason to assume that this expression can itself express these multiple non-equivalent things rather than only one thing. Saul Kripke expresses this idea as follows:

“Bank” is ambiguous; we would expect the ambiguity to be disambiguated by separate and unrelated words in some other languages. Why should the two separate senses be reproduced in languages unrelated to English? First, then, we can consult our linguistic intuitions, independently of any empirical investigation. Would we be surprised to find languages that used two separate words for the two alleged senses of a given word? If so, then, to that extent our linguistic intuitions are really intuitions of a unitary concept, rather than of a word that expresses two distinct and unrelated senses. Second, we can ask empirically whether languages are in fact found that contain distinct words expressing the allegedly distinct senses. If no such language is found, once again this is evidence that a unitary account of the word or phrase in question should be sought. (Kripke 1977, 268)

Many philosophers have applied these considerations concerning translatability to the expression ‘knows how to’, most prominently with regard to German. I contend that (8) can be translated as (9 a–c) into German:⁷

- (8) Linda knows how to perform a backflip.
- (9) (a) Linda weiß, wie man einen Rückwärtssalto macht.
 (b) Linda kann einen Rückwärtssalto machen.
 (c) Linda versteht sich darauf, einen Rückwärtssalto zu machen.

⁷ Gregor Damschen has suggested that ‘to know how to A’ can also be expressed in German as ‘S weiß A zu tun’ (cf. Damschen 2009, 281). Thus, (8) would be translated as ‘Linda weiß einen Rückwärtssalto zu machen’, which would also support my argument. While I have not been able to engage in a proper linguistic study of these issues, I find this sentence very untypical. I generally suspect that such constructions occur only very rarely, probably mostly with respect to activities affecting one’s social relations to others, paradigmatically in expressions like ‘Sie weiß sich zu wehren’ or ‘Er weiß zu gefallen’. Again, there is no conclusive evidence either way. But at the very least, Damschen is wrong to claim that German is just like French in this respect, where (8) would indeed be translated without any interrogative particle as ‘Linda sait faire un saut périlleux en arrière’. I will come back to the case of French in § 8.2.

The translation of (8) as (9a) is certainly correct. But it does not explicitly express the sense of ‘knows how to’ in which it expresses competence. Instead, it merely asserts that Linda has knowledge *about* how to perform a backflip. This is supported by the fact that a literal translation of (9a) back into English would not lead to (8), but instead to ‘Linda knows how one performs a backflip’. After the German verb ‘*wissen*’ and an interrogative particle like ‘*wie*’, it is grammatically impossible to use an infinite verb phrase, unlike in English after ‘to know’ and ‘how’. Instead, this verb phrase needs to be finite and therefore requires a subject, in this case, the generic ‘*man*’ (cf. Wurmbrand 2001, 107).

By contrast, the translation of (8) as (9b) is also often correct. For example, the title of Ryle’s chapter, “Knowing how and Knowing that”, is typically not translated as “*Wissen, wie und Wissen, dass*” but instead as “*Wissen und Können*” (cf. Rosefeldt 2004, 377). But this translation fails to express two things explicitly – first, that Linda has an ability to perform a backflip rather than that there is only a mere possibility of her doing so (cf. § 1.4), and second, that her ability is not a *mere* ability, but an intelligent ability, a competence (cf. § 1.5).⁸ Again, this is supported by the fact that a literal translation of (9b) back into English would not lead to (8), but instead to ‘Linda can perform a backflip’.

A translation which makes these things clearer than (9b) is given in (9c). This German translation has so far not been noted in the debate about know-how, but I think that this is indeed the most faithful translation of ‘knows how to’ when it expresses the concept of know-how. Like the English construction ‘knows how to’, the German expression ‘*sich verstehen auf*’ takes an infinite verb phrase as an argument. And like this English construction, it uses a verb denoting a canonical epistemic relation – ‘*verstehen*’, ‘to understand’ – with respect to the activity denoted by the infinite verb phrase. It thereby makes very explicit that what is attributed is an epistemic achievement.

I contend that (9a) and (9c) are both equally plausible translations of (8). The English expression ‘knowing how to do something’ can be translated into German both as ‘*wissen, wie man etwas tut*’ and as ‘*sich darauf verstehen, etwas zu tun*’. The former expresses knowledge *about* an activity, an understanding of what it takes to do well in that activity, whereas the

⁸ Andreas Ditter has argued that the German ‘*können*’ has at least one sense in which it expresses a specifically cognitive ability along the lines of know-how, and that this sense is etymologically primary (cf. Ditter 2016, 508). But since there are other senses as well, this is compatible with my claim that ‘*können*’ does not explicitly express the fact that the possibility in question is one which stems from a genuine competence.

latter expresses the skill or competence to engage in that activity oneself, guided by such an understanding. Thus, the Argument from Translation clearly supports the bifurcational hypothesis.

But this is certainly only a defeasible argument. It is an open question how other languages are supposed to be understood in this respect, and in any case, translations cannot be entirely devoid of background theories and assumptions. For example, one may plausibly object that I only propose (9c) as a translation of (8) because it is particularly congenial to the positive account of know-how I offer, an account partly in terms of the notion of understanding which occurs explicitly in (9c). And even if I reassure my critics that I honestly think that these points are independent, it is hard to see how this debate is supposed to be solved.

But despite the fact that my argument may not be absolutely convincing on its own, it is nevertheless strong in the dialectical context of the current debate about know-how. This is because the general principle behind the Argument from Translation is explicitly endorsed by the *opponents* of my view who hold that the English expression 'knows how to' uniformly expresses propositional knowledge.

Jason Stanley, for example, employs a similar Argument from Translation when he correctly remarks that the English verb 'to know' has two distinct senses, a propositional knowledge sense on the one hand, and a sense of *objectual acquaintance* on the other hand (cf. Stanley 2011b, 36–37). Stanley's own examples are the following:

- (10) (a) John knows Bill.
- (b) John knows the mayor of Boston.

I will come back to the important notion of objectual knowledge in §8.2. For now, the important thing is how one may justify the view that objectual knowledge and propositional knowledge are different. Stanley points to the fact that 'knows' in these sentences would be translated into German with '*kennen*' rather than with '*wissen*' – or with '*connaître*' rather than '*savoir*' in French. As he argues, this is evidence for the claim that 'knows' can express the two non-equivalent concepts of objectual acquaintance and of propositional knowledge (cf. Stanley 2011b, 37).

I completely agree with this argument. But the same kind of argument also shows that 'knows', when followed by 'how to', can even express a third concept, competence. This is what follows from the fact that 'knowing how to do something' can be translated into German both as '*wissen, wie man etwas tut*' and as '*sich darauf verstehen, etwas zu tun*'. Unfortunately, this point is missed by Stanley and many other philosophers who explicitly

assert that the German translation of ‘knowing how to do something’ is *always* ‘*wissen, wie man etwas tut*’ (cf. Stanley 2011b, 37; Brogaard 2011, 138). However, my above discussion of (9 a–c) has shown that this is false.

In fact, some philosophers endorse the even stronger claim that the occurrence of ‘knows’ in ‘knows how to A’ and in ‘knows that p’ is translated with the *same* verb in *all* languages other than English (cf. Stanley & Williamson 2001, 437; Stanley 2011b, 37; Brogaard 2011, 138). As I have argued, this is already falsified by German. But it is nevertheless important also to look at other languages, especially if my argument concerning German should turn out to be unconvincing, after all.

While I shall leave a detailed cross-linguistic analysis to those with genuine expertise in linguistic typology, there already is strong evidence that a number of languages translate ‘knows how to’ with different non-equivalent expressions – where one sense entails ability and one sense does not. The languages for which this has already been shown are Russian, Cantonese, the Native American language Montana Salish, and Turkish (cf. Rumfitt 2003, 164; Abbott 2013, 10; Ditter 2016, 503–506). Crucially, *each* of these languages provides the resources for a version of the Argument from Translation entirely analogous to the one I made for the case of German.

This concludes my argument in support of the bifurcation hypothesis. The Argument from Translation may be defeasible, but at the very least, it shifts the burden of proof back to the opponents of the bifurcation hypothesis. After all, it is incoherent to accept an argument for the substantial distinction between ‘knows’ in the sense of objectual acquaintance and ‘knows’ in the sense of propositional knowledge, but to simultaneously *deny* an entirely *analogous* argument which establishes a further substantially different concept of know-how or competence.⁹

§ 7.4 Ambiguity Tests

I have supported the bifurcation hypothesis, i.e. the view that ‘knows how to’ can express the distinct concepts of genuine know-how or competence on the one hand and of mere understanding of or knowledge about an activity on the other hand. The most natural way to understand this hypothesis would be to claim that this expression is simply semantically ambiguous. However, Bengson & Moffett have claimed that there are a number of established ways for testing the ambiguity of an expression, all of which show that ‘knows how to’ is *not* ambiguous (cf. Bengson & Moffett 2007, 38–40).

⁹ Andreas Ditter has meanwhile proposed a similar argument (cf. Ditter 2016, 511–512).

These tests are taken from a seminal paper by Arnold Zwicky and Jerold Sadock (1975). In this section, I will show that these tests fail to establish this conclusion and instead suggest that ‘knows how to’ may plausibly be taken to be ambiguous after all.

These considerations should not be understood as a partisan philosophical criticism of standard methodology in linguistics. It is already well-established that the tests proposed and discussed by Zwicky & Sadock are far from conclusive. In an overview article on the debate about ambiguity, Adam Sennet comments:

These tests generally depend on the presence or lack of interpretations and on judgments regarding the ridiculousness of interpretation (the absurdity of the meaning is known as *zeugma*—though it should probably be known as *sylllepsis*). These judgments can be difficult to make especially in tricky philosophical cases, so expect that the tests may be of less help than we might hope for at first. (Sennet 2016, sect. 4)

However, one still has to make the case that know-how is indeed such a ‘tricky philosophical case’. General considerations aside, I shall therefore address these tests, as applied by Bengson & Moffett, in detail.¹⁰

The first such test deals with what linguists call ‘conjunction reduction’ (cf. Zwicky & Sadock 1975, 17–18), and it is the most prominent one in the debate about ‘knows how to’. According to Bengson & Moffett,

[(11 a)] is *zeugmatic*, indicating that the verb in question (‘drove’) standardly expresses multiple non-equivalent concepts. On the other hand, [(11 b)] is non-zeugmatic: (Bengson & Moffett 2007, 39)

- (11) (a) Irina drove her trainer crazy and her mother to the airport.
- (b) Irina knows how to do a quintuple salchow and elementary addition.

¹⁰ I only address three of the four tests discussed by Bengson & Moffett, leaving out the one involving verb phrase deletion (Bengson & Moffett 2007, 39; Bengson *et al.* 2009, 393). My reason for doing so is that I fail to see where this test is endorsed by Zwicky & Sadock (1975). At the very least, these authors certainly do not do so on the page mentioned by Bengson & Moffett. Instead, Zwicky & Sadock even appear to explicitly contradict Bengson & Moffett’s view on that page, saying that the deletion of verb phrases – as in Bengson & Moffett’s example ‘I didn’t see her duck, but Irina did’ – “excludes crossed understandings” of the ambiguous expression ‘her duck’ (Zwicky & Sadock 1975, 19). In other words, the use of ‘her duck’ in this sentence cannot refer to an animal in one half of that sentence and to a quick evasive movement in the other. By contrast, Bengson & Moffett hold that there *is* a ‘grammatically anomalous reading’ of this sentence on which this is in fact the case, and they claim that since such a reading is not available for ‘I don’t know how to do a quintuple salchow, but Irina does’, ‘knows how to’ is not ambiguous (cf. Bengson & Moffett 2007, 39). Again, Zwicky & Sadock simply seem to disagree with the premise of this test. As do I.

To begin with, there are two general problems with this ambiguity test. First, as has been pointed out since the seminal work of Zwicky & Sadock, zeugmaticity is gradual and context-dependent (cf. Sennet 2016, sect. 4.6–4.7; Lewandowska-Tomaszczyk 2007). It is therefore unclear how much weight such tests actually carry. I come back to this later, on page 232. Second, and likewise, Bengson & Moffett simply assert which of these sentences they assess as zeugmatic and which as non-zeugmatic. But it seems likely that intuitions about zeugmaticity and about ambiguity are not independent from each other but are strongly correlated instead. It is therefore unclear if such a test constitutes a genuine argument for ambiguity from the independent phenomenon of zeugmaticity or merely an illustration of the intuition of ambiguity via a corresponding intuition of zeugmaticity.

Still, I shall bracket these principal problems from now on and assume for the sake of argument that we do have sufficiently clear intuitions of zeugmaticity. Even so, I think there are clear cases of completely acceptable sentences involving conjunction reduction where it is nevertheless true that the verb in question expresses two non-equivalent concepts. For example:

(12) Irina knows Berlin and that winter can be quite harsh there.

A widely shared argument for the zeugmaticity of this sentence again employs the crucial principle behind the Argument from Translation discussed in § 7.3. According to this view, ‘knows’ in the first half of (12) expresses an epistemic relation of acquaintance. Irina knows Berlin in the sense that she is acquainted with this city – what one would express with ‘*kennen*’ rather than ‘*wissen*’ in German or with ‘*connaître*’ rather than ‘*savoir*’ in French (cf. page 222). By contrast, the eluded ‘knows’ in the second half of that sentence does not express the relation of acquaintance with the proposition that winter can be quite harsh in Berlin. Irina is not *merely* acquainted with this proposition, she *knows* it to be true. Thus, objectual acquaintance and propositional knowledge are two distinct epistemic relations which can both be expressed in English with the aid of the verb ‘to know’.

As I have already stressed, this view is shared very widely.¹¹ For my purposes, it is crucial that Jason Stanley (2011b, 36–37) and, in a later paper, John Bengson and Marc Moffett (2011c, 178–180) also endorse this position explicitly. The latter point out that, “[a]s is well known, ‘that’-complements can be conjoined with complements that denote vastly different types of entity, including propositions, properties and objects” (Bengson & Moffett

¹¹ Stanley & Williamson reject this view, but with a much too complicated example of an ungrammatical sentence, “Hannah knows that penguins waddle, and Bill, Ted.” (Stanley & Williamson 2001, 437) Whatever the precise reason for the ungrammaticality of this sentence, it is very plausible to hold that (12) is grammatical.

2011c, 178–179 fn. 34), referring to a seminal linguistic paper by Ivan Sag and colleagues and to two discussions of this point for 'knows how to' (cf. Sag *et al.* 1985; Roberts 2009; Ginzburg 2011; Abbott 2013, 8–9).

On this basis, my argument runs as follows. Bengson & Moffett have argued that (11 b) is not zeugmatic, but that 'knows how to' could only express multiple non-equivalent concepts if (11 b) were zeugmatic after all. However, they agree with the widespread view that (12) is not zeugmatic either, despite the fact that 'knows' in this sentence *does* express the multiple non-equivalent concepts of objectual acquaintance and propositional knowledge. But if the ambiguity of 'knows' does not require the zeugmaticity of (12), then there is no reason to demand that the ambiguity of 'knows how to' should require the zeugmaticity of (11 b).

This brings me to the next ambiguity test employed by Bengson & Moffett which involves potential contradictions (cf. Zwicky & Sadock 1975, 7–8). Bengson & Moffett write (cf. Bengson *et al.* 2009, 393):

Clearly [(13 a)] has a reading on which it is not contradictory, indicating that one of the terms occurring in the sentence (in this case, 'bank') standardly expresses multiple non-equivalent concepts. On the other hand, there is no reading of [(13 b)] on which it is not contradictory: (Bengson & Moffett 2007, 39)

- (13) (a) Irina deposited her check in the bank, but she didn't deposit her check in the bank.
- (b) Irina knows how to do a quintuple salchow, but she doesn't know how to do a quintuple salchow.

There are two things to say in reply to this charge.

First, I contend that how unacceptable a sentence like (13 b) appears crucially depends on further contextually salient information. Bengson & Moffett do not comment on this problem, despite the fact that Zwicky & Sadock stress this very clearly when presenting this kind of test. They point out that sentences like their own example (14 a) "are not contradictions" and that "additional information brings this out", as illustrated by (14 b) (Zwicky & Sadock 1975, 7). I will come back to this later, on page 233.

- (14) (a) That dog isn't a dog.
- (b) That dog isn't a dog; it's a bitch.

Second, the ease with which an unproblematic reading of (13 a) is available may also in part depend on the fact that it involves an ambiguous *noun*, 'bank', whereas the allegedly ambiguous expression in (13 b) is the complex *verbal* expression 'knows how to'. It seems that there is a general preference

for the acceptability of such sentences when they involve ambiguous nouns as opposed to verbs. To illustrate, the information conveyed in (15 a) and (15 b) is identical, but (15 b) seems more acceptable than (15 a).

- (15) (a) He answered her question, but he didn't answer her question; he didn't address her point at all.
- (b) His answer to her question wasn't an answer to her question; he didn't address her point at all.

In tandem, these two points cast serious doubt on the claim that “there is no reading of [(13 b)] on which it is not contradictory” (Bengson & Moffett 2007, 39). Adding some additional information to this sentence, and formulating a second version involving the noun ‘know-how’ yields the following:

- (16) (a) Irina knows how to do a quintuple salchow, but she doesn't know how to do a quintuple salchow; she has a perfect understanding of what she needs to do to perform a quintuple salchow, but she is not competent to do one herself.
- (b) Irina's know-how to do a quintuple salchow isn't know-how to do a quintuple salchow; she has a perfect understanding of what she needs to do to perform a quintuple salchow, but she is not competent to do one herself.

To the extent that (16 a) or (16 b) are acceptable and non-contradictory, (13 b) is equally acceptable and non-contradictory. And arguably, (16 b) is perfectly fine and (16 a) may plausibly be okay, at least in some situations. Certainly, these sentences are not as clearly unproblematic as (14 b), ‘dog’ being much more clearly ambiguous than both ‘know-how’ and ‘knows how to’. But I think that they are at least sufficiently unproblematic to cast doubt on the way in which Bengson & Moffett appeal to this test.

According to the final ambiguity test in Bengson & Moffett's repertoire, if ‘knows how to’ was ambiguous and thus standardly expressed, say, two non-equivalent concepts, then [(17)] would have four distinct readings arising from various combinations of the two non-equivalent concepts allegedly expressed by ‘knows how to’. Theoretical prejudices aside, [(17)] does not have four distinct readings. (Bengson & Moffett 2007, 39–40)

- (17) Irina knows how to do a quintuple salchow, and she knows how to add.

Suppose that the two concepts between which ‘knows how to’ is ambiguous are indeed competence and mere understanding. If so, then this example

is very ill-chosen. As quoted in *Quintuple Salchow* on page 150, Bengson & Moffett point out themselves that nobody currently has the competence to perform a quintuple salchow and that all anybody currently has is an understanding of how to perform it. And as discussed on page 212, Bengson & Moffett point out themselves that it is impossible to have a mere understanding of addition since the understanding of this activity already entails the competence to engage in it. Thus, (17) has only one plausible reading – the one according to which Irina has the mere understanding of how a quintuple salchow is to be performed and furthermore possesses the competence to add.

Evidently, Bengson & Moffett's third ambiguity test should only be discussed with respect to sentences where each activity admits of the difference between genuine competence and mere understanding. For example:

(18) Irina knows how to do a salchow, and she knows how to dance.

I contend that, in principle, this sentence can indeed be read in four ways:

- (19) (a) Irina has the competence to do a salchow, and she has the competence to dance.
- (b) Irina understands how to do a salchow without being competent at doing so herself, and she understands how to dance without being competent at doing so herself.
- (c) Irina has the competence to do a salchow, and she understands how to dance without being competent at doing so herself.
- (d) Irina understands how to do a salchow without being competent at doing so herself, and she has the competence to dance.

While all of these readings seem to be available in principle, it seems to me that (19a) and (19b), those which ascribe the same kind of epistemic state in both conjuncts, are the most natural ways to read (18). But in suitable contexts with further information, (19c) and (19d) may be equally available. At the very least, it is not obvious that they are *never* available.

In sum, I have discussed and rejected Bengson & Moffett's arguments against the view that 'knows how to' is ambiguous. As already mentioned, all of these considerations depend on the availability and strength of intuitions about what can or cannot be said or what does and what does not sound strange. It seems almost equally problematic for me to attack the arguments by Bengson & Moffett than it is for them to make them in the first place. Thus, much will depend on further considerations such as those discussed in the other sections of this chapter.

§ 7.5 The Polysemy of ‘knows how to’

I have argued that there is no conclusive evidence against the view that ‘knows how to’ is ambiguous. Instead, it may very well be ambiguous, after all, especially given the considerations in support of the bifurcation hypothesis in § 7.3. However, semantic ambiguity is only one way of cashing out this hypothesis. In this section, I shall consider such other views and eventually present my favorite account, the view that ‘knows how to’ is polysemous.

To begin with, Ephraim Glick has suggested what I would like to dub a ‘bifurcation of kinds’, i.e. the view that there are different kinds of know-how (Glick 2012, 120–122). He argues that we may understand this view in analogy with “the claim that there are two kinds of memory, working memory and long-term memory, or the claim that there are two kinds of pain, emotional and bodily.” (Glick 2012, 121) Along these lines, I may also hold that we should distinguish two different sub-concepts of the concept of know-how – genuine competence and the mere understanding of or knowledge about an activity – but that this does not require any ambiguity in the English expression ‘knows how to’. As with ‘pain’, Glick writes, “[t]here is simply a theoretically significant division within the extension of an ordinary expression.” (Glick 2012, 121)

Similarly, Barbara Abbott has called attention to what I propose to call a ‘contextual bifurcation’. She pointed out that one and the same word may express different concepts since it is sensitive to the conversational context, particularly to the kinds of objects under discussion. She writes:

[A] word like *red* is interpreted differently when we are talking about apples (the skin is red), grapefruits (the edible part is red), hair (it may be the same color as things which we call “orange”), and so forth. The point here is that even if the verb *know* in English is unambiguous, that does not mean that it could not be used to talk about two (or more) very different kinds of knowledge. (Abbott 2013, 15)

Thus, it is perfectly possible to defend the bifurcational hypothesis without positing any semantic ambiguity of ‘knows how to’.

While the Rylean responsibilist account defended in Part One of this book is neutral with respect to all of these options, I take it that the most plausible account is still another one – a polysemy view of ‘knows how to’. This is what I shall discuss and defend in the remainder of this section.

Polysemy is closely related to ambiguity, but still an importantly different phenomenon. Roughly, an expression is semantically ambiguous when

it has multiple non-equivalent meanings. This includes cases of homonymy like the paradigm example ‘bank’, which has two non-equivalent meanings which have nothing to do with each other whatsoever. By contrast, a polysemous expression has multiple non-equivalent meanings which stand in a sufficiently close relationship. This rough characterization is far from sufficient, of course, among other things because it leaves entirely open what *kind* of relation this is supposed to be, and what it means for such a relation to be sufficiently close. Still, this sketch will be sufficient for my present purposes. And as far as I can see, it seems to be common ground in an otherwise rather complicated debate in philosophy of language and linguistic semantics about how best to cash out this notion (cf. e.g. Falkum & Vicente 2015; Sennet 2016; Lewandowska-Tomaszczyk 2007).

A standard example of polysemy is the expression ‘book’. Consider the following sentences.

- (20) (a) This book is too heavy to carry around all day.
 (b) This book is too difficult to read on the train.

In (20 a), ‘this book’ is used to refer to a specific edition or copy – a physical object. By contrast, the occurrence of ‘this book’ in (20 b) refers to the content of the text – to the abstract work which can be instantiated in many different ways. These are clearly different meanings of ‘book’, but these different meanings are also closely related. A ‘book’ as a specific physical object *instantiates* a ‘book’ as an abstract work.

Such a relation of instantiation is *one* example of a kind of relation between two different meanings which indicates polysemy. Other kinds of relation include constitution – e.g. ‘wood’ for either the material or a small forest –, causal proximity – e.g. an artist’s name for either the person or her work –, and many others. For my present purposes, I can leave open how these kinds of relations should be systematized and explained.¹² Still, I hope that examples like these provide a sufficiently clear intuitive grip on the notion of polysemy.

With this background, we are in a position to consider the question of polysemy with respect to ‘knows’ and ‘knows how to’. In fact, it has already been argued that ‘knows’ is polysemous – i.e. that this verb expresses the related, but distinct meanings of objectual acquaintance and propositional knowledge. Barbara Abbott has already suggested this view (cf. Abbott

¹² The three families I mentioned – instantiation, constitution, and causal proximity – are briefly, but illuminatingly discussed by Viebahn & Vetter (2016) who go on to apply these lessons to the semantics of modal expressions. There is also strong cross-linguistic evidence for various further and more specific kinds of relationships between the distinct meanings of polysemous expressions (cf. e.g. Srinivasan & Rabagliati 2015).

2013, 7) and Michael Kremer has made a convincing case not only for the its truth, but also for the interpretive claim that Gilbert Ryle also held this view (Kremer 2016, 5–7; cf. Ryle 1945b). Thus, treating ‘knows’ as polysemous is particularly congenial to my project in Part One, the development and defense of a Rylean account of know-how.

However, my current concern is not merely with ‘knows’ in general, but with ‘knows how to’ in particular. § 7.3 has argued that a single sentences with ‘knows how to’ such as (8) on page 220 can express both of the distinct concepts of genuine know-how or competence on the one hand and understanding of an activity on the other hand. This, however, already strongly suggests the view that ‘knows how to’ is indeed polysemous rather than ambiguous. After all, genuine competence and an understanding of an activity are very closely related indeed.

What is the *nature* of the relation between these distinct meanings? In fact, having an understanding what it takes to do well in an activity is a *necessary condition* for, indeed a *proper part* of, having the competence to engage in this activity. Those who have such an understanding form a *proper subset* of those who have the relevant competence. Given this close relationship, if it is indeed correct that ‘knows how to’ can express both of these meanings, then it is overwhelmingly plausible to treat this expression as polysemous.

One might point out that this semantic diagnosis crucially requires a substantial account of know-how, understanding and other notions such as the view developed in Part One of this book. But this is not problematic at all. As the debate about polysemy suggests, the formation and functioning of polysemy crucially involves appeal to further, external information. For example, Vyvyan Evans argues that “polysemy arises, in large measure, from the sorts of non-linguistic knowledge we possess, and which we draw upon during language understanding.” (Evans 2015, 122) Thus, one may well have objections against the Rylean responsibilist account of know-how in general and the necessity of ability for know-how in particular. But whatever one’s substantial commitments are, in philosophy and elsewhere, they can always be one among many factors in discussing questions of polysemy in particular, and of semantics in general.

The polysemy of ‘knows how to’ also allows a fresh look at some of my earlier discussions in this chapter. First, it suggests a way in which pragmatic considerations such as those in § 7.2 may be connected with, and in fact lead to, the semantic phenomenon of polysemy. And second, ambiguity tests such as those discussed in § 7.4 appear to be less promising when dealing with polysemy.

As for the first point, §7.2 has argued that, even if 'knows how to' uniformly expresses the concept of competence, the acceptability of semantically problematic sentences can be explained with the aid of pragmatic considerations regarding contextually salient alternatives. Thus, even if 'knows how to' expresses competence, and competence only, some uses of this expression may still conversationally implicate mere understanding.

This consideration is mirrored precisely by the view of many intellectualists. According to them, 'knows how to' does not semantically express competence or skill, but only a mere understanding or propositional knowledge about an activity. Still, 'knows how to' may nevertheless be used to attribute genuine competence, because this may be what is *meant*, even if not *said*. In this vein, Bengson & Moffett explicitly hold that this is a conversational implicature (cf. e.g. Bengson & Moffett 2007, 35).

Thus, there are two rival semantic claims, each with the potential to explain some of the intuitive appeal of the opposing view by appealing to conversational implicatures. However, an excellent way to make sense of this result is precisely to assume that 'knows how to' is indeed polysemous.

To see this, one may pick *any* of the two rival accounts just sketched. On such an assumption, the pragmatic mechanisms which allow for the charitable reinterpretations just rehearsed may, over time, *create* just the form of polysemy I have advocated here. Given the crucial role of non-linguistic knowledge pointed out by linguists like Evans (2015), such mechanisms are well-established and discussed. The *pragmatic* phenomenon of conversational implicatures may, over time, *create* the *semantic* phenomenon of ambiguity or polysemy, in a process called 'pragmatic strengthening' (cf. Traugott 1989; Falkum 2015; Viebahn & Vetter 2016, 5).

Thus, whatever one's *initial* view about the *semantics* of 'knows how to' may be, the *pragmatics* of the use of this expression clearly invites a reconsideration of these initial semantic views, and strongly suggests to assume that 'knows how to' is polysemous.

The second aspect which the polysemy of 'knows how to' promises to illuminate concerns the ambiguity tests in §7.4. In fact, it supports my earlier arguments at two points.¹³

On the one hand, the first test looked for suitable cases of zeugma as a condition for ambiguity, as in (11 a) on page 224. However, I have argued that there is no such requirement for ambiguity. If the expression in question is polysemous, then the expectation to detect zeugmaticity is even less

¹³ These considerations are partly parallel to an extended discussion by Emanuel Viebahn, who argues that standard tests for ambiguity are generally inadequate when dealing with polysemous rather than homonymous expressions (cf. Viebahn 2016).

plausible. For example, ‘book’ in (21) refers to the abstract work in the first conjunct and to a token physical copy in the second conjunct, without any strange air of zeugmaticity.

- (21) This book is a classic of world literature and it is overdue at the library.

On the other hand, the second test involving contradictions, as in (13a) on page 226, has turned out to be inconclusive since other contextually salient information may resolve the relevant contradiction, as already pointed out by Zwicky & Sadock, from whom Bengson & Moffett borrow these tests (cf. Zwicky & Sadock 1975, 7). However, Zwicky & Sadock’s example to support this, (14a–b) on page 226, crucially draws on a *polysemous* expression, ‘dog’. Clearly, ‘dog’ may mean ‘male dog’ or ‘dog of any sex’, where the latter is a necessary condition of the former – just like, on my account, having an understanding of an activity is a necessary condition of having the competence to engage in it.

I have defended the view that ‘knows how to’ is polysemous and applied it to shed light on some earlier considerations. To conclude this section, I would like to discuss an important worry with this view.

This worry can be presented as an argument on behalf of my intellectual opponents like Bengson & Moffett who argue that ‘knows how to’ does not express the concept of competence, but merely the concept of an understanding of or knowledge about an activity, while genuine competence may still be what is *meant* even if not said – i.e. it may be a conversational implicature. The current discussion of polysemy has shown that these concepts are closely related, indeed that one of them is *necessary* for the other one. But then, the argument goes, we are able to avoid the idea of a bifurcation altogether. Instead, we may apply the principle of theoretical parsimony and argue that the expression in question only expresses the more basic of the two concepts. This line of argument would maintain the view that ‘knows how to’ semantically expresses the knowledge *about* or understanding of an activity, independently of the question of ability and genuine competence.

What to make of this? I agree that theoretical parsimony is an important virtue. But I do not think that it outweighs all the other considerations discussed so far. I would like to briefly highlight three of these aspects.

First, the intuitive pre-theoretic case in § 7.1 relied on the fact that a typical authoritative dictionary and thesaurus of the English language, Merriam-Webster Online, explicitly includes a competence-sense of ‘to know’ and ‘knowledge’. However, that which has already made its way into the lexicon of a language cannot be *merely* pragmatic.

Second, the Argument from Translation in § 7.3 also suggests that the phenomenon is a semantic one. At the very least, arguments of this pattern are widely taken to establish semantic claims, even by intellectualists about know-how. Then, however, it would require an *additional* argument why this particular case should be different. Without such a further argument, maintaining the view that 'knows how to' does not semantically express competence, but only conversationally implicates it, would be *ad hoc*.

Finally, as just discussed on page 232, such pragmatic mechanisms and implicatures are bound to change the semantics and lead to the very polysemy I prefer. Thus, even if there may be no knock-down evidence on either side, partly because of the complicated distinction and division of labour between semantics and pragmatics, I maintain that the polysemy of 'knows how to' is the overall most plausible view.

Thus, part of my disagreement with intellectualists is merely verbal. But there is a further part to this disagreement. Regardless of such linguistic questions, everybody should agree that there is a clear distinction between full-blown competence and the mere understanding of or knowledge about an activity. And this distinction is indeed accepted by intellectualists. However, chapter 9 will argue that *only* the concept of full-blown competence can perform the crucial job of explaining intelligent practice – the job for which the Rylean concept of know-how is designed. Crucially, my intellectualist rivals *agree* with this explanatory aim. If it is true that this concept is indeed expressed by 'knows how to', then this constitutes a powerful reason to maintain that 'knows how to' genuinely expresses the concept of competence. In fact, then, the Wittgensteinian motto quoted at the beginning of Part Two of this book, on page 145, is quite apt.

Chapter 8

From Language to Intellectualism

In chapter §7.1, I have discussed questions of the linguistic analysis of ‘knows how to’, but entirely omitted what is arguably the most important approach in the linguistic analysis of ‘knows how to’ – the linguistic theory which Stanley & Williamson (2001) have most prominently appealed to in order to establish the intellectualist view that know-how is a species of propositional knowledge. In this chapter, I follow them on the path from language to intellectualism, partly building on the considerations from chapter §7.1, but with a clear focus on the intellectualist metaphysics of know-how. After all, Stanley has correctly noted:

Discussions of semantics are often in fact discussions of metaphysics, carried out in the formal mode. [...] The ability to move smoothly between the material mode and the formal mode is a consequence of the fact that modern semantic theories take the form of inductive characterizations of truth in a language. (Stanley 2011b, 144)

While semanticists may indeed be doing metaphysics in this way, it should be clear that “the important question is whether they are doing *good* metaphysics.” (Habgood-Coote 2017, 5). This chapter will therefore discuss the linguistic and the metaphysical side of the issue, as well as their relationship.

In a first step, §8.1 presents the relevant linguistic approach, what I shall dub the ‘Standard Linguistic Account’ of the syntax and semantics of English sentences of the form ‘S knows how to A’. §8.2 will discuss linguistic objections against this account, most prominently in the form of alternative analyses of the relevant knowledge-ascriptions and in the form of problems from cross-linguistic data, i.e. data from languages other than English.

Then, §8.3 will discuss the way in which the formal mode is transformed into the material mode, i.e. the argument which leads from a linguistic analysis of ‘knows how to’ to a metaphysics of know-how. As I shall argue, this

argument fails to establish an intellectualist account of the nature of know-how. In fact, § 8.4 will argue that what this argument establishes is entirely compatible with the Rylean responsibilist account of know-how established in Part One of this book. While propositionalist intellectualism will be my most prominent target in § 8.3 and § 8.4, I will also make parallel points about the objectualist intellectualism championed by Bengson & Moffett.

Finally, § 8.5 will argue that there are good reasons to prefer Rylean responsibilism over its intellectualist rivals. While part of these considerations have meanwhile also been endorsed by intellectualists (cf. Stanley & Williamson 2016), I contend that this constitutes no convincing defense of intellectualism, but instead suggests further reasons to endorse an alternative view such as the one advocated in Part One of this book. This is the project of the final § 8.6 of this chapter.

§ 8.1 The Standard Linguistic Account

According to a widespread view in the debate about know-how, the semantics of the ascriptions of know-how in English already shows that know-how is a species of propositional knowledge since these involve what linguists call ‘embedded questions’. Elsewhere (cf. Löwenstein 2011a), I have dubbed this view ‘linguistic intellectualism’ and this kind of argument the ‘Argument from Linguistics’, and I have discussed the most prominent champions of this view at the time, Stanley & Williamson (2001). In this section, as well as in § 8.2 and § 8.3, I will in part rely on this previous work in assessing the most recent defense of this view by Jason Stanley (2011b; 2011c).¹

Gilbert Ryle notes that champions of intellectualism hold that “the primary exercise of minds consists in finding the answers to questions” (Ryle 1949, 27). Jason Stanley’s view fits this bill precisely. The core idea is that know-how consists in knowing an answer to a question, namely the question which is syntactically embedded in the sentence attributing know-how. For example, if Gregor knows how to ride a bicycle, linguists tell us that the expression ‘how to ride a bicycle’ is an embedded version of a question like “How can one ride a bicycle?” Then, to say that Gregor knows how to ride a bicycle is just to say that Gregor knows an answer to such a question, that he knows that such-and-such would be a way for him to ride a bicycle. Since this knowledge is propositional knowledge, Stanley concludes that know-how is a species of propositional knowledge.

¹ While Stanley will be the main target of my discussion here, my considerations also extend to cognate versions defended by philosophers like David G. Brown (1970; 1974), Jaako Hintikka (1975; 1992) and David Braun (2006; 2011).

The core of this argument is the following:

Our view of ascriptions of knowledge-how is very straightforward. It is just that the standard linguistic account of the syntax and semantics of embedded questions is correct. (Stanley & Williamson 2001, 431)

For lack of a better term, I shall turn this description into a proper name of the linguistic position Stanley defends – the ‘Standard Linguistic Account’. In this section, I shall present and explain this view.

The Standard Linguistic Account is indeed very widespread among semanticists – linguists and philosophers alike.² I shall not discuss all details of this view, but instead only propose a sketch of those elements of the theory that are relevant for the problems under discussion.

Stanley’s first step concerns the uniformity of knowledge ascriptions involving interrogative particles such as ‘who’, ‘where’, and so forth, including ‘how’. Such knowledge is commonly referred to as ‘knowledge-*wh*’, even if the interrogative particle crucial for the topic of this book is spelled ‘how’ rather than ‘*whow*’. Stanley offers the following examples (2011b, 36):

- (22) (a) John knows whether Mary came to the party.
 (b) John knows why Obama won.
 (c) Hannah knows what Obama will do in office.
 (d) Hannah knows who Obama is.
 (e) Hannah knows what she is pointing at.
 (f) Hannah knows how Obama will govern.

On the standard view, the so-called *wh*-complements in (22 a–f) denote embedded questions – syntactically embedded versions of the corresponding free-standing questions which may be expressed as follows:

- (23) (a) Did Mary come to the party?
 (b) Why did Obama win?
 (c) What will Obama do in office?
 (d) Who is Obama?
 (e) What is Hannah pointing at?
 (f) How will Obama govern?

Cutting a long story short, the knowledge attributed in (22 a–f) is then taken to be the knowledge of at least one contextually relevant proposition

² Stanley relies on the seminal work by Charles L. Hamblin (1958), Lauri Karttunen (1977), and Jeroen Groenendijk and Martin Stokhof (1982; 1984). Still, the details and foundations of this account are subject to an ongoing debate (cf. e.g. Brogaard 2008a; Stout 2010; Krifka 2011; Parent 2014; Cross & Roelofsen 2014) and there are notable rivals (cf. e.g. Roberts 2009; Brogaard 2011; Ginzburg 2011; Michaelis 2011).

which is an answer to the corresponding embedded question in (23 a–f). This view involves commitments in a number of ongoing debates which I shall bracket for my purposes (cf. Parent 2014; Cross & Roelofsen 2014).³

To illustrate, the knowledge attributed in (22 a) is taken to be a contextually relevant proposition which answers the question in (23 a) – say, that, yes, Mary did come to the party. Likewise, the knowledge attributed in (22 b) is taken to be a contextually relevant proposition which answers the question in (23 b) – say, that Obama won because he was, at the time, an inspirational political leader. And so forth.

Of course, the examples in (22 a–f) do not include the crucial kind of case which is relevant for the question of know-how. Clearly, the only example involving the word ‘how’ – (22 f) – does not ascribe know-how to Hannah, but knowledge *about* how somebody *else* will do something. As I have already highlighted in the Introduction, know-how is typically expressed in English by saying that somebody knows how *to do* something – that is, with the verb ‘to know’ followed by an *infinitive*. But this does not make a difference for Stanley. Infinitival *wh*-complements can also be understood in terms of embedded questions. Here are his examples (Stanley 2011b, 113):

- (24) (a) John knows who to call on case of an emergency.
 (b) John knows where to buy an Italian newspaper.
 (c) John knows when to call a doctor.
 (d) John knows whether to call a doctor.
 (e) John knows how to solve the problem.
 (f) John knows what to do.

What are the free-standing questions which correspond to the embedded infinitival *wh*-complements in (24 a–f)?

³ To illustrate, there are at least the following five points of debate. *First*, there is the question if an ascription of knowledge-*wh* merely entails knowledge of at least one proposition which is an answer to the embedded question (what has come to be known as the ‘mention-some-reading’) or if it sometimes or always furthermore entails knowledge of all propositions which are such answers (the so-called ‘mention-all-reading’ in the version of ‘weak exhaustivity’) or if it sometimes or always even entails knowledge of all propositions which are such answers coupled with no belief in a false answer (the ‘mention-all-reading’ in the version of ‘strong exhaustivity’) (cf. e.g. Groenendijk & Stokhof 1982; Stanley 2011b, 115–122). *Second*, this view involves the commitment that the proposition which answers the embedded question needs to be contextually relevant rather than merely a true answer (cf. e.g. Braun 2006). *Third*, there is the question if it is sufficient to know propositions which are, as a matter of fact, answers to the relevant questions or if it is furthermore necessary to know these propositions *as* answers to those questions (cf. e.g. Kallestrup 2009; Stout 2010; Schaffer 2007; Schaffer 2009a). *Fourth*, there is the question what counts as an answer to a question in the relevant sense (cf. e.g. Sgaravatti & Zardini 2008). *Fifth*, it is an important type-theoretical question whether the relata of knowledge-*wh* are propositions or *sets* of propositions (cf. e.g. Schroeder 2012; Stanley 2012b).

As Stanley argues, such constructions introduce two complications (cf. Stanley & Williamson 2001, 422–425). On the one hand, infinitival *wh*-complements involve a modal element which has both an ‘ability- or dispositional’ reading to be paraphrased with ‘can’ and a deontic reading to be paraphrased with ‘ought to’ (cf. Stanley 2011b, 111–114; Bhatt 2006, 122). On the other hand, they introduce what linguists call the unpronounced pronoun ‘PRO’ which occurs after the interrogative particle and is sometimes explicitly inserted in order to make the structure of such a sentence explicit. These occurrences of ‘PRO’ have two readings, a first one under which it is equivalent with the generic ‘one’ and a second under which it is anaphorically dependent on an therefore co-extensional with the subject of the embedding sentence – i.e. ‘John’ in (24 a–f) (cf. Stanley 2011b, 70–76).⁴

This yields a total of four possible readings of such questions and of the corresponding *wh*-complements. For example, (24 b) can be read as:

- (25) (a) John knows where *he ought to* buy an Italian newspaper.
- (b) John knows where *one ought to* buy an Italian newspaper.
- (c) John knows where *he can* buy an Italian newspaper.
- (d) John knows where *one can* buy an Italian newspaper.

Again, the idea is that the knowledge attributed to John in (24 b) is, in each of these readings, the knowledge of a contextually relevant proposition which answers the corresponding embedded question. Given that the interrogative particle in (24 b) is ‘where’, each such proposition refers to a *place*. On reading (25 a), this may be the proposition that *he ought to* buy an Italian newspaper at the train station, on reading (25 d), it may be the proposition that *one can* buy an Italian newspaper at the train station, and so forth.

The very same considerations are then applied to uses of ‘to know’ followed by ‘how’ and an infinitive. To take the example explicitly discussed by Stanley & Williamson (2001, 424–425), (26) has the readings in (27 a–d).

- (26) Hannah knows how PRO to ride a bicycle.
- (27) (a) Hannah knows how she ought to ride a bicycle.
- (b) Hannah knows how one ought to ride a bicycle.
- (c) Hannah knows how she can ride a bicycle.
- (d) Hannah knows how one can ride a bicycle.

⁴ As Stanley points out, this reading requires a *de se* mode of presentation on which (24 a) attributes to John the knowledge who he *himself* can call in the case of an emergency, rather the knowledge what somebody who merely happens to be himself can do (cf. Stanley 2011b, 83–94). I have discussed such *de se* modes of presentation in § 4.3 and pointed out their crucial role intellectual guidance in § 4.4.

While the interrogative particle ‘where’ in (24 b) introduces quantification over places, the interrogative particle ‘how’ in (26) introduces quantification over *ways of doing something* – in this case, ways of riding a bicycle (cf. Stanley & Williamson 2001, 427). This leads to an analysis of (27 a–d) as:

- (28) (a) Hannah knows, of some way *w*, that *w* is how she ought to ride a bicycle.
- (b) Hannah knows, of some way *w*, that *w* is how one ought to ride a bicycle.
- (c) Hannah knows, of some way *w*, that *w* is how she can ride a bicycle.
- (d) Hannah knows, of some way *w*, that *w* is how one can ride a bicycle.

Since Stanley stresses that all of these readings of (26) are possible, I would like to note that this already commits him to the view that ‘knows how to’ can express several substantially different things. Arguably, this is itself a clear case of polysemy since the two different modal senses – deontic and abilitative – are distinct, but conceptually related, as Viebahn & Vetter (2016) show for modal expressions in general. Even if this falls short of the full form of polysemy I advocated in § 7.5, it shows how important this issue is even for Stanley, who explicitly rejects that ‘knows how to’ is ambiguous.

However, Stanley holds that (28 c) is the “paradigm reading” of (26) (Stanley & Williamson 2001, 425), its “most natural interpretation” (Stanley 2011b, 114). This is the reading he goes on to single out as the canonical expression of know-how. I will also grant this simplification until § 8.4.

I shall furthermore bracket some important further remarks about the relevant modes of presentation and the modal nature of these propositions because they are not relevant at this point. These will be introduced in § 9.2, where I discuss the question how this version of propositionalist intellectualism can explain the notion of intellectual guidance.

I conclude that Stanley’s view of ‘knows how to’ – the Standard Linguistic Account – holds that “S knows how to A” is true just in case S knows, for some way of A-ing *w*, that *w* is how she can A (cf. Stanley 2011b, 122).

§ 8.2 Uniformity and Universality

I have presented the Standard Linguistic Account of the syntax and semantics of ‘knows how to’. In this section, I would like to discuss some of the objections which have been raised against this view. I shall here consider

only those two objections which I take to be most decisive and most important – the question of the uniformity of all knowledge-*wh*, and the question of the universality of this account with respect to other languages.⁵ A further problem, the fact that ascriptions of know-how are gradable, unlike ascriptions of propositional knowledge, will be discussed in § 8.5.⁶

The first objection concerns something which is very dear to Stanley, the uniformity of ascriptions of knowledge-*wh* with infinitival -complements such as those quoted in (24 a–f) on page 238. He writes:

It is a common assumption between the Rylean and the Intellectualist that sentences involving constructions like “know where + infinitive”, “know when + infinitive”, “know why + infinitive”, etc. all can be defined in terms of propositional knowledge. But given that ascriptions of knowing-how in English look so similar to such ascriptions, it is hard to see how they could ascribe a different kind of mental state. This provides a powerful argument in favor of the conclusion that our ordinary folk notion of -how is a species of propositional knowledge. (Stanley 2011c, 208)

Stanley makes a strong case for the view that these sentences should all be treated in the same way since their form is exactly the same. But the further premise that all of these examples other than the ones involving ‘how’ ‘can be defined in terms of propositional knowledge’ is far from ‘a

⁵ There are at least two further objections with which I sympathize, but which I cannot discuss here. First, Kent Johnson has objected that the Standard Linguistic Account is committed to the view that ‘how’ is a *sui generis* existential quantifier with a number of peculiar features (cf. Johnson 2006, 25–26 fn. 3). Second, John Collins has proposed a general argument to the effect that the appeal to syntax in theories like the Standard Linguistic Account does not have philosophical consequences which are as straightforward as philosophers like Stanley suggest (cf. Collins 2007).

⁶ One further objection stems from Bengson & Moffett. They hold that infinitival *wh*-complements probably do not denote propositions, or questions, or sets of propositions (cf. footnote 3 on page 238 on these options). They suggest that this view requires that not only (29 a) is a perfectly grammatical sentence, which is true, but that at least one of (29 b–d) is also grammatical, which is false (cf. Bengson & Moffett 2011c, 182, 182 fn. 43). But Joshua Habgood-Coote has already shown that these criteria are inadequate (cf. Habgood-Coote 2017, 9–10). However, there seems to be an even deeper issue here. Bengson & Moffett assume that what is picked up anaphorically by ‘it’ at the end of (29 b–d) is ‘how to swim’. Instead, I think that ‘it’ refers back only to the infinitive ‘to swim’. This can be seen in (29 e–f), which I have added to their list. Thus, this objection against the Standard Linguistic Account fails.

- (29) (a) Michael knows that *w* is a way to swim; so it must be true.
 (b) Michael knows how to swim; so it must be true.
 (c) Michael knows how to swim; it is easily answered.
 (d) Michael knows how to swim; it is nonempty.
 (e) Tom knows how to swim; so it cannot be too difficult.
 (f) Tom knows why to swim; after all, he sees its health benefits.

common assumption'. I shall discuss two alternative ways in which one may account for the uniformity of such sentences.

First, one may hold that knowledge-*wh* is not the *knowing* of an answer to the embedded question, but the capacity to *choose* such an answer. Arguments for this view have been proposed by Meghan Masto (2010) and Katalin Farkas (2016b; 2016a).⁷ Since this view is very close to what I have already presented and only adds the extra step of a capacity for answering, I will not discuss this proposal in more detail at this point. However, my argument for the compatibility of Rylean responsibilism with the Standard Linguistic Account will be very much in the spirit of this idea (cf. § 8.4).

Second, one may account for the uniformity of ascriptions of knowledge-*wh* not in terms of *propositional* knowledge, but in terms of *objectual* knowledge. After all, as discussed in § 7.3, this can also be expressed by the verb 'to know', and Stanley explicitly agrees that this is the case.

Let me spell out this option more carefully. Consider, for example, the following schematic sentences:

- (30) (a) x knows how to A.
 (b) x knows where to A.
 (c) x knows why to A.
 (d) x knows when to A.

John Bengson and Marc Moffett have suggested that (30 a–d) can be paraphrased as (31 a–d) (cf. Bengson & Moffett 2011c, 180), and one may even consider alternative paraphrases, e.g. by replacing 'the' in (31 a–d) with 'a'.

- (31) (a) x knows the way (in which) to A.
 (b) x knows the location (at which) to A.
 (c) x knows the reason (for which) to A.
 (d) x knows the time (at which) to A.

On this view, *wh*-complements may be understood as what linguists call 'free relatives' rather than, as the Standard Linguistic Account would have it, as embedded questions or as the propositions or the sets of propositions which constitute answers to these questions. Thus, on this alternative view, 'how to A', 'where to A' etc. do not denote the questions 'How to A?', 'Where to A?', etc. or the propositions that w is a way to A, that p is a

⁷ Unlike the argument in Masto (2010) and Farkas (2016b), the considerations in Farkas (2016a) crucially involve the hypothesis of the extended mind (cf. Clark & Chalmers, 1998). But this idea is not essential to my present point. For a discussion of the possibility of extended know-how, see Carter & Czarnecki (2016), who show how know-how may be extended even if anti-intellectualism is true and know-how is just a disposition.

place at which to A, etc. or even sets of propositions of this form. Instead, these are interpreted as nominals, and they denote the properties of being a way to A, a place at which to A, and so forth.

Of course, this alternative view would need to be spelled out much further.⁸ And over and above the specifically linguistic questions, it is also still unclear what exactly knowledge in the sense of objectual acquaintance is supposed to be, whether and where it varies between the different properties invoked by the different interrogative particles, and so forth. But at least some proposals along these lines are already on the table (cf. e.g. Bengson & Moffett 2011c; Brogaard 2011; Abbott 2013). In the specific case of ‘knows how to’, it is perfectly possible and arguably even most natural to say that “to stand in the know relation to a technique is to be competent with that technique” (Abbott 2013, 5). Thus, such an alternative linguistic analysis has at least some sufficient *prima facie* plausibility, even if it faces a number of problems (cf. Habgood-Coote 2017).⁹

But such an alternative view also promises certain merits over the account in terms of embedded questions. One such merit has been discussed by Kent Bach.¹⁰ The view just presented would do better at explaining the difference between knowledge *how* to do something and knowledge *about* how to do something. It would explain “why it makes perfectly good sense to say that a person can know a lot about how to play golf, how to write a philosophy paper, or how to improvise at the keyboard” (Bach 2012).

It seems odd to construe such knowledge *about* how to do something as meta-level knowledge *about* questions or propositions or sets of propositions. It is much more natural to say that this is knowledge *about ways* to do these things. On this view, knowledge *about how* to A clearly turns out to be propositional, and part of the explanation of this fact is that knowledge *how* to do something (without the ‘about’) is objectual rather than propositional. Unfortunately, Stanley himself does not discuss the distinction between ‘S knows how to A’ and ‘S knows about how to A’ – a distinction which has already played an important role in § 7.2. Even worse, he threatens

⁸ For a general argument for the importance and viability of complementing propositional attitudes with attitudes towards objects, see Grzankowski (2014).

⁹ One particularly striking problem is the possibility that the best linguistic account of ascriptions of objectual knowledge like (31 a–d) may analyze them in terms of concealed questions, i.e. as (30 a–d). After all, to know *Jen* just is to know *who Jen is*, right? This would threaten to undercut the very distinction between objectual knowledge and knowledge-*wh*. For discussion of such approaches, see Brogaard (2008b) and Habgood-Coote (2017). Note, however, that this option is not easily available to those who, like Stanley, accept the Argument from Translation (cf. § 7.3) which established a clear difference in the kind of knowledge ascribed here.

¹⁰ For a related argument, see Abbott (2013, 5–6).

to blur this distinction by, perhaps unwittingly, switching back and forth between the expressions ‘knowledge how’ and ‘knowledge of how’ in his own writings.¹¹ In § 9.2, I will come back to this problem.

The second objection against the Standard Linguistic Account concerns the plurality of natural languages. As already discussed in part in § 7.3, there are at least some and probably quite a number of natural languages in which the English construction ‘knows how to’ has no exact analogue. One example is provided by languages like German where, in fact, there is no such thing as an infinitival *wh*-complement (cf. § 7.3).

A second important case, pointed out by Ian Rumfitt (2003), concerns Romance languages like French where infinitives are used in the relevant constructions which translate ‘knows how to’, but where interrogative particles are neither required nor even admissible. As Rumfitt remarks, (32) translates as (33 a), but (33 b) is ungrammatical, even though ‘*comment*’ is the interrogative particle which translates ‘how’.

(32) Pierre knows how to swim.

(33) (a) Pierre sait nager.

(b) * Pierre sait comment nager.

A further example along these lines can also be found outside the Romance languages, in Modern Greek, where corresponding constructions do not employ an interrogative particle either (cf. Douskos 2013, 2331–2332).

Third, as discussed in more detail in § 7.3, there is a number of languages which translate ‘knows how to A’ and ‘knows that p’ with different and clearly non-equivalent expressions. And finally, there are thousands of languages which have not been explicitly discussed with respect to these questions. In sum, this casts serious doubt on the universal applicability of the Standard Linguistic Account.

In reply, Stanley has reacted to the second of these four points, the one concerning languages like French, and argued that we should postulate, next to the unpronounced constituent ‘PRO’, a further unpronounced constituent which quantifies over ways, just like ‘how’ does in English (cf. Stanley 2011b, 138–139; Stanley 2011c, 229). One may plausibly have general reservations about stipulating too many such unpronounced elements

¹¹ One of the best examples can be found on the first page of the preface of Stanley’s book *Know How* which begins with ‘knowing how’ and ends with ‘knowledge of how’, on the clear implicit assumption that these are equivalent: “The thesis of this book is that knowing how to do something is the same thing as knowing a fact. It follows that learning how to do something is learning a fact. For example, when you learned how to swim, what happened is that you learned some facts about swimming. Knowledge of these facts is what gave you knowledge of how to swim.” (Stanley 2011b, vii)

(cf. Abbott 2013; Douskos 2013). But even beside this point, and even beside the first point concerning languages like German, the other two considerations remain most pressing. If there are languages which explicitly use different non-equivalent verbs for the English occurrences of ‘know’ in ‘knows how to A’ and in ‘knows that p’, then it is at least very unclear if the analysis Stanley defends for English can be upheld as universal.

I conclude that the Standard Linguistic Account, while indeed the most widespread view of ‘knows how to’, nevertheless faces serious problems.

§ 8.3 The Limits of the Linguistic Approach

As already indicated, Stanley employs the Standard Linguistic Account presented in § 8.1 in order to derive a *metaphysical* view about know-how. Since his semantic view is given in terms of truth conditions, one can apply the disquotational scheme to infer a metaphysical claim. While Stanley is very explicit about this step of the argument (cf. Stanley 2011b, 144), there seems to be no discussion of his further inference from the bisubjunction thus obtained to a metaphysics of the *nature* of know-how, according to which “knowing how to do something consists in knowing the answer to a question” (cf. Stanley 2011b, 131). In this section, I will discuss a number of possibilities how this argument can be resisted.¹²

Stanley’s argument can be summarized as follows:

(H) *The Argument from Linguistics*

- (H1) ‘S knows how to A’ is true just in case S knows, for some way of A-ing w, that w is how she can A.
- (H2) S knows how to A just in case S knows, for some way of A-ing w, that w is how she can A. *from* (H1)
- (H3) S’ knowledge how to A consists in S’ propositional knowledge, for some way of A-ing w, that w is how she can A. *from* (H2)

An obvious motivation for scepticism about the Argument from Linguistics stems from the fact that, as discussed in § 8.2, there is significant controversy about these linguistic theories and even about the exact linguistic data. As Barbara Abbott suggests, one should not rely on such shaky grounds when justifying a substantive metaphysical view (cf. Abbott 2013, 12–13). While

¹² I shall bracket a number of further proposals. For example, Paulo Santorio (2016) has suggested to model a way to reject the argument from linguistic considerations to propositionalist intellectualism on norm-expressivism in metaethics. This is especially interesting in the context of Rylean responsibilism where norms and normativity play a crucial role. But I will remain neutral on the question of expressivism here.

I am sympathetic with this worry, it is also clear that it goes both ways. The same problem recurs when one supports a metaphysical claim not with a semantic, but, well, with a further metaphysical claim.

But there are also other and more specific objections to the Argument from Linguistics. Most prominently, it has been suggested that we should not start with some contingent natural language expression of know-how, but simply with scientific facts (cf. e.g. Wallis 2008; Devitt 2011). It is widely held that cognitive scientists have established the independence of know-how from propositional knowledge. And according to the present line of objection, we should not infer the falsity of this claim simply on grounds of the Standard Linguistic Account. But unfortunately, § 6.5 and § 6.6 have shown that these considerations are not fully supported by the evidence.

Stanley has also given a further response to these worries. He has insisted that there is no general problem with inferring (H2) from (H1) by using the disquotational scheme. He argues that this is not a case where this scheme should be restricted because the sentences in question contain self-referential or other problematic expressions (cf. Stanley 2011b, 144–146). This step should be seen as entirely uncontroversial.

This defense is certainly correct, but it is not particularly charitable. As Jessica Brown and Barbara Abbott have independently argued, the point that science is relevant for an account of the metaphysics of know-how should be understood as the claim that science is also important for the best semantic account of the relevant expressions (cf. Abbott 2013, 17; Brown 2013, 3–4). As Brown points out, this is what has happened in determining the truth of the semantic claim that “x is water” is true just in case x is H₂O. And already in the very paper criticized by Stanley, Michael Devitt explicitly writes that he takes the scientific evidence straightforwardly as evidence against the Standard Linguistic Account (cf. Devitt 2011, 215).

In sum, these considerations make clear that the disquotational scheme works both ways. Stanley is correct to insist that (H2) stands and falls with (H1). But his critics should therefore be understood as, correctly, holding that (H1) stands and falls with (H2).

So far, I have discussed two kinds of objections against the Argument from Linguistics – worries about (H1), the Standard Linguistic Account in the first premise of the argument, which have turned out to be substantial, and worries about the inference from this semantic claim to the metaphysical claim (H2), which have turned out to be either inconsistent or disguised attacks on the semantic claim (H1) itself. I shall now formulate two further objections against this argument – one which disputes a crucial detail in the initial semantic claim (H1), and one which disputes the inference from the

first metaphysical claim (H2) to the second, reductive metaphysical claim (H3). As I shall argue, these two objections both rely on the fact that the connections between the truth of ‘S knows how to A’, and S’ relations to corresponding propositions may be just as the Standard Linguistic Account demands, but grounded in a very different metaphysics of know-how.

The first of these objections relies on an observation by Ephraim Glick (2011). His surprisingly simple yet decisive point is that something may be ‘knowledge which has a proposition as a relatum’, and in this sense ‘propositional knowledge’, but not propositional knowledge in the substantive sense of knowledge that something is the case (cf. Glick 2011, 412). Further, having a proposition as a relatum is itself nothing particularly demanding since propositions are abstract entities posited on the grounds of our explanatory aims.¹³ This may simply be a question of convention and even include ability. Glick writes:

For any action of φ ing, we could map S’s ability to φ onto the proposition that S φ s, and instead of saying that S is able to φ , we could say that S “ables that he φ s”. If we had this linguistic convention, we might note that “abling” is a relation to a proposition, but of course, by hypothesis, we would be talking about the same thing we actually talk about with ability attributions. (Glick 2011, 413)

Glick beautifully points out a gap in the Standard Linguistic Account, as stated in (H1). The linguistic theory which allegedly supports this view does not entail anything about the *nature* of the relation between the person to whom know-how is ascribed and the propositions in terms of which the *wh*-complement is analyzed. True, it must be a relation denoted by the verb ‘to know’. But it is an open question which kind of relation this is. In § 8.3, I have mentioned that this relation may also be the capacity to answer the relevant question (cf. Mastro 2010; Farkas 2016b), or the direct objectual acquaintance with the ways of doing something which play a role in these propositions (cf. Bengson & Moffett 2011c).

Before discussing the consequences of this point in more detail, I shall now turn to a further objection and go on to discuss both of these in tandem.

This second objection consists in rejecting the inference from (H2) to (H3). That is, one may agree that somebody has know-how just in case they also have correlated propositional knowledge, but point out that this view does not entail the view that know-how *consists* in this propositional knowledge. Combined with the first objection, one may reject the inference to the claim that know-how *consists* in a relation expressed by ‘knows’ to the relevant propositions, even if this relation is not propositional knowledge.

¹³ For the question of propositions and propositional attitudes, see footnote 7 on page 67.

The general possibility of this view should already be obvious from my initial presentation of the Argument from Linguistics. After all, (H2) states a metaphysical *correlation*, according to which S knows how to A just in case she also possesses knowledge that w is a way for her to A. By contrast, (H3) states a metaphysical *reduction* of the know-how on the left hand side of this biconditional to the propositional knowledge on its right hand side. However, (H2) does not entail (H3) and is, in fact, compatible with a number of other claims. Any metaphysical correlation of the form (I) – like (H2) – is compatible with views of all of the forms expressed in (I1–3).

(I) Every x is F just in case x is G.

(I1) F consists in G.

(I2) G consists in F.

(I3) F and G are distinct, and both F and G are grounded in H.

(I3) employs the notion of an ontological ground which has gained some prominence in contemporary ontology and metaontology (cf. e.g. Schaffer 2009b; Audi 2012; Fine 2012). But the basic idea behind this concept can be stated rather easily. To say that something is ontologically grounded in something else is to say that the former exists *in virtue of* the latter.

To illustrate, conjunctions of two claims of the forms (I) and (I3) are plausibly true in the case of the properties of having mass and having volume. Something has mass just in case it also occupies space, but these properties are distinct. Arguably, both having mass and having volume are grounded in having matter. A further case concerns the relations of being better and of there being a reason to favor something over something else. Arguably, x is better than y just in case there is a reason to favor x over y. However, it still remains an open question if the relation of being better makes it the case that there is such a reason, or if it is the other way around, or even if both are grounded in something else.

Thus, the Argument from Linguistics fails. If the Standard Linguistic Account is true at all, then it still fails to establish propositionalist intellectualism but only supports a much weaker claim instead.¹⁴ The argument

¹⁴ I should briefly note that there is a partly analogous problem about what the Standard Linguistic Account predicts for sentences like ‘Tom knows what it is like to be a bat’ which are crucial for the famous ‘knowledge argument’ against physicalism (cf. e.g. Nagel 1974; Jackson 1982; Jackson 1986; Levin 1986; Nemirow 1990; Mellor 1993; Alter 2001; Nida-Rümelin 2009; Howell 2011). In line with my present suggestions, such sentences may also be only superficially or only derivatively propositional and in fact grounded in something else, for example in the acquaintance with the phenomenal character of what is described in the infinitival verb phrase – here, with being a bat. But I cannot deal with this here. For discussion, see Stanley & Williamson (2001), Snowdon (2003), Cath (2009), Glick (2011), Löwenstein (2011a) and Tye (2011).

either stops immediately after the first step and merely establishes (H2) on page 245 which I repeat below. Or it even falls short of this, establishing (H2') instead, which leaves what kind of knowledge relation open.

- (H2) S knows how to A just in case S knows, for some way of A-ing w, that w is how she can A.
- (H2') S knows how to A just in case S stands in some relation, expressed by 'knows', to the proposition that, for some way of A-ing w, w is how she can A.

In § 8.4, I shall come back to this consequence of the Standard Linguistic Account. But before doing so, I would like to note that my considerations concerning the Argument from Linguistics for propositionalist intellectualism are entirely general. They apply equally to any other attempt to derive a metaphysics of the nature of something from a linguistic account in terms of truth-conditions. Even if one accepts the first, disquotational step, the result is merely a metaphysical correlation, but not an account of the nature of what is under discussion. As pointed out on the preceding page with respect to (I) and (I1–3), such a metaphysical correlation is compatible with different metaphysical views.

To illustrate, consider again the linguistic alternative to the Standard Linguistic Account discussed in § 8.2. If one analyzes the relevant knowledge ascriptions in terms of objectual acquaintance with ways of doing something, one may also offer a variant of the Argument from Linguistics. Bengson & Moffett, who have offered these linguistic alternatives, are very cautious here. They correctly note:

However, it is not clear to what extent the metaphysical distinction between propositions and ways of acting currently at issue corresponds to the linguistic distinction between embedded questions and free relatives. (Bengson & Moffett 2011c, 182 fn. 42)

But even if it may not be explicitly endorsed by anyone, such a variant of the Argument from Linguistics is interesting in its own right. This version may be understood as follows:

(J) *The Argument from Linguistics (Objectualist Variant)*

- (J1) 'S knows how to A' is true just in case S knows w, where w is some way of A-ing in which she can A.
- (J2) S knows how to A just in case S knows w, where w is some way of A-ing in which she can A. *from (J1)*
- (J3) S' knowledge how to A consists in S' objectual knowledge of w, where w is some way of A-ing in which she can A. *from (J2)*

As already indicated, the steps in this argument in favor of objectualist intellectualism are no more plausible than those in the original Argument from Linguistics in (H) in favor of propositionalist intellectualism.

First, the inference from (J1) to (J2) is comparatively unproblematic. Still, one may be cautious about what kind of knowledge relation is strictly necessary. Just like Glick's point against Stanley, one may hold that the objectual knowledge relation to a way of doing something may be constituted by something else – maybe by a kind of ability as Glick suggests against Stanley, maybe by propositional knowledge of a proposition involving a concept of such a way as stated in (H2), or maybe by objectual acquaintance with such a proposition as suggested in (H2').¹⁵

Second, however, it is false that (J2) entails (J3). Know-how may be metaphysically correlated with having objectual knowledge of ways of doing something, but this does not entail that this is what know-how consists in. Know-how may just as well consist in something else, as long as it grounds and explains the fact that it is correlated with such objectual knowledge.

§ 8.4 Common Ground with Intellectualism

I have argued that the Argument from Linguistics fails to establish propositionalist intellectualism, and I have pointed out that a possible analogous argument for objectualist intellectualism fails for the very same reasons. For the sole purpose of defending Rylean responsibilism, this result is entirely sufficient. However, I will now go on to argue that my account is perfectly compatible with the linguistic theories advanced by intellectualists, such as the Standard Linguistic Account, and that it straightforwardly *predicts* the surviving metaphysical consequences of these theories.

I shall argue that the crucial connections between know-how on the one hand and propositional knowledge or objectual knowledge on the other hand follow directly from Rylean responsibilism. These can be derived from the intellectual elements I have established in § 4.1, § 4.2 and § 4.3 – that is, from the guiding states of understanding an activity which consist in conceptual capacities to assess individual performances of this activity. Thus, just like there is agreement between Rylean responsibilism and anti-intellectualism when it comes to the crucial role of genuine ability, there is also agreement between Rylean responsibilism and intellectualism when it comes to the crucial role of intellectual states such as propositional knowledge.

¹⁵ The latter two options will be especially interesting to those who analyze objectual knowledge ascriptions in terms of concealed questions (cf. footnote 9 on page 243).

My argument will be congenial to an objection against Stanley's propositionalist intellectualism which has been proposed by Imogen Dickie (2012). According to her, all that Stanley is able to show is that "a skilled Φ -ing manifests the agent's knowledge of a $\langle w \text{ is a way to } \Phi \rangle$ proposition" (Dickie 2012, 740). As she points out, one can sketch a plausible notion of skill, according to which what Stanley shows is compatible with the following contradictory accounts:

Intellectualism: knowledge before skill—In a case of skilled Φ -ing, S chooses w as a way to Φ *because* S knows that w is a way to Φ ; S is a skilled Φ -er iff S knows a range of suitable $\langle w \text{ is a way to } \Phi \rangle$ propositions. (Dickie 2012, 741)

Anti-intellectualism: skill before knowledge—S is a skilled Φ -er iff S's intentions to Φ are non-lucky selectors of non-lucky means to their fulfilment; a skilled Φ -ing manifests propositional knowledge *because* it is the appropriately generated Φ -ing of a skilled Φ -er. (Dickie 2012, 741)

I am very sympathetic with Dickie's sketch of an account of skill, here summarized as the point that "S's intentions to Φ are non-lucky selectors of non-lucky means to their fulfilment" (Dickie 2012, 741). While I have proposed a more nuanced picture of this concept over the course of Part One of this book,¹⁶ the basic idea of my present argument is the same as Dickie's. If the Standard Linguistic Account is correct, then what grounds the truth of 'S knows that A' may just as well be S' skill rather than her propositional knowledge.¹⁷ However, I shall offer a more thorough discussion of these considerations in this section.

To establish this, it is crucial to return to chapter 7. There, I have argued at length that 'knows how to' can both express genuine competence and mere knowledge *about* or mere understanding of an activity, most plausibly because 'knows how to' is polysemous (cf. § 7.5). Presently, I have stated the surviving metaphysical correlation between know-how and propositional

¹⁶ For example, I have clearly distinguished between a mere ability, which may also fit Dickie's short characterization, and a full-blown piece of know-how (cf. § 1.5), and I have taken into account that the exercise of competences does not always involve intentions and intentional action (cf. chapter 3).

¹⁷ In his reply to Dickie's paper, Stanley unfortunately fails to address this insight of hers altogether (cf. Stanley 2012b). Instead, he tackles an *additional* argument proposed by Dickie which purports to show that, *if* the Standard Linguistic Account is compatible with both intellectualism and the sketched anti-intellectualist alternative, *then* there are reasons to prefer the latter. These considerations concern the question how the myriad routes to acquiring a skill can be explained if skill is to be explained solely in terms of propositional knowledge and its application (cf. Dickie 2012, 741–742; Stanley 2012b, 763–765). But these are indeed additional questions on which I comment elsewhere (cf. e.g. § 2.5, § 4.3 and § 8.5). Dickie's basic point about the compatibility of the anti-intellectualist alternative with the Standard Linguistic Account remains untouched by Stanley's reply.

knowledge which results from the Standard Linguistic Account as (H2) or even (H2') on page 249. And I have stated the surviving metaphysical correlation between know-how and objectual knowledge which results from an alternative view with appeal to free relatives as (J2) on page 249. Thus, the claim I would like to defend can be stated as follows:

- (K) 'S knows how to A' is true just in case S possesses *either* the competence to A, which involves guidance by an understanding of A-ing, *or* a *mere* understanding of A-ing without such guidance.
- (K1) The understanding of A-ing which is crucial in both cases grounds (H2'), i.e. it grounds the fact that there is some relation, expressed by 'knows', between S and the proposition that w is how she can A, for some way of A-ing w. Even more, it grounds (H2), i.e. the fact that S has genuine propositional knowledge of this.
- (K2) The understanding of A-ing which is crucial in both cases grounds (J2), i.e. it grounds the fact that S has objectual knowledge of w, where w is some way of A-ing in which she can A.

To establish this, I shall begin with the notion of a way of doing something which figures prominently in both propositionalist intellectualism (cf. Stanley 2011b; Stanley 2011c; Stanley & Williamson 2001) and objectualist intellectualism (cf. 2007; 2011c). This notion is deliberately unspecific. The only explicit remark in the intellectualist literature is this:

The propositions that concern us will contain *ways of engaging in actions*. To be more precise, we shall take ways to be properties of token events. [...] But we shall not have much more of substance to say about the metaphysics of ways [...]. (Stanley & Williamson 2001, 427)

In the end, I take it to be plausible to identify the ways of doing something appealed to by intellectualists with precisely the types of performances I introduced in § 4.3 and abbreviated as 'X' rather than 'w'. Both notions can be applied on different levels of granularity of individuation. Furthermore, intellectualists explicitly hold that the relevant concepts of ways of doing something are typically possessed implicitly or demonstratively rather than on the basis of a sophisticated description.¹⁸ In § 2.5 and § 4.3, I have made the same point about the concepts involved the propositions involved in learning and in assessing an activity. Thus, what I identified there as

¹⁸ For discussion of this point with regard to propositionalist intellectualism, see Stanley and Dickie (Stanley 2011b, 167–173; Stanley 2012b, 765–766; Dickie 2012, 743–745) and with regard to objectualist intellectualism, see Bengson & Moffett (Bengson & Moffett 2007, 51; Bengson & Moffett 2011c, 189–192).

concepts of types of individual acts, can be seen just as a notational variation on the talk of ‘ways of doing something’.

Against this background, I shall now argue that the crucial role of the epistemic states intellectualists appeal to can be derived from Rylean responsibilism. The key to this argument is that the state of understanding of an activity which is crucial for both genuine competence and a mere understanding of an activity consists in the capacity to assess an activity (cf. § 4.1). Having this capacity requires the relevant epistemic states, and *vice versa*. The exercise of the assessment capacity leads to those states as individual assessments. And the only way in which one may arrive at these assessments is precisely by having and exercising the relevant capacity. Thus, (K) and its sub-claims (K 1) and (K 2) are true.

This argument is most straightforward when it comes to objectualist intellectualism, i.e. with respect to (K 2). Cutting a long story short, objectualist intellectualism holds that knowing how to do something consists in having an understanding of at least one way of engaging in this activity, and that one possesses such an understanding just in case one has reasonably mastered the concept of this way of acting (cf. Bengson & Moffett 2007, 50–54; Bengson & Moffett 2011c, 185–192). They state this view as follows:

Having objectual knowledge of a way *w* of φ -ing while grasping a correct and complete conception of *w* is necessary and sufficient for knowing how to φ . (Bengson & Moffett 2011c, 187)

To know how to φ is to stand in an objectual *understanding* relation to a way *w* of φ -ing. (Bengson & Moffett 2011c, 189)

Thus, objectualist intellectualism appeals to states of objectual understanding of individual ways of doing something, i.e. to *conceptions* of such ways. My considerations in § 4.3 have shown that having a conception of the *whole* of an activity, in the sense of being competent in assessing individual acts of such this activity, also entails at least a minimum of conceptions of specific *types* of acts and their quality as acts of engaging in the activity in question.

The parallel case for propositionalist intellectualism, i.e. with respect to (K 1), is more complicated. To begin with, recall that this statement involves a crucial simplification. The analysis of ‘knows how to’ in Stanley’s propositionalist intellectualism is, as Stanley explicitly admits, an analysis of only *one* of four different readings of this expression. To gain a full statement of what the Standard Linguistic Account actually entails, we must therefore go back to page 239 above and to (26) with the four readings Stanley gives in (27 a–d). This simplification in (K 1) can be amended by replacing (H 2) with the following statement:

- (L) ‘S knows how to A’ is true just in case there is some relation, expressed by ‘knows’, between S and at least one of the propositions in (L 1–4).
- (L 1) For some way of A-ing w, w is how S ought to A.
- (L 2) For some way of A-ing w, w is how one ought to A.
- (L 3) For some way of A-ing w, w is how S can A.
- (L 4) For some way of A-ing w, w is how one can A.

Thus, my claim in (K 1) is more wide-ranging. I hold that knowledge-relations to the propositions in (L 1–4) are predicted by Rylean responsibilism.

The key to this argument has been presented in § 4.3. The relevant relation expressed by ‘knows’ which, according to (L), obtains between the subjects S of true sentences of the form ‘S knows how to A’ on the one hand and at least one of the propositions in (L 1–4) on the other hand, may indeed be the relation of full-blown propositional knowledge. It may simply consist in the relations of propositional knowledge established in § 4.3. That is, if ‘knows how to’ refers to either a full-blown competence or a mere understanding of an activity, then the fact that the understanding involved in both of these states consists in an assessment capacity entails that, by *exercising* this capacity, the person in question gains propositional knowledge of precisely the kind of propositions which are involved in the Standard Linguistic Account. That is, exercising this assessment capacity can be understood as answering the very question which is embedded in the relevant ascription of know-how, just like Meghan Masto (2010) and Katalin Farkas (2016b) have suggested with respect to all knowledge-*wh* (cf. § 8.2).

I shall argue that the propositions in (L 1–4) follow from the propositions involved in an assessment capacity. To do so, I will rely on two important statements (C) on page 118 and (D) on page 120. I repeat both below.

- (C) Every capacity to assess A-ing involves knowledge of at least some principles of assessment of acts and situations of the forms (C 1–2).
- (C 1) An act of the type X in circumstances of the type C is a performance of A-ing of quality Q.
- (C 2) 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.

On the one hand, I will appeal to (C) in order to show that (C 1–2), the propositional knowledge involved in every understanding of an activity in the form of correct assessments of performances and situations, entails what the Standard Linguistic Account predicts as (L 2) and (L 4).

On the other hand, I will appeal to (D) in order to argue that (D 1–2), the special kind of essentially indexical propositional knowledge which

is only involved in cases of genuine competence, entails what the Standard Linguistic Account predicts as (L 1) and (L 3), which, according to Stanley, also require a *de se* mode of presentation (cf. § 8.1).

- (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.

First, I have already argued that we can identify ways of doing something with precisely the types of performances I introduced in (C) and abbreviated as 'X' rather than 'w'. Further, the propositions I introduced in (C 1–2) and (D 1–2) explicitly mention the quality of the relevant acts as performances of A-ing, but (L 1–4) lacks such a specification altogether. However, Stanley suggests that the individual ways w over which (L 1–4) quantify are to be individuated in a sufficiently fine-grained way, such that differences in proficiency can be mapped onto differences in the qualities of these ways of acting (cf. § 8.5). Therefore, it must already be possible to characterize the quality of the performances which instantiate the relevant way w mentioned in (L 1–4). And if so, then it is equally unproblematic to leave out a specific characterization of this quality, which is present in (C 1–2) and (D 1–2), and simply leave an implicit evaluation along the lines of 'sufficiently well', as (L 1–4) evidently presuppose. In sum, the difference between saying that performances of the type X count as good A-ings and saying that what instantiates w is a good way to A is largely terminological.

Finally, there are explicit modal elements in (L 1–4), but not in (C 1–2) or (D 1–2). However, just as the performance's quality was explicit only in the latter and implicitly present in the former, these modalities are implicitly present in (C 1–2) and (D 1–2). This is particularly obvious in the case of the modal 'can'. If, with (C 1), an act of the type X – an act which instantiates the way w – counts as an A-ing, then it follows that w is how *one can* A – i.e. (L 4) – and if, with (D 1), my own act x – an act which instantiates the way w – counts as an A-ing, then it follows that w is how *I can* A – i.e. (L 3). The deontic modal 'ought to' can be derived on the basis of assessments of options as what, *pro tanto*, one *ought to do* if one is evaluating one's acts with respect to A-ing, and if the relevant option is not

outweighed by other and better options. If, with (C2), there is an option to perform an act of the type X – an act which instantiates the way *w* – which counts as an A-ing, then it follows that *pro tanto*, *w* is how *one* ought to A – i.e. (L2) – and if, with (D2), there is an option for me to perform some act *x* – an act which instantiates the way *w* – which counts as an A-ing, then it follows that, *pro tanto*, *w* is how *I* ought to A – i.e. (L1).

Thus, the consequences of the Standard Linguistic Account stated in (L) can be explained very well by Rylean responsibilism.

§ 8.5 Gradability and Novelty

In § 8.4 have argued that understanding an activity in the relevant sense, i.e. having the capacity to assess performances of it, stands and falls with the very epistemic states which, according to intellectualists, *constitute* know-how. But this crucial connection does not establish such a claim of constitution (cf. § 8.3). Alternatively, I have suggested that both full-blown know-how and a mere understanding of an activity are what grounds and explains these epistemic states. In this section, I will argue that this result suggests an argument for preferring the Rylean responsibilist account of know-how over the intellectualist ones. One argument for this conclusion has already been presented in § 6.3 and § 6.4 with respect to the problem of epistemic luck. Now, I shall offer a crucial further argument for the same conclusion, one which concerns the gradability of know-how and the acquisition of situation-specific knowledge on the fly.

The gradability of know-how has already been stressed by Ryle:

We never speak of a person having partial knowledge of a fact or truth, save in the special sense of his having knowledge of a part of a body of facts or truths. A boy can be said to have partial knowledge of the counties of England, if he knows some of them and does not know others. But he could not be said to have incomplete knowledge of Sussex being an English county. Either he knows this fact or he does not know it. On the other hand, it 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, and even the champion has still much to learn. (Ryle 1949, 57–58)

In my own account of know-how, I have argued that the states of know-how – competences – are *themselves* gradeable, partly because know-how is a kind of ability and abilities are gradable (cf. § 1.4) and partly because know-how involves a state of understanding, which can itself be understood

as a capacity and therefore as gradable (cf. § 4.1). By contrast, since the vast majority of philosophers would agree that states of propositional knowledge are indeed *not* gradable¹⁹ – either one knows a fact or one does not know it – this constitutes a *prima facie* problem for propositionalist intellectualism.²⁰ And even if objectualist intellectualists can correctly claim that states of understanding a way of doing something *are* gradable, the fact that they start out with *individual* ways of acting still leads to some similar difficulties.

Consider the two kinds of gradability mentioned by Ryle. Taking his own examples and adding my own labels, Ryle's point is that 'knows how to' can be qualified in terms of *comparative quality*, as in (34 a), and in terms of *partial exhaustivity*, as in (34 c). And we should add that there is also a third kind of gradability akin to the first one, *absolute quality* as in (34 b).²¹

- (34) (a) The champion knows how to play chess better than the student.
- (b) The grandmaster knows how to play chess extraordinarily well.
- (c) After a short lesson, the child knows in part how to play chess.

Stanley has briefly suggested some ways to explain (34 a) and (34 c).²²

First, he proposes to understand cases of comparative quality in terms of comparisons of the quality of the way of acting which is known, such that (34 a) can be understood as (35 a) (cf. Stanley 2011b, 34). Second, while it has been suggested that Stanley cannot account for cases of absolute quality (cf. Bengson & Moffett 2011c, 183 fn. 45), I take it that one can simply understand these in terms of the absolute quality of the way of acting which is known, somewhat on the model of cases of comparative quality, such that (34 b) is understood as (35 b). Finally, as for cases of partial exhaustivity, Stanley seems to suggest that a way of acting can also be understood only as a partial way to do something, i.e. that (34 c) may be paraphrased as (35 c) (cf. Stanley 2011b, 34).²³

¹⁹ The minority position that propositional knowledge is gradable after all has notably been defended by Sosa (2009) and Hetherington (2011a), among others.

²⁰ Of course, one may resort to what I shall call the gambit of detaching skill in § 8.6 and argue that *know-how* is not gradable, but *skill* is (cf. Stanley & Williamson 2016, 9). This requires an additional defense of the view that know-how is indeed *not* gradable. An argument for this conclusion is forthcoming in Pavese (2017).

²¹ By contrast, it would be ungrammatical to say "The boy knows in part that Sussex is a county" or "The geographer knows that Sussex is a county extraordinarily well" or even "The teacher knows that Sussex is a county better than the boy".

²² For a discussion of comparatives and degree constructions, see Schwarzschild (2008).

²³ This point relies on the idea that a proposition can be a partial answer to the question which is embedded in a sentence expressing know-how (cf. § 8.1). Stanley suggests that one can formally define a notion of a proposition's being a partial answer to a question, and that the propositional knowledge ascribed in this way only qualifies as such a partial answer to the corresponding embedded question (cf. Stanley 2011b, 34).

- (35) (a) The champion knows, for some way w_1 , that w_1 is a way for her to play chess, and the student knows, for some way w_2 , that w_2 is a way for her to play chess, where w_1 is a better way to do so than w_2 .
- (b) The grandmaster knows, for some way w , that w is a way for her to play chess, and w is an extraordinarily good way to do so.
- (c) After a short lesson, the child knows, for some way w , that w is a partial way for her to play chess.

Having discussed this topic at the beginning of his book *Know How*, Stanley also proposes a further way to account for the gradability of ‘knows how to’ at the end of that book, a way which nicely fits Ryle’s idea that somebody’s propositional knowledge can be gradable “in the special sense of his having knowledge of a part of a body of facts or truths.” (Ryle 1949, 57)

Of course, when we say that a skilled outfielder knows how to field a fly ball, we do not mean that he knows, of at least one way to field a fly ball, that it gives him counterfactual success in fielding fly balls. [...] What we assert when we assert of a skilled outfielder that he knows how to field fly balls is that he knows *all* of a range of relevant ways that give him counterfactual success in fielding fly balls. Hence, to say of an outfielder in baseball that he knows how to catch a fly ball is to impart to him knowledge of *many* propositions of the form ‘ w is a way for him to field a fly ball’. (Stanley 2011b, 183)

In these terms, (34 a–b) can also be understood as (36 a–b):

- (36) (a) The champion knows for n ways w_i , that w_i is a way for her to play chess, and the student knows, for m ways w_j , that w_j is a way for her to play chess where $n > m$.
- (b) The grandmaster knows, for all relevant ways w_i , that w_i is a way for her to play chess.

Evidently, these two accounts of the gradability of ‘knows how to’ can also be combined. The quality of the ways of acting known as in (35 a–b) and the quantity of these ways as in (36 a–b) can both play a role in explaining the quality of the know-how ascribed. And equally evidently, these two approaches are also available to *objectualist* intellectualism. This offers them the opportunity not only to say that know-how is gradable in the sense that one may have a better understanding of one and the same way of acting than somebody else, but that one may also understand more such ways, and that one may have an understanding of other and better ways.

Thus, intellectualists have resources to accommodate the gradability of know-how. But I take it that Rylean responsibilism still does better.

To see this, consider what these intellectualist explanations entail for the nature of know-how. In the beginning, the claim was that knowledge how to A consists in propositional or objectual knowledge of at least one way of A-ing. In the end, know-how has to consist in a whole set of such knowledge states, which is sufficiently large, as well as in further information (presumably, further knowledge) about the comparative quality of the relevant ways. However, the intuition of the gradability of know-how is the intuition of the gradability of *one single state*. Of course, it is perfectly possible to try to retain this intuition by saying that the single state of know-how must be analyzed as a set of further single states, where the size of this set can then be measured and brought together with measurements of the individual elements of the set. But a much more straightforward way to preserve the intuition would be to maintain the idea that know-how is a single state – a capacity to make the very situation-specific assessments under discussion. Given that I have already argued that this is what grounds and explains these states of propositional or objectual knowledge, this option is easily available and very attractive.

There is also a further argument for this conclusion. The capacity to come to know the relevant ways of doing something is *temporally prior* to actually knowing them. In fact, Stanley implicitly accepts this point himself:

[T]he fact that expertise requires fluid responses to novel situations has no bearing on the thesis that skilled action is acting on the basis of reasons. Skilled action may involve fluid *acquisition* of reasons for acting in novel situations, reasons that are only accessible to one when one is in that situation. The fact that some reasons for acting can only be entertained when one is in the situation is fully consistent with the agent's acting for those reasons. (Stanley 2011b, 182)

Clearly, Stanley's point is that propositional knowledge can guide and explain an expert's conduct even despite the fact that she comes to acquire this knowledge only in the process of performing rather than beforehand. In a later co-authored paper, he makes these points even more clearly:

Having skill in sailing is a state that requires having different knowledge states on different occasions, since knowing what to do to initiate an action at sailing will involve knowing one set of facts under stormy weather conditions, and another set of facts under calm weather conditions. [...] The same kind of knowledge that is used to initiate an activity can also be injected at anytime in the ongoing course of that activity. For example, a tennis player changes her mind and switches from a groundstroke to a drop shot based on the position of the opponent. Such cases of learning are also knowledge. (Stanley & Krakauer 2013, 5)

Thus, the better somebody's know-how is the better is her ability to adjust to novel and previously unknown circumstances. This is the phenomenon

of control I discussed in § 4.4. But Stanley's explicit story presupposes that some of the relevant propositions about how to perform are *not* known *before* the agent encounters the relevant situation. She has to learn these facts *on the fly* and immediately go on to act on them. In terms of objectual understanding rather than propositional knowledge, she does not possess conceptions of all of the relevant ways of acting beforehand, but she comes to understand them in the particular situation.

The crucial point is that, intuitively, learning these things is *not* acquiring *new know-how*, but only *exercising* know-how which was possessed all along, albeit maybe in new ways. The fact that an agent already *has* know-how is precisely what *explains* the fact that she is able to *learn* these facts in the first place.

One might object that there is indeed a sense in which an agent acquires new know-how in such a situation – say, the knowledge how to score a point in tennis against the current opponent right here and now, or the knowledge how to sail safely under precisely those conditions which currently obtain. But this is of not help, and not merely because it relies on an artificially fine-grained level of individuating states of know-how. The problem is that whenever these fine-grained pieces of know-how are acquired on the fly, the reason for this must still consist in the fact that the agent was able to acquire them – i.e. that she had the competence to adjust to these conditions, too. Again, this requires that she had the relevant know-how all along, even if it may adjust to a genuinely novel situation.

Thus, the knowledge of the situation-specific ways which an actor comes to possess on the fly cannot be *identical* with her know-how, but can at most be a *part* of her know-how. In contrast, her understanding of the whole of the activity of A-ing is something she *did* have all along. And this understanding is nothing but the capacity to acquire the relevant knowledge.

This conclusion also extends to Stanley's declared view of the linguistic uniformity of ascriptions of knowledge how with other forms of knowledge-*wh* in English (cf. § 8.2). In a *précis* of his book *Know How*, he writes:²⁴

If certain behavior convinces us that someone knows *how* to catch a Frisbee, that very same behavior would convince us that the person knows *when* to stretch out his arm to catch a Frisbee. Similarly, if someone knows how to get home, then she knows where to go to get home, and vice-versa. The features that convinced Ryle that knowledge how is non-propositional, such as automaticity of action, are present with all states of knowledge-wh. So we should seek a unified account. (Stanley 2012a, 735)

²⁴ For similar considerations, see Stanley & Williamson (2016, 2).

The assessments of self and situation which I have shown to be the core form of the propositional knowledge involved in responsibly controlling oneself explicitly include references not only to ways of acting, but also to situations, i.e. to combinations of time and place, including the social context if relevant for the activity in question. In this sense, knowing when to stretch out one's arm to catch a frisbee is involved in the competence to catch a frisbee because it relies on the underlying competence to assess frisbee-catching. By exercising this underlying competence, one will gain knowledge about individual situations in which stretching out one's arm will be a good or a bad thing to do in order to catch a frisbee. And in one's actual performance, one then acts in the light of such propositional knowledge-when. Similarly, of course, for propositional knowledge-where.

This adds further weight to the view that knowledge-*wh* may in general consist in the *capacity* to answer the embedded question (cf. Masto 2010; Farkas 2016b). And it shows that the way in which the unity of knowledge-*wh* is involved in the explanation of intelligent performances reveals that what Stanley distinguishes as knowledge-how, -when and -where on the basis of the syntax and semantics of English *is already unified* in the situation-specificness of the exercise of assessment-competences.

§ 8.6 The Gambit of Detaching Skill

I have argued that Rylean responsibilism both accomodates and supersedes intellectualism, first, in that it fully accounts for the many specific knowledge states involved in knowing how to do something but, second, in that it grounds and explains these knowledge states in an assessment capacity, giving a better account of the gradability of know-how and of the way in which competent actors adjust to novel situations. As it happens, it seems as if the core moral of these arguments from § 8.5 have meanwhile been accepted by intellectualists, namely by Jason Stanley and Timothy Williamson.

This proposal involves the gambit of detaching skill from know-how. While I have used 'know-how', 'skill' and 'competence' synonymously over the course of Part One of this book, intellectualists have always drawn a sharp distinction between know-how on the one hand and the largely interchangeable concepts of skill and competence on the other hand (cf. Pavese 2016b). However, if skill and know-how are distinct concepts, one may try to repair the deficits of one's conception of know-how by proposing a suitable conception of the distinct concept of skill. This is precisely what Stanley & Williamson (2016) have now proposed.

First, Stanley & Williamson argue that skill requires precisely the kind of propositional knowledge they take to constitute know-how, as well as other such kinds of knowledge (cf. Stanley & Krakauer 2013). They write:

Skill is intimately connected to a kind of knowledge which philosophers have typically, though misleadingly, called “knowing how”. If one is skilled at chess, one knows how to employ one’s knight; if one is skilled at tennis, one knows how to serve. But the knowledge states connected to skill are not exclusively or even mainly expressed by the “how” construction. Skill at scoring in basketball requires knowing when to leap. Skill at driving to the basket requires knowing where to go when confronted with a defender. Skill at conversation requires knowing whether or not to interrupt. (Stanley & Williamson 2016, 2)

But what, then, is the connection between skills and these states of propositional knowledge? They write:

Our claim about skills is straightforward: it is that *skills are a kind of disposition to know*. More specifically, to be skilled at the action type of φ -ing is to be disposed to form knowledge appropriate for guiding tokens of φ -ing. So to be skilled at returning serve in tennis is to be disposed to have knowledge appropriate for guiding returns of serve, to be skilled at driving to the basket in basketball is to be disposed to have knowledge appropriate for guiding one’s movement to the basket, and so on. (Stanley & Williamson 2016, 3)

This development of the intellectualist view is very plausible, and obviously very close to the position I have advocated myself. For example, these considerations allow for an account of way in which competent actors adjust to novel situations which is nearly identical to the one I have discussed in § 4.4 and § 8.5 (cf. Stanley & Williamson 2016, 6–7; Stanley & Krakauer 2013). However, there is still a number of problems with this view. In the remainder of this section, I will discuss these issues, and argue that Stanley & Williamson’s newest proposal can hardly be understood as a defense against critics of intellectualism, but constitutes a substantial concession on their part which paves the way for a more heartfelt endorsement of a rival view such as the Rylean responsibilist account discussed here.

The first problem with the view that skills are dispositions to know is that it leads to a counterintuitive view of what it is to exercise a skill. As also discussed in § 2.1, the most intuitive view would be that the skill to, say, play squash manifests in performances of playing squash. The skill to A manifests itself by being exercised, i.e. in A-ing. Clearly, this is not what Stanley & Williamson are able to say. The relevant disposition to know is instead exercised by judgments about the relevant activity.

However, Stanley & Williamson want to have their cake and eat it, too. They introduce the following pair of distinct notions of the manifestation of skill, distinguished by indices:

Manifestation₁: A skill manifests₁ in knowledge states.

Manifestation₂: A skill manifests₂ in actions guided by knowledge states that are manifestations₁ of that skill.

More generally, a disposition to ϕ manifests₁ only in ϕ -ing. It may manifest₂ in all sorts of other ways. For example, the fragility of a vase may manifest₂, but not manifest₁, in the label ‘Handle with care’ on the crate in which it is being transported.

We shall say that an action *exhibits* a skill if and only if it is a manifestation₂ of that skill. That is, an action exhibits a skill if and only if it is guided by the knowledge states that are *direct* manifestations, i.e. manifestations₁, of that skill. We will also sometimes speak of the exercise or execution of a skill, by which we also mean the manifestation₂ of that skill. (Stanley & Williamson 2016, 5)

This distinction of different notions of manifestation is not intrinsically problematic, of course. There are indeed many different ideas and phenomena to be distinguished here, and I have offered a stratification of these myself (cf. e.g. § 2.1). However, I take it that Stanley & Williamson’s proposal is not very natural.

To see this, it is instructive to consider their own example concerning the fragility of a vase. The analogy with respect to skill is this. The fragility of a vase manifests in the *primary* sense in the breaking of the vase, and it manifests *derivatively* in its being transported in a crate labeled ‘Handle with care’. Analogously, the skill to play squash manifests in the *primary* sense in judgments about what to do when and how in playing squash, and it manifests *derivatively* in acts of playing squash. While Stanley & Williamson are clearly committed to this view, I have followed Ryle in arguing that the primary, indeed the sole manifestation of a skill is in performances of the very activity the skill is a skill to engage in (cf. § 2.1). This order of priority is exactly reversed here.

Further, I have proposed to distinguish not between two ways of how one and the same skill manifests itself, but between how a *skill as a whole* manifests itself and how a *proper part* of this skill manifests itself, namely the relevant state of understanding of the activity in question, which is itself an assessment skill (cf. § 4.1). Conceptually, however, there was no need to distinguish between different notions of manifestation. A competence to A manifests itself in A-ing. The relevant difference was a difference in what it is that manifests itself. Again, I contend that this proposal is more plausible than the distinction offered by Stanley & Williamson.

So far, however, my objections may be easily brushed off as mere matters of conceptual aesthetics. Maybe so. But there are deeper issues. One of them is directly addressed by Stanley & Williamson:

One might worry that because we make skill a disposition to acquire knowledge, we thereby put skill *before* knowledge in explaining intelligent action, and so vindicate anti-intellectualism. We agree with Ryle, for example, that for virtually any φ , skill at φ -ing is a *multi-track disposition* (Ryle, 1949, p. 44). For example, there may be no finite non-indexical specification of the disposition that is skill at driving to the basket in basketball. Nevertheless, even if there is no such specification, this does not mean that “skill comes before knowledge”, as Imogen Dickie (2012) has argued. Skill at φ -ing is a state whose nature is constituted through the knowledge relation. (Stanley & Williamson 2016, 9)

Unfortunately, I do not see how this answer can genuinely solve Stanley & Williamson’s problem. There is a very real sense in which the disposition to acquire the relevant knowledge states – Stanley & Williamson’s ‘skill’ – is prior to these states themselves. This is where they come from, this is how they are grounded and explained. As I have discussed at length in § 8.5, this clearly vindicates Dickie’s suggestion that skill comes before knowledge.

The only direct hint as to what Stanley & Williamson mean to reply is their insistence in the last sentence of the quoted paragraph: “Skill at φ -ing is a state whose nature is constituted through the knowledge relation.” (Stanley & Williamson 2016, 9) But what does it mean to be ‘constituted through the knowledge relation’? The only clear connection between Stanley & Williamson’s ‘skill’ and knowledge is that the former is a disposition to have the latter. But this does not mean that the former is ‘*constituted through*’ the latter. After all, the fragility of a vase is also not ‘constituted through’ events of its breaking. It is a potentiality for such events, but not itself such an event. Thus, skill as a disposition for knowledge clearly puts the potentiality to acquire knowledge prior to the knowledge itself. The core of the account is the disposition, not its manifestations.

This raises a further problem. What more can be said about this disposition? Stanley & Williamson commit themselves to very little. Instead, they even suggest that a further explanation may lead to problems:

The view in this paper is that skills are dispositions. The view is not well-expressed by saying that “skills are competences”, because “competence” is easily read as involving skill. In such a sense of competence, the view that skills are competences to know is in danger of a regress, since competences would involve skills. Thus in effect it says, unpromisingly, that skill in φ -ing involves skill in acquiring knowledge relevant to φ -ing. By the same token, skill in acquiring knowledge relevant to φ -ing involves skill in acquiring knowledge relevant to acquiring knowledge relevant

to φ -ing, and so on. By contrast, on our view, although skill in φ -ing *is* the disposition to have knowledge appropriate to guiding φ -ing, it is not in general *skill* in having or acquiring such knowledge. (Stanley & Williamson 2016, 4)

Thus, the only clear position taken by Stanley & Williamson is the negative one that the relevant dispositions are at least not always themselves skills or competences, for otherwise, this would lead to a vicious regress.

In this respect, the account of intellectual guidance established in chapter 4 is much more substantial. I have argued that the relevant dispositions are not mere dispositions, but genuine competences because they fulfill all the relevant criteria of intelligence (cf. § 4.1). This, I take it, constitute a *prima facie* problem for Stanley & Williamson. They must be able to show that the hallmark of competence – intelligent practice and normative guidance – is not or not always present in the case of the relevant dispositions to know. However, it is not clear if this is a viable option. After all, these dispositions are dispositions make full-blown propositional judgments when and where they are correct. Arguably, this is a simply a paradigm case of genuine competence.

This leaves the problem of regress mentioned by Stanley & Williamson. However, given the conceptual nature of the relevant capacities (cf. § 4.2), I have argued that thus problem can be solved if one appreciates the self-reflexivity of conceptual competences (cf. § 4.6 and § 4.7).

Thus, the gambit of detaching skill and know-how in the way proposed by Stanley & Williamson is a substantial rapprochement between intellectualism and its critics. As I have argued, however, it paves the way for abandoning intellectualism in favor of the Rylean responsibilist account I have offered. And while one may certainly draw a sharp distinction between know-how and skill, my considerations concernin the polysemy of the English expression ‘knows how to’ in § 7.5 may suggest a change in conceptual policy, as well.

However, this discussion has still stopped short of the heart of the problem of know-how – the project of explaining intelligent practice. Stanley & Williamson also address this topic in their paper on skill. They write:

The relation between the knowledge that manifests₁ a skill and the action that is done on the basis of that knowledge is *Guidance*:

Guidance: Any skilled action is guided by knowledge that manifests₁ possession of skill at that activity.

So, in a very clear sense, skilled action is action guided by propositional knowledge, the propositional knowledge that is revelatory of that agent’s skill. (Stanley & Williamson 2016, 6)

Bracketing the other issues I have discussed in this section, I am in broad agreement with this claim, as should be clear from my own account of guidance as responsible control in §4.4. However, the next and final chapter of this book will argue that there are systematic problems in the way in which intellectualism tries to account for such guidance by knowledge states. Despite my large agreement with intellectualism when it comes to the role of individual states of knowledge of ways of acting, and even in part with regard to their genesis and grounding, I will argue that there are serious problems for intellectualism in putting them to practice.

Chapter 9

Intellectualism in Practice

The hallmark of know-how is that it explains intelligent practice. I have spelled out this core criterion of explanatory adequacy in chapter 1, following Ryle's notion of intelligent practice. Proponents of intellectualism are also committed to this aim. Chapter 8 has presented intellectualism in detail and in various forms, identifying both common ground with Rylean responsibilism and points where my proposal may be preferable. But I have not mentioned the way in which intellectualism attempts to meet the crucial explanatory task in the debate about know-how – accounting for intelligent practice. This is the topic of the current, concluding chapter of this book.

In the first pair of sections, I look at what proponents of intellectualism have offered as positive explanations of how exactly know-how is related to intelligent practice. §9.1 discusses objectualist intellectualism and the idea of action-guiding states of understanding, a notion which evidently shares many features with my own account of intellectual guidance in chapter 4. Then, I discuss propositionalist intellectualism, with a particular focus on the view that certain states of propositional knowledge are practical in that they involve practical modes of presentation of ways of doing something (cf. §9.2). In both cases, I shall argue that intellectualism can only succeed in accounting for intelligent practice if they make a crucial step away from the intellectualist identification of know-how with states of propositional knowledge or objectual understanding.

In the remainder of this chapter, I discuss Ryle's famous regress argument against intellectualism, an objection which purports to establish the principal impossibility of explaining intelligent practice with the sole appeal to standing states of knowledge. §9.3 discusses the way in which Ryle presents his argument, or arguments, and identifies its core form. In §9.4, I defend Ryle's Regress against the objection that it conflates intentional and

automatic performances. Given my independent considerations from chapter 3, this will turn out to clarify and even strengthen Ryle's argument. § 9.5 will go on to discuss how intellectualists have responded to Ryle's Regress and maintain that it withstands these criticisms.

In the concluding § 9.6, I show how this regress problem can be solved, drawing on my own account of intellectual guidance from chapter 4, an account which also reserves an essential role for propositional knowledge. As I shall suggest, some of the ways in which intellectualists have reacted to the problem of Ryle's Regress already pave the way for intellectualism to transform itself into something closer to Rylean responsibilism.

§ 9.1 Action-Guiding Understanding

John Bengson and Mark Moffett are very explicit about the fact that the concept of know-how has its point in what they call "the philosophical theory of intelligence" (cf. Bengson & Moffett 2011b). And with their own objectualist intellectualist account of know-how, they claim to be able to explain "the practical import of know-how, that is, the role of know-how in the production and explanation of action" (Bengson & Moffett 2007, 53).

The first formulation of the envisaged explanation is this:

[A]ccording to our analysis, the concept of knowing how to φ is the concept of a non-ability entailing epistemic success state that could guide action. (Bengson & Moffett 2007, 53)

However, the fact that something *could* guide somebody's intelligent performances does not show that it is the explanans of such intelligent acts when these actually occur. It could guide intelligent action, but it remains an open question where and when it does so, indeed whether it does so at all. Bengson & Moffett are aware of this worry. They write:

Of course, understanding a way of performing an action does not invariably produce action in any particular individual. Nevertheless, it is clearly *apt* to do so. (Bengson & Moffett 2007, 53)

Equivalently, Bengson & Moffett say that the understanding is "poised" to guide successful, intelligent practice, and they spell out this notion of aptness or poisedness to guide as follows:

Knowledge how to φ is a state σ such that: if x is in σ , then it is possible for there to be some individual y such that y 's exercise of σ underlies and explains y 's successfully and intentionally φ -ing—that is, σ *guides* y in successfully, intentionally φ -ing. (Bengson & Moffett 2011c, 177)

This statement is what Bengson & Moffett call the “*action-guidingness* connection” (Bengson & Moffett 2011c, 177). But can it solve the problem of practical import?¹ In order to discuss this question, let me also quote two of the statements with which Bengson & Moffett clarify their view:

y’s exercise of σ must be the explainer (not simply an element in, or enabler of, a complete explanation) of *y*’s intentionally and successfully φ -ing. (Bengson & Moffett 2011c, 177 fn. 32)

[The *action-guidingness* connection] does not require a subject who knows how to φ to possess a power to φ ; it requires only that there be *some* subject who does. (Bengson & Moffett 2011c, 177 fn. 33)

Unfortunately, these two claims turn out to be inconsistent.

To see this, consider again the difference between the coach and his gymnastics student described in *Bela Karoli* on page 54. According to Bengson & Moffett, both Bela Karoli and his student know how to perform a standing layout on beam, but only the student possesses the ability to do so herself. However, if it is true that the student’s intelligent success is explained by her exercise of her know-how, and *only* by her exercise of her know-how, (which her coach possesses, too), then it follows immediately that her actual *power* or *ability* to engage in this activity does not explain anything at all. But this conclusion is absurd. It is *precisely* such an ability which explains the difference between these cases – an intelligent ability to rely on an understanding which offers not merely the counterfactual possibility to guide one’s performances, but which *does* in fact do so.

Of course, chapter 7 has firmly established that part of my disagreement with Bengson & Moffett is entirely verbal. What I distinguished as genuine competence and a mere understanding of an activity may both be ‘know-how’ on their account. But the present point is about the explanatory aims of these concepts. We want to understand the difference between those who can and do intelligently engage in an activity on the one hand and those who cannot and do not do so on the other hand, and we want to do so in terms of the concept of know-how. But then it cannot be true that our *complete* explanation fails to distinguish between these two at all. If so, we have simply failed to explain what we wanted to explain.

I have argued that, *pace* Bengson & Moffett’s above ‘*action-guidingness* connection’, know-how is not merely something which *could* guide intelligent acts if the actual power to perform them were added. Instead, I hold

¹ I shall bracket the fact that Bengson & Moffett focus too narrowly on intentional action here since I have discussed the crucial role of unintentional, entirely automatic performances in detail in chapter 3.

that know-how is something the exercise of which *already constitutes* an intelligent performance. Bengson & Moffett are correct to insist that this requires guidance by an understanding, but they are wrong to believe that the mere aptness of a state to play this role already suffices for know-how.

In contrast, chapter 4 has offered a holistic account of know-how as the *whole* of the capacity to engage in an activity on the basis of an understanding of what it takes to do well in that activity. Given an explication of this notion of understanding as an assessment capacity (cf. § 4.1), I have offered an account of what exactly it means to be guided by such an understanding, namely that one needs to exercise this assessment capacity and to act in the light of these assessments (cf. § 4.4). This crucial element of responsible control is what Bela Karoli lacks, but her student exhibits.

Despite these differences, there is a large common ground between Rylean responsibilism and Bengson & Moffett's objectualist intellectualism (cf. § 4.2 and § 8.4). Thus, I hope that my account will be seen as a way to preserve the insights of this view and to offer a solution for its core problem.

§ 9.2 Practical Modes of Presentation

Like Bengson & Moffett, Jason Stanley also explicitly endorses the criterion that know-how is what explains intelligent practice. Propositionalist intellectualism is therefore committed to the same explanatory aim. After what I already quoted from Stanley & Williamson (2016) on page 265, here are some examples of how Stanley expresses this commitment:

A particular action of catching a fly ball is a skilled action, rather than a reflex, because it is guided by knowledge, the knowledge of how to catch a fly ball. (Stanley 2011b, 130)

What makes an action an exercise of skill, rather than a mere reflex, is the fact that it is guided by the intellectual apprehension of truths. (Stanley 2011b, 174)

That someone skilled at an activity knows how to do that activity is as good a candidate as any to be a conceptual truth. It is therefore no surprise that everyone who discusses skilled action, from Ryle forwards, agrees that skilled action requires knowledge how. The debate has been about the nature of knowledge how. I have argued that skilled action is action guided by knowledge how, and that knowing how to do something amounts to knowing a fact. Skilled action is action guided by knowledge of facts. (Stanley 2011b, 175)

In this passage, Stanley correctly argues that skilled action must be explained in terms of know-how, but falsely equates know-how with the guiding propositional knowledge involved in skilled action. This is in keeping

with what I called the gambit of detaching skill in § 8.6. In contrast, chapter 4 has offered a holistic account of know-how as the *whole* of the capacity to engage in an activity while acting in the light of propositional knowledge. In this sense, I contend that Rylean responsibilism can also preserve the insights of propositionalist intellectualism, but avoid its problems.

Stanley is well aware of the fact that he needs to offer an explanation of the practical import of the propositional knowledge with which he identifies know-how. To do so, he amends the linguistic account discussed in § 8.1 with a number of further considerations which he introduces with the aid of a *prima facie* counterexample to his view.²

Suppose that Hannah doesn't know much about bicycles, and has certainly never ridden one. Susan points to John, and tells Hannah that John is riding his bicycle in a way in which Hannah could use to ride a bicycle. Since Hannah trusts Susan, it seems that Hannah thereby comes to know that that way is a way in which she could ride a bicycle, i.e. [(37a)] is true. Nevertheless, it seems that Hannah does not know how to ride a bicycle, i.e. [(37b)] is false: (Stanley 2011b, 126)

- (37) (a) Hannah knows that that way is a way in which she could ride a bicycle.
 (b) Hannah knows how to ride a bicycle.

Of course, Stanley is aware of the fact that, on his account as I have laid it out so far, (37a) *entails* (37b). After all, (37a) has the form stated in (L3) on page 254, namely 'For some way of A-ing w, w is how S can A'.³ And it was precisely the claim of propositionalist intellectualism that this is what constitutes genuine know-how. Stanley's problem then consists in the fact that this is clearly false in the present example. In order to solve this problem, he introduce a *further distinction* with respect to propositional knowledge of the form (L3). Thus, there are *two kinds* of Hannah's knowing, of a way w, that w is how she can ride a bicycle.

But how should this further distinction be explained? Stanley offers two different options here – an appeal to the kind of Fregean mode of presentation of the relevant way of doing something, and an appeal to the kind of modality involved in the infinitive in the ascription of know-how.

I shall begin with the second of these proposals where Stanley relies on the seminal work on modality by Angelika Kratzer (1977) and points out the context-sensitivity of the modal elements involved in these sentences:

² This counterexample was already discussed earlier (cf. Stanley & Williamson 2001, 429–430), but that treatment only includes one of the two answers Stanley offers later.

³ True, (37a) involves 'could' instead of 'can', but Stanley's problem remains the same if this is corrected. On this topic, see also footnote 4 on the following page.

Though the modals in [(37 a)] and [(37 b)] have the same force – they are kinds of dispositional, or ability modals – they are interpreted via distinct modal parameters. In [(37 a)], the modal parameter is one that takes the world of evaluation, and yields a set of propositions that characterize Hannah’s physical state after training for some time with a bicycle. In contrast, the natural modal parameter for the envisaged utterance of [(37 b)] is one that takes the world of evaluation, and yields a set of propositions that characterize Hannah’s physical state at the moment. That is why the two utterances express different propositions – because the modals in the two sentences are interpreted via distinct modal parameters. (Stanley 2011b, 126)

This is a very straightforward explanation of the intuitive difference between these two sentences and of the way in which it is linguistically represented. But as Ephraim Glick has already pointed out in detail, Stanley fails to make clear *why* it is that the modal expressions in (37 a) and (37 b) are supposed to be interpreted differently (cf. Glick 2013, 19–20). In particular, the expression ‘could’ in (37 a) which suggests a more distant possibility than (37 b) is simply a red herring. For the same contrast recurs when (37 a) is formulated with ‘can’ or even in a way which involves the *very same* modal expressions as (37 b).⁴ Thus, what Stanley offers here is not an account of the relevant difference, but merely a restatement of what we wanted to have an explanation for.

I would also like to note that, even if this can somehow be made to work, it could solve Stanley’s problem only by giving away a core element of intellectualism. For according to the account on offer, (37 b) attributes to Hannah an epistemic state which entails that she has the actual ability ride a bicycle given her actual “physical state at the moment” (Stanley 2011b, 126). After all, Hannah knows that she is able to ride a bicycle in a certain way, which requires the truth of the proposition that she can do so, and thereby the actual existence of Hannah’s ability to ride a bicycle.⁵ But

⁴ Such a version of (37 a) is “Hannah knows that that is how to ride a bicycle” (Glick 2013, 20), which differs from (37 b) only in that it inserts ‘that that is’ between ‘knows’ and ‘how’. And as a matter of fact, the version of (37 a) in Stanley’s first statement of this problem also avoided the misleading expression ‘could’: “Hannah knows that that way is a way for her to ride a bicycle.” (Stanley & Williamson 2001, 429)

⁵ In this vein, Stanley has also argued that there is still a “difference between explicit ability modals and ascriptions of knowing how” in that “ascriptions of knowing how tolerate cases in which there is success only in more distant situations.” (Stanley 2011b, 127) This, he claims, allows him to explain cases like *Ski Instructor* on page 150 where know-how does not require actual ability (cf. § 5.1 and § 5.2), but merely some form of counterfactual or more distant possibility of success. But as Ephraim Glick has already pointed out (cf. Glick 2013, 20), such differences in the ‘distance’ of the relevant possible worlds also occur with *direct* ability ascriptions as with ‘can’ or ‘is able to’. Thus, Stanley has failed to establish that ‘knows how to’ does not involve genuine ability or even some form of less ‘close’ ability.

then, the aim of explaining intelligent practice with appeal to know-how would only be successful because something else is introduced, an *actual ability* to ride a bicycle, which helps to explain what it means to possess the kind of propositional knowledge to be identified with know-how.

But maybe there is another way to explain the difference between (37 a) and (37 b). In Stanley's joint paper with Timothy Williamson (cf. Stanley & Williamson 2001, 430–431), he already makes the following suggestion. While (37 a) ascribes to Hannah knowledge of a way of riding a bicycle under a *purely demonstrative* mode of presentation, (37 b) requires that Hannah has knowledge of this proposition and conceives of this way of riding a bicycle under a *practical* mode of presentation. Thus, the full account of the semantics of 'knows how to' in the reading (L3) requires that the relevant ways of acting are entertained under such practical modes of presentation.

But what are practical modes of presentation? I agree that it is entirely in line with established semantic theories of modes of presentation in general that there should also be a practical guise of entertaining a proposition and that these are just as unproblematic as modes of presentation in general,⁶ for example because of analogies with indexical modes of presentation.⁷ But it still constitutes a problem for Stanley's view that it remains unclear what exactly such practical modes of presentation are other than that they are stipulated to make the difference between (37 a) and (37 b). The only positive thing which is explicitly mentioned about practical modes of presentation is that they are connected with dispositions:

Thinking of a person as oneself entails being disposed to behave in certain ways, or form certain beliefs, given relevant input from that person. Similarly, thinking of a place as *here* entails being disposed to behave in certain ways, or form certain beliefs, given relevant input from that place. Analogously, thinking of a way under a practical mode of presentation undoubtedly entails the possession of certain complex dispositions. It is for this reason that there are intricate connections between knowing-how and dispositional states. (Stanley & Williamson 2001, 429)

But unfortunately, nothing further is said about the nature of these dispositions, leaving the account nearly as unclear as before (cf. Schiffer 2002; Schröder 2013; Glick 2013). How can this be spelled out more fruitfully?

⁶ For further discussion of this, see Stanley & Williamson (2001, 428–429), Stanley (2011b, 123–125), Stanley (2011c, 210–212), Stalnaker (2012, 758–761), Stanley (2012b, 774–778) and Pavese (2015b).

⁷ Alva Noë has objected that Stanley & Williamson's analogy with indexical modes of presentation is "plainly circular" because we have "no independent reason to believe that the complement clauses in [(37 a)] and [(37 b)] express the same proposition" (Noë 2005, 288). But as Löwenstein (2011, 281–282) and Glick (2013, 3–4) have noted, this is mistaken. The semantic theory discussed in § 8.1, albeit controversial, constitutes such an independent reason.

One way to go would be in analogy to an idea I have just discussed in § 9.1, namely what Bengson & Moffett have called the ‘aptness’ or ‘poisedness’ of a piece of knowledge to guide practical conduct. In fact, they even suggest that “there is no need to invoke practical modes of presentation, since an independently motivated condition (understanding) for know-how provides an adequate solution.” (Bengson & Moffett 2007, 49 fn. 32) But this option only leads back to my earlier worries. On this interpretation, a practical mode of presentation merely makes it the case that there is a counterfactual possibility for one to intelligently engage in the activity in question, provided one also develops the actual ability to do so. But it does not explain the difference between this mere counterfactual possibility and actual intelligent conduct.

This suggests that practical modes of presentation must go hand in hand with having the actual ability to instantiate this way of acting. And indeed, this has already been noted by many commentators on this proposal (cf. Koethe 2002; Rosefeldt 2004; Jung & Newen 2010).

In response, Stanley may appeal to cases such as *Amputee Pianist* on page 155. Here, it is *prima facie* plausible to assume that, such a pianist retains a practical mode of presentation of the ways in which she used to play the piano, at least for a while. But there are three reasons why this is of no help. First, it does not solve the problem what *constitutes* this practical mode of presentation in the absence of actual competence. Second, my discussion in § 5.2 suggests that there may still be a sense in which we assume there to be an ability, and this may precisely be what explains the plausibility of assuming that the practical mode of presentation is retained, as well. And finally, whatever such practical modes of presentations are, this proposal would concede that all the support for their existence would *still* be derived from the fact that the master pianist *used* to have the actual competence to play the piano. The conceptual dependence of practical modes of presentation on ability remains.

A second response, as Stanley explicitly considers (cf. Stanley 2011b, 129), appeals to a different kind of case, Hawley’s *Impressive Skating* on page 179. Here, Hawley’s *alter ego* has the right mode of presentation of the right way to impress the kids at the skate park, but only because she has the full-blown know-how to perform certain skateboard tricks. But she does not know that performing these tricks would impress the kids. This is an example of practical luck rather than genuine know-how to impress the kids (cf. § 5.4), one which involves opacity (cf. § 6.1). But this clearly fails to show that practical modes of presentation are possible without ability. At most, it shows that they are possible without the *corresponding* know-how.

Thus, the result of the appeal to practical modes of presentation seems to be the same as the with the appeal to differences in the kind of modality involved in ascriptions of know-how. In both cases, the explanation had to assume the presence of the relevant ability. And again, this would lead to the problematic result that the aim of explaining intelligent practice with appeal to know-how would only be successful because something else is introduced, an *actual ability*, which helps to explain what it means to possess the kind of propositional knowledge to be identified with know-how.

This is also the upshot of the most recent and most elaborate defense of practical modes of presentation by Carlotta Pavese. After making a compelling case for the view that propositional attitudes are crucial for know-how (cf. Pavese 2015a), she argues that their crucial connection to practice indeed consists in Fregean practical senses (cf. Pavese 2015b). This notion can, according to Pavese, be made rigorously precise and theoretically fruitful on the basis of an analogy with computer programs:

[W]e find examples of practical senses in the semantic values assigned to programs by operational semantics for programming languages. As I argue, such operational semantic values are naturally construed as modes of presentation for ways to execute tasks — modes of presentation that exist independently of thinkers and that determine referents. As such, these operational semantic values qualify as Fregean senses. (Pavese 2015b, 2)

I cannot discuss this promising proposal here in detail. But I would like to point out that it also assumes actual abilities in the explanation of practical modes of presentation. To see this, consider what Pavese writes about the possibility that a certain program with certain operational semantic values, or “OSV”s, for short, may still not be executable:

Of course, there may be programs that are not *compatible* with a particular system. Those are programs that the system simply cannot “access” in the sense that they are the operational meanings of program texts that the interpreter cannot interpret. [...] But if the interpreter comes to understand the OSV of a program text, then the system must thereby be endowed with a certain set of abilities and that is so in virtue of what OSVs *are*. (Pavese 2015b, 12)

Thus, the relevant computer system must already have certain *basic* abilities – to add, for example – in order to allow for a program to endow it with some further ability – say, to multiply – as well as the *general* capacity to interpret these kinds of operational semantic values. Thus, “by appealing to practical senses, we can explain how the ability to perform a complex task arises from the ability to perform its parts.” (Pavese 2015b, 16)

It may be entirely correct that this proposal solves a number of important problems. But, applied to the original explanatory projects in the debate about know-how, the consequences seem to remain the same as before. If defenders of propositionalist intellectualism appeal to practical senses, then the project of explaining intelligent practice with appeal to know-how would only be successful because something else is introduced, *actual abilities*, which help to explain what it means to possess the kind of propositional knowledge to be identified with know-how.

The core problem with all of these considerations about practical modes of presentation has been identified most clearly by Ephraim Glick (2013):⁸

[T]he opponent [...] will grant that there are “action-based” thoughts that one is able to grasp iff one knows how to do a certain action. But the opponent might simply maintain that acquiring know-how results in the accessibility of those new thoughts, rather than vice versa. If there is an argument for the claim that the order of explanation goes the way the Intellectualist wants, I don’t see it in the literature. (Glick 2013, 6)

Thus, practical modes of presentation may indeed exist, and they may indeed be as important as Pavese (2015b) has stressed. But the way in which these should be explained will crucially involve an appeal to genuine ability and practice.

In fact, I take it that Rylean responsibilism is precisely an account of know-how which shows what it would mean for propositions to be essentially connected to practical guidance. Such propositional knowledge must be *organically integrated in the whole of a competence* – both as the *result* of a capacity to assess one’s options and performances, and as the *material* of the capacity to act in the light of such assessments. Such an organic integration is perfectly suited to account for what it means for the relevant assessments to involve practical modes of presentation.

§ 9.3 Ryle’s Regress

I have argued that the positive intellectualist proposals on offer have failed to explain the practical import of know-how. In this section, I will argue that there is a further reason why intellectualism cannot account for the fact that know-how explains intelligent practice. This can be established by the most famous and most unclear argument in the debate about know-how, Ryle’s Regress. This is the topic of the remainder of this chapter.

⁸ Pavese responds to this paper by Glick, but only with respect to a different problem, leaving the present objection untouched (Pavese 2015b, 19).

I shall begin my discussion of this argument by quoting Ryle's own presentations of it in detail before I propose an explicitly reconstructed version of this argument in its strongest form. In § 9.5, I will then go on to discuss how Ryle's Regress can be avoided, and argue that this requires abandoning intellectualism.⁹

Ryle presents his regress objections in a number of guises, and in various places. The two clearest statements of his argument are the following:

I argue that the prevailing doctrine leads to vicious regresses, and these in two directions. (1) If the intelligence exhibited in any act, practical or theoretical, is to be credited to the occurrence of some ulterior act of intelligently considering regulative propositions, no intelligent act, practical or otherwise, could ever begin. [...] (2) If a deed, to be intelligent, has to be guided by the consideration of a regulative proposition, the gap between that consideration and the practical application of the regulation has to be bridged by some go-between process which cannot by the pre-supposed definition itself be an exercise of intelligence and cannot, by definition, be the resultant deed. [...] Consistency requires, therefore, that this schizophrenic broker must again be subdivided into one bit which contemplates but does not execute, one which executes but does not contemplate and a third which reconciles these irreconcilables. And so on for ever. (Ryle 1945a, 2–3)

The crucial objection to the intellectualist legend is this. The consideration of propositions is itself an operation the execution of which can be more or less intelligent, less or more stupid. But if, for any operation to be intelligently executed, a prior theoretical operation had first to be performed and performed intelligently, it would be a logical impossibility for anyone ever to break into the circle. Let us consider some salient points at which this regress would arise. According to the legend, whenever an agent does anything intelligently, his act is preceded and steered by another internal act of considering a regulative proposition. But what makes him consider the one maxim which is appropriate to his practical problem rather than any of the thousands which are not? [...] Intelligently reflecting how to act is, among other things, considering what is pertinent and disregarding what is inappropriate. Must we then say that for the hero's reflections how to act to be intelligent he must first reflect how best to reflect how to act? The endlessness of this implied regress shows that the application of the criterion of appropriateness does not entail the occurrence of a process of considering this criterion. Next, [...] how am I led to make a suitable application of the reason to the particular

⁹ The literature on Ryle's Regress is at least as extensive as the literature on know-how since this argument is also important for general questions of the nature of mental states beyond know-how. I will only refer to those parts of the literature which are particularly important for the problem of know-how, and even in this realm, I will have to leave a number of pertinent interpretations aside because it would occupy too much space to disentangle precisely where my view differs or overlaps with them (cf. Parry 1980; Koethe 2002; Schiffer 2002; Noë 2005; Hetherington 2006; J. Williams 2008; Damschen 2009; Fantl 2011; Tsai 2011a; Weatherson 2016).

situation which my action is to meet? For the reason, or maxim, is inevitably a proposition of some generality. It cannot embody specifications to fit every detail of the particular state of affairs. Clearly, once more, I must be sensible and not stupid, and this good sense cannot itself be a product of the intellectual acknowledgements of any general principle. [...] Knowing how to apply maxims cannot be reduced to, or derived from, the acceptance of those or any other maxims. (Ryle 1949, 31)

I take it that the general idea of Ryle's argument is very clear. He begins with the premise that know-how is what explains intelligent performances in the sense that it explains the very *intelligence* of them. However, if propositional knowledge is supposed to perform this explanatory role, then there are two open questions of how this should be possible.

First, how is the right piece of propositional knowledge selected? Since this must also be an intelligent act, its performance must also be explained in terms of know-how, and it would therefore require selecting the right piece of propositional knowledge *about* how to select the first one, and so on.

Second, when selected, how is the right piece of propositional knowledge practically applied? Since this must also be an intelligent act, its performance must also be explained in terms of know-how, and it would therefore require applying the right piece of propositional knowledge *about* how to apply the first one, and so on.

However, Ryle's argument does not only apply to propositionalist intellectualism, but to *any* form of intellectualism. The very same considerations can also be used to undermine the view that know-how consists in conceptions of ways of acting – objectualist intellectualism. The difference between the propositional knowledge that some way *w* is a way of doing something and the objectual knowledge or understanding of such a way is negligible in this context. To encompass both of these views, I shall henceforth employ the neutral notion 'knowledge of *w*', which is intended to include both propositional knowledge and an understanding of *w*. Further, Ryle's Regress also applies to the revisionary versions of intellectualism mentioned at the end of § 6.3 because it only relies on the point that know-how is identified with a standing epistemic state of believing or otherwise relying on a proposition or of conceiving of something in a certain way, independently of the question whether or not, and in which sense, this epistemic state qualifies as knowledge or as a correct conception. But I will leave this implicit.

I shall now present the way in which I propose to understand this famous two-fold regress argument in detail, perhaps with excessive technical rigor. This version of Ryle's Regress strongly relies, but still improves on earlier presentations of this argument (cf. Löwenstein 2011a, 291–293; Löwenstein

2013, 366–368; Fridland 2013, 886–887, 889–890), among other things because it extends the argument to objectualist intellectualism, and because it gives sufficient weight to the problem of selecting the right way, thereby allowing for a defense of the argument in the form of Ryle's original two-fold rather than only one-fold regress. But I will omit more nuanced references to this earlier work for the sake of better readability.¹⁰

(M) *Ryle's Regress*

- (M 1) If S intelligently ϕ s, then S' ϕ -ing is intelligent because S knows how to ϕ .
- (M 2) Intellectualism: S' knowing how to ϕ consists in S' conceptions or propositional knowledge of ways of ϕ -ing $w_1(\phi)$ – $w_n(\phi)$. *assumption*
- (M 3) *Thus*: If S intelligently ϕ s, then S' ϕ -ing is intelligent because S has conceptions or propositional knowledge of ways of ϕ -ing $w_1(\phi)$ – $w_n(\phi)$. *from (M 1-2)*
- (M 4) If S' ϕ -ing is intelligent because S has conceptions or propositional knowledge of ways of ϕ -ing $w_1(\phi)$ – $w_n(\phi)$, then S' ϕ -ing is intelligent because
- (a) S intelligently $\sigma(\phi)$ s, i.e. selects $w_i(\phi)$ from $w_1(\phi)$ – $w_n(\phi)$, and
 - (b) S intelligently $\alpha(\phi)$ s, i.e. applies $w_i(\phi)$ in her ϕ ing.
- (M 5) *Thus*: If S intelligently ϕ s, then S' ϕ -ing is intelligent because
- (a) S intelligently selects $w_i(\phi)$ from $w_1(\phi)$ – $w_n(\phi)$, and intelligently selects $w_j(\sigma(\phi))$ from $w_1(\sigma(\phi))$ – $w_m(\sigma(\phi))$, and intelligently selects $w_k(\sigma(\sigma(\phi)))$ from $w_1(\sigma(\sigma(\phi)))$ – $w_o(\sigma(\sigma(\phi)))$, ...
 - (b) S intelligently applies $w_i(\phi)$ in her ϕ ing, and intelligently applies $w_j(\alpha(\phi))$ in her applying of $w_i(\phi)$, and intelligently applies $w_k(\alpha(\alpha(\phi)))$ in her applying of $w_j(\alpha(\phi))$, ... *from (M 3-4)*
- (M 6) It is false that, if S intelligently ϕ s, then S' ϕ -ing is intelligent because S intelligently performs infinitely many further acts.
- (M 7) *Thus*: (M 2) is false, i.e. intellectualism is false. *from (M 5-6)*

Evidently, the crux of this argument lies with premise (M 4). And as far as I can see, the debate about Ryle's Regress is a debate exclusively about the plausibility of this premise, while the other premises and inferences are explicitly or implicitly accepted, or at least unchallenged.

¹⁰ This reconstruction bears on a number of interesting general questions about the nature and form of infinite regress arguments (cf. e.g. Sanford 1984; Nolan 2001; Wieland 2012; Wieland 2013). I have recently proposed a uniform account of regress problems elsewhere (cf. Löwenstein 2016) which also sheds light on the present case. But I shall maintain a different way of presenting the argument here for expository reasons.

A natural intuitive motivation for something along the lines of (M4) is that possessing knowledge of a way of doing something is always one thing, while acting or even being able to act on such knowledge of a way of acting is quite another thing. In Ryle's words:

It is of first-rate importance to notice from the start that stupidity is not the same thing, or the same sort of thing, as ignorance. There is no incompatibility between being well-informed and being silly. (Ryle 1949, 26)

Ryle exemplifies this point with a number of examples of people who possess lots of propositional knowledge about an activity, but nevertheless fail to intelligently engage in it – paradigmatically, with a chess player and a logic student who possess lots of knowledge about how to play chess and draw inferences, but nevertheless fail at these tasks (cf. Ryle 1945a, 5–7).

As I stated it, premise (M4) involves two requirements which I shall call the *selection requirement* in (M4a) and the *application requirement* in (M4b). The first of these leads to a regress in selecting a suitable way, as stated in (M5a), and the second to a regress in applying this way, as stated in (M5b). However, both (M5a) and (M5b) are *independently* sufficient to derive the conclusion (M7) on the basis of the final premise (M6) which excludes infinite regresses in the explanation of intelligent practice. Thus, to defend Ryle's Regress, it will be sufficient to show that *at least part* of premise (M4) is true, i.e. that at least *one* of the two requirements in (M4a) and (M4b) holds. In what follows, I shall defend both of these requirements, but even if one of these defenses should be unsuccessful, I can resort to the other one.

Before doing so, I would like to highlight that a large part of the debate about Ryle's Regress construes the selection requirement in a rather peculiar way. Since this point spells out what Ryle refers to as 'considering propositions', and sometimes misleadingly as 'contemplating propositions', several commentators have interpreted this idea as the requirement that one needs to cognitively formulate, and in this sense 'contemplate', the single relevant proposition which is later applied (cf. Stanley & Williamson 2001, 415). And there are indeed passages where Ryle seems to be proposing such a notion of contemplation (cf. Ryle 1945a, 2–4). Now, to the extent that such a requirement is intelligible in the first place, I agree with Ryle's critics that this requirement does not hold at all. But as Ryle's own presentation of his argument in the above quotation should make clear, the core point here is about *selecting* the right way of acting.

§ 9.4 Selecting and Applying

I have argued that the core of Ryle's Regress lies in (M 4) on page 279 – the double requirement of selecting the right way of doing something from one's knowledge of several such ways, and of applying this knowledge in practice. In this section, I shall spell out more clearly what this point involves, and how it should be understood.

As a matter of fact, there are many objections against Ryle's Regress which target precisely these two requirements – the selection requirement and the application requirement. This is because these objections are concerned with the notion of doing something *intelligently*, which plays a crucial role in both of these requirements, and even in the first premise of the argument, (M 1). In this respect, a standard criticism is this:

[F]or premise [(M 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)

And if doing something intelligently is doing so *intentionally*, then there is a clear problem for premise (M 4), which is nicely brought out by a much-cited passage by Carl Ginet and by the standard conclusion from these considerations in a quotation from Jason Stanley.

Ginet's Door (Ginet 1975, 7)

I exercise (or manifest) my knowledge *that* one can get the door open by turning the knob and pushing it (as well as my knowledge *that* there is a door there) by performing that operation quite automatically as I leave the room; and I may do this, of course, without formulating (in my mind or out loud) that proposition or any other relevant proposition.

The reasonable Intellectualist about intelligent action will hold that an action is intelligent in virtue of being guided by propositional knowledge, but deny that this entails that intelligent action requires a prior act of self-avowing the propositional knowledge that guides one's actions. (Stanley 2011b, 14)

As the growing literature on Ryle's Regress has already shown in detail,¹¹ Ryle's critics can easily, and correctly, dismiss both the idea that selecting the right way of acting is an intentional action, and that applying a way of acting is an intentional action. Instead, these things can perfectly well

¹¹ For discussion of such versions of Ryle's Regress, see Löwenstein (2011a, 288–291) and Cath (2013, 263–265, 372–373, 376). Alva Noë has suggested that these performances may not be conscious, but nevertheless intentional (cf. Noë 2005, 282). While I find this idea rather implausible (cf. Cath 2013, 265–266), it seems to me that Noë's remark still gestures at the very notion of intelligent performances I defend as well.

happen entirely automatically, and this immediately halts the regress. Thus, if intelligent performances are intentional actions, then both the selection and the application requirement fail and premise (M 4) is false.

Ryle's own way of presenting the regress certainly invites this reading, probably partly because his argument is clouded by the somewhat excessive polemic of which he is well aware (cf. Ryle 1949, 10–11). However, the core of Ryle's argument can be freed from the problem of intentional action.¹² I contend that the fact that the relevant performances of selecting and applying are automatic is a red herring when it comes to the question of intelligence and know-how. In § 3.1 and § 3.4, I have offered extensive arguments for the view that entirely automatic performances can also count as genuine exercises of know-how – arguments which are entirely independent from the problem of Ryle's Regress. Thus, premise (M 4) cannot be rejected on grounds of the obvious phenomenological fact that things like considering, selecting or applying propositions are often conducted automatically and unreflectively.

These considerations can be spelled out more fully in tandem with my groundwork about the concept of intelligence from chapters 1 and 2. As I have argued, intelligent performances are performances which are answerable to norms of good conduct, things which one may perform well or badly, better or worse, etc. (cf. § 1.1). What is more, they meet these standards because they are guided by an understanding of what it takes to do well in this activity (cf. § 1.5). Then, it is an unequivocal phenomenological fact that not every such intelligent performance is intentional, and that these can also occur entirely automatically (cf. § 3.1). But even such completely automatic performances may nevertheless be intelligent in virtue of the fact that they are the actualizations of routines and automatisms which have been, which are, and which continue to be shaped in the light of an understanding of what they ought to be (cf. § 3.4). This requires that it is *possible* to intentionally do what one sometimes does automatically, yet intelligently, because this possibility underwrites the idea that one can *practice* one's competences in order to improve them (cf. § 2.4 and § 3.3). I contend that all of these characterizations of the notion of doing something intelligently are also true in the case of the activities of selecting the right way of doing something, and of applying this way in practice, thereby positively supporting the selection requirement and the application requirement in premise (M 4) of Ryle's Regress.

¹² This crucial point is often not brought out explicitly enough even in otherwise brilliant defenses of Ryle's Regress (cf. e.g. Hornsby 2005; Wiggins 2009; Wiggins 2012). For an exemplary clarification of these matters, however, see Hornsby (2011, 94 fn. 15).

On the one hand, when confronted with a number of ways of doing something and with the task of choosing and practically applying one such way, one clearly ought to choose the best way available, or take the best option at hand. But for one's performance to be genuinely intelligent, it is not enough that one is merely reliable in selecting the best option. Instead, one must choose the best option because one understands that this is what one ought to do. True, all of this selecting may happen entirely automatically, and it even does so very often. But these automatic routines of choosing the best option in the situation at hand are intentionally practiced, and practiced partly by intentionally reflecting about one's options and choosing the best one. And even if such a selection is automatic most of the time, one can certainly also select intentionally. One way to see this is in the light of the fact that one may even intentionally choose an option *below* the optimal one – say, because one wants the opponent to win the game or because one wants to demonstrate how *not* to perform something to a student.

On the other hand, when a certain way of acting has been chosen as to be applied, i.e. when an option has been singled out as to be pursued, then it is still an open question if one manages to apply this way of acting in one's actual performance. Again, there is a clear sense in which one may do better or worse at this task, and there is also a clear sense in which performing intelligently requires that one is not merely reliable at doing so, but that one further manages to do well at applying this way of acting because one understands that this is what one ought to do. And again, all of this applying can perfectly well happen entirely automatically, as it often does. But such automatic routines of applying a way of acting chosen as the option to be pursued require intentional practice. As with the automaticity of selecting, one can also intentionally intervene in the automaticity of applying a way of acting, and it is even possible to intentionally apply a way of acting *worse* than one otherwise could – say, because one wants to save one's energy for later tasks or because one wants to fool somebody about one's true capabilities.

As discussed in § 3.1, Ryle certainly struggled with the idea of intelligent routine and automatic know-how. But in the context of his regress objection, he already anticipated this point at least to the extent that he discussed the idea of the *implicit* selection and application of knowledge:

There is a not unfashionable shuffle which tries to circumvent these considerations by saying that the intelligent reasoner who has not been taught logic knows the logicians' formulæ "implicitly" but not "explicitly"; or that the ordinary virtuous person has "implicit" but not "explicit" knowledge of the rules of right conduct; the skilful but untheoretical chess-player "implicitly" acknowledges a lot of strategic

and tactical maxims, though he never formulates them and might not recognise them if they were imparted to him by some Clausewitz of the game. This shuffle assumes that knowledge-how must be reducible to knowledge-that, while conceding that no operations of acknowledging-that need be actually found occurring. It fails to explain how, even if such acknowledgements did occur, their maker might still be a fool in his performance. All this intellectualist legend must be rejected, not merely because it tells psychological myths but because the myths are not of the right type to account for the facts which they are invented to explain. However many strata of knowledge-that are postulated, the same crux always recurs that a fool might have all that knowledge without knowing how to perform [...]. (Ryle 1945a, 7–8)

My defense of premise (M4) evidently does not involve a presupposition Ryle makes in this passage, namely that the knowledge of the relevant ways of acting must be knowledge of ‘general maxims’ or ‘regulative propositions’. It seems clear that Ryle envisaged an intellectualist opponent who equates know-how with propositional knowledge of fully descriptive and rather general rules and maxims. Critics of Ryle’s Regress are therefore absolutely correct in complaining that such an intellectualist position is implausible on independent grounds. As discussed in §8.4, it is a shared contention between intellectualism and my Rylean responsibilist proposal that, first, knowledge or a conception of the relevant ways of acting may also be demonstrative, and that it, second, cannot only take the form of *general* maxims and regulations, but also capture the fine-grained specificities of an actor’s performances and of a given situation in which she finds herself.

But Ryle’s Regress remains fully intact despite these corrections, and this in both directions of the regress. This has been spelled out most clearly by Ellen Fridland (cf. Fridland 2013, 886–890).¹³ The more coarse-grained one’s individuation of these ways, the more pressing is the application requirement and the application regress. And the more fine-grained this individuation is, the more pressing becomes the selection requirement and the selection regress. As Ryle said, “[h]owever many strata of knowledge-that are postulated, the same crux always recurs” (Ryle 1945a, 8).

§ 9.5 Avoiding Ryle’s Regress

In §9.3 and §9.4, I have spelled out the strongest version of Ryle’s Regress and distinguished it from less plausible cognate ideas which are also present in Ryle’s presentation of this argument. As I have shown, Ryle’s Regress

¹³ For a recent critical discussion of this argument, see Clarke (2016).

poses a severe challenge for intellectualism, and indeed for any view which holds that standing epistemic states like propositional or objectual knowledge are what explains intelligent practice. In this section, I will consider how intellectualists have responded to this argument, but argue that this problem remains unsolved.

As far as I can see, there are two ways in which intellectualists have responded to this problem. Both are prominently defended by Jason Stanley.

The first of these replies consists in two ideas. First, ways of acting are individuated in a very fine-grained way, and second, both selecting and applying a suitable way are claimed to be entirely automatic performances which are *therefore* not intelligent (cf. Bengson & Moffett 2011b, 26). Stanley expresses this point very succinctly in response to Josefa Toribio (2008):

The expert golfer knows *many* propositions of the form 'w is a way to get the ball to the green'. [...] Her expertise consists not just in the possession of this large body of propositional knowledge, but also in the fact that the automatic mechanisms responsible for applying standing epistemic states of an agent are well-aligned to her propositional knowledge about golf. (Stanley 2011b, 185)

As Stanley suggests, all that such automatic mechanisms actually do simply comes down to the triggering of a suitable mental representation:

Triggering a representation can certainly be done *poorly* or *well*. But this does not show that it can be done *intelligently* or *stupidly*. [...] Since triggering a representation is something we do automatically, [...] [premise (M4)] results in a manifest implausibility. (Stanley 2011b, 16)

As I have shown in § 9.4, the remark that the relevant acts are automatic is of no help here. Not only is Stanley's inference from the presence of automaticity to the absence of intelligence a straightforward fallacy – in fact, a fallacy which is ironically akin to what I called the 'phenomenological fallacy' committed by anti-intellectualists such as Hubert Dreyfus (cf. § 3.6). Instead, I have presented clear *positive* reasons to understand selection and application as intelligent in precisely the sense required in premise (M4).

However, there is more to be said here on Stanley's behalf. This can be found in what I have omitted in the above quotation, where Stanley writes:

In the vocabulary of Fodor (1983), triggering representations is something done by an *input system* rather than by a *central system*, by a module rather than by a central processor. (Stanley 2011b, 16)

Thus, Stanley conceives of selecting and applying ways of acting as something which is performed not by a person, but by a *sub-personal module*.

In his most recent defense of this view, co-authored with John Krakauer, Stanley explains such mechanisms in terms of the notion of ‘motor acuity’, a concept analogous to what cognitive scientists have called ‘perceptual acuity’ (cf. Stanley & Krakauer 2013). They offer an argument for the view that these mechanisms of selecting and applying knowledge are not intelligent, which would show that premise (M4) in Ryle’s Regress is false:

It is implausible to think of perceptual discrimination as a kind of action at all, much less an intentional action. The fact that capacities such as perceptual acuity do not characteristically manifest themselves as intentional actions explains why it is incorrect to think of such capacities as skills. (Stanley & Krakauer 2013, 6)

And they hold the same to be true about what they call ‘motor acuity’:¹⁴

Motor skill tasks have an acuity component that is directly analogous to perceptual acuity. (Stanley & Krakauer 2013, 9)

Thus, their argument is that the selecting and applying of individual ways of acting is not intelligent because it is never intentional and takes place entirely entirely on the sub-personal level.¹⁵

There are several principal problems and controversies concerning the sub-personal level of explanation which I find persuasive, but which I cannot discuss here.¹⁶ But there are even more straightforward problems with Stanley’s idea, some of which have already been pointed out in detail by Ellen Fridland (cf. Fridland 2014b, 9–12).

My most important worry about this proposal is that, on the face of it, it is simply phenomenologically absurd to hold that competent actors *never* intentionally choose a way of acting and that they *never* intentionally apply it in practice. As I have explained in §9.3, this is perfectly possible, and this is part of what explains the fact that competent actors can make

¹⁴ For my present purposes, I shall bracket the obvious problem that not all competences are motor skills and that this kind of argument therefore applies to part of the general topic of know-how at best (cf. *Ryle’s Range of Cases* on page 14). Instead, I will grant that an analogous argument may also be possible in the case of other competences.

¹⁵ In this context, some critics of Stanley’s view have argued that propositionalist intellectualism relies on further more specifically Fodorian views, including his famous account of the so-called ‘Language of Thought’ (cf. Roth & Cummins 2011). However, Stanley has correctly argued that his view is sufficiently independent from the question how a person *represents* propositional knowledge (cf. Stanley 2011a). For a related discussion, see Bartels & May (2015a; 2015b) and Glauer (2015).

¹⁶ I am sympathetic with objections along the lines of what Anthony Kenny calls the ‘hommunculus fallacy’ (cf. e.g. Wittgenstein 1953, §281; Kenny 1991; Keil 2003; Hornsby 2000; Tanney 2011). Zoe Drayson has offered an insightful general discussion of the distinction between the personal and the sub-personal level of explanation which is rather critical of these objections (cf. Drayson 2012; Drayson 2014). But I contend that these general issues are independent from my specific concern with intellectualism.

voluntary mistakes or intentionally do worse than they otherwise could. In fact, Stanley & Krakauer are very explicit about this feature of skills and competences, and they correctly point out that this is at least part of what distinguishes abilities which are in this sense intelligent from other mere abilities or dispositions (cf. Stanley & Krakauer 2013, 3–4). But this clearly shows that, unlike information-processing in general, the possession, selection and application of *knowledge* must always take place at the *personal* rather than the sub-personal level (cf. Pavese 2015b, 16).

In sum, it is inconsistent to hold, first, that the exercise of a competence *does* allow for intentional mistakes and that, second, the exercise of a competence *just is* the selection and application of a specific way of acting, where this does *not* allow for intentional mistakes since it happens sub-personally and therefore not intentionally. Given that it is indeed possible to make intentional mistakes, we must firmly remain on the personal level of description. Ryle's Regress is therefore untouched by these considerations.

This brings me to the second idea which Stanley has offered in defense against Ryle's Regress. This is an argument from symmetry between all kinds of knowledge states. What I take to be the clearest statement of this argument (cf. Stanley 2011b, 17, 26; Stanley 2012a, 733) reads as follows:

[I]t is undeniable that knowledge of how to do something is a standing epistemic state of an agent; both those who deny and those who accept its propositional nature agree with this truism. But for *any* of the standing epistemic states of an agent, she needs to possess automatic mechanisms that are responsible for applying them to particular situations. Even if knowing how to do something were an ability or a complex of dispositions, an agent needs to have automatic mechanisms that are responsible for the application of the ability or the complex of dispositions to the particular situation at hand. (Stanley 2011b, 185)

The problem with this argument is that the 'truism' Stanley mentions at the beginning of this passage is a truism only because it concerns knowledge *of* how to do something rather than know-how. It is precisely the *point* of the anti-intellectualist position, as well as of the Rylean responsibilist view defended here, that know-how is *not* a standing epistemic state in this sense. Instead, know-how is a specific kind of *ability* – however, an ability which involves coming to possess standing epistemic states and acting in the light of them.¹⁷ The Rylean responsibilism I defend does not conceive of know-how as identical with a standing epistemic state, but as a competence to

¹⁷ Ellen Fridland already points in precisely this direction: “[P]erhaps we need to construe knowledge how as intelligent processes, but not standing epistemic states. We will need to incorporate the mechanism of application or triggering into this kind of state rather than allow it to require an independent process of selection. This it seems is the most viable option for the anti-intellectualist to pursue.” (Fridland 2013, 883 fn. 7)

perform intelligent acts *directly*. The intelligence of these acts is explained by the fact that they are exercises of a competence.

But what about Stanley's point that even the 'application' of an *ability* to a particular situation requires automatic mechanisms? To the extent that I understand the notion of 'applying an ability', this simply comes down to the notion of exercising an ability in a given situation. And there is certainly a sense in which this requires automatic mechanisms, for example in the form of sub-personal muscular and brain activity. But the crucial point is that these mechanisms are *part* of the exercise of the ability and thereby part of the intelligent performance *itself*. Of course, there is a distinction between merely possessing a competence on the one hand and actually exercising it on the other hand. But this distinction is not the same as the distinction between possessing standing epistemic states on the one hand and selecting and applying them on the other hand. Let me explain.

We can explain the difference between the mere possession of a competence and the actual performance of an intelligent act simply in terms of the fact that the agent in question exercises her competence. True, since this happens either automatically or intentionally, it presupposes further explanatory elements such as the triggering of a routine reflex on the one hand or a suitable reason, desire or intention on the other hand. But as discussed at length in § 3.2, these further explanatory elements do not explain the very *intelligence* of this performance, but merely *that* the performance occurs in the first place. Instead, the fact that what the agent exercises is a competence directly explains her performance as an intelligent act.

By contrast, we cannot explain the difference between the mere possession of a standing epistemic state and an intelligent performance in the same way because both an automatic and an intentional selection and application of a standing epistemic state is compatible with the performance in question's being not intelligent at all. As Ryle aptly said, "the same crux always recurs" (Ryle 1945a, 8) – to ensure the intelligence of this performance, the selection and application of the relevant standing epistemic state must also be intelligent. Again, Ryle's Regress is fully intact.

Stanley briefly considers supporting his argument from symmetry with the aid of the functionalist position that states of belief or propositional knowledge are themselves *dispositions* which he quotes from Robert Stalnaker (cf. Stanley 2011b, 17–18; Stalnaker 1987, 15). If this is true, then my argument for the conclusion that there is no explanatory gap between know-how and its exercise seems to translate immediately to the conclusion that there is also no explanatory gap between belief or propositional knowledge and their exercise. After all, the same is true for all dispositions and their

manifestations. And independently of the question whether functionalism is indeed true in general (cf. Levin 2013), there seem to be good reasons for the view that beliefs are a kind of disposition (cf. e.g. Schwitzgebel 2002).

In the present context, however, the crucial problem is that this fails to pose an argument against Ryle's view. Against Stanley's interpretation of Ryle, according to which he views belief and propositional knowledge as 'behaviorally inert' (cf. e.g. Stanley 2011b, 11), Ryle explicitly accounts for these at least in part in terms of dispositions and tendencies (cf. e.g. Ryle 1949, 128).¹⁸ The problem of Ryle's Regress in the version I spelled out in § 9.3 does not concern the question of the manifestation of propositional attitudes *in general*, but only the question how their manifestation can explain the specific phenomenon of intelligent practice.

Crucially, the manifestations and the conditions of manifestation of the dispositions which we may identify with beliefs or states of propositional knowledge are not the same as they are for competences. This has also been pointed out very clearly in a recent paper by Brian Weatherson (2016). For belief or propositional knowledge, it is entirely sufficient to have manifestations which consist in verbal avowals or in acts of thinking about practice. Given such manifestations, we clearly have a belief. Thus, the intellectualist has to argue that there are certain kinds of beliefs which have just the right manifestation conditions – i.e. which manifest themselves in practical ways. However, § 9.2 has already argued that this project fails, or that it lends further credit to my Rylean responsibilist alternative.

§ 9.6 Guidance for Intellectualists

Over the course of this chapter, I have argued that intellectualism fails at the very explanatory project which is crucial for the concept of know-how. It fails to fully account for intelligent practice. However, I have repeatedly stressed that intellectualism may only be a short step away from abandoning the identification of know-how with states of propositional or objectual knowledge and instead endorsing a view closer to Rylean responsibilism. In this concluding section, I shall advertize this option one final time.

First, Rylean responsibilism is immune to the problem posed by Ryle's Regress (cf. § 9.3). To see this, it is important to note this argument does not show that it is *false* that standing epistemic states such as conceptions or propositional knowledge are involved in competences. Ryle himself makes explicitly clear that such knowledge is indeed a part of competence:

¹⁸ An excellent discussion of this can be found in Scheffler (1968) and Kremer (2016).

A man knowing little or nothing of medical science could not be a good surgeon, but excellence at surgery is not the same thing as knowledge of medical science; nor is it a simple product of it. The surgeon must indeed have learned from instruction, or by his own instructions or observations, a great number of truths; but he must also have learned by practice a great number of aptitudes. Even where efficient practice is the deliberate application of considered prescriptions, the intelligence involved in putting the prescriptions into practice is not identical with that involved in intellectually grasping the prescriptions. (Ryle 1949, 48–49)

The crucial point of Ryle's Regress is what Ryle reiterates in the last sentence of this quotation. Competence cannot be equated with the possession of standing epistemic states. This is because the capacity to intelligently act in the light of standing epistemic states, in the specific way needed for a given activity, is not entailed by their mere possession.

But this does not show that there are thousands of specific sub-competences responsible for the selection and application of standing epistemic states, or even that there is one mysterious master-capacity for doing so, which is always involved, irrespectively of the activity in question. As I have argued in § 4.5, there is a much more plausible alternative. The competence to intelligently do something in the light of the relevant standing epistemic states is *identical* with the competence to engage in this activity.

Let me illustrate this with the example of *Ginet's Door* on page 281. Why is it so intuitive and unproblematic that I can act in the light of "my knowledge *that* one can get the door open by turning the knob and pushing it (as well as my knowledge *that* there is a door there) by performing that operation quite automatically as I leave the room"? (Ginet 1975, 7) The reason is that I act in the light of these and other standing epistemic states *in exercising my competence to open the door*. The fact that these standing epistemic states can unproblematically bear on my performances in the right way does not show that they provide *the whole* of the explanation of my act as an intelligent act. This crucial point is easily missed by those who support Ginet's point against Ryle's Regress (cf. § 9.4). As friends and foes of intellectualism agree, the explanation of intelligent practice is the touchstone of the concept of know-how. Intellectualists are entirely correct to insist that standing epistemic states are *relevant* in this explanation. But they are mistaken in thinking that they *complete* the explanatory task.

This also sheds light on an important fact which may have confused part of the debate about Ryle's Regress and the consequences of Ginet's point. Intelligent performances are always manifestations of *both* abilities *and* of standing epistemic states such as propositional knowledge. They are exercises of competences, but the idea of an exercise of a competence is partly

explained in terms of the idea that one acts in the light of one's knowledge of facts. As David Wiggins has beautifully brought out in a number of detailed examples, it may be very difficult to set a precise boundary as to what of an agent's knowledge is propositional knowledge of facts and what of her knowledge is practical know-how (cf. Wiggins 2012, 108–116). But this does not threaten the point of Ryle's Regress. The knowledge we need to explain intelligent acts cannot *only* consist in propositional knowledge.

As already stressed in § 8.4, I view the Rylean responsibilist proposal defended in this book as a way of avoiding the problems of intellectualism while preserving its insights. Partly on the basis of the linguistic problem whether the concept of know-how is equated with the concept of competence or merely with a *part* of this concept – like knowledge *about* an activity – (cf. chapter 7), intellectualists have focused too narrowly on such standing epistemic states in understanding the whole of know-how. When explicitly discussing the concept of competence or skill, however, they are sometimes very clear about the point that this is not *only* about knowledge of how to do something. § 8.6 has already discussed this point with respect to the dispositions to acquire the relevant propositional knowledge which Stanley & Williamson (2016) call 'skills'. Likewise, Stanley & Krakauer (2013) describe the way in which such knowledge bears on practice as follows:

[S]killed action is action guided by ongoing accrual and improving application of knowledge of facts about an activity, though skill is not exhausted by such knowledge. (2013, 2)

I am entirely in agreement with this claim. The problem with intellectualism is that it has failed to give a satisfactory account of what competences involve *besides* standing epistemic states. As I have shown in § 9.1 and § 9.2, intellectualists should endorse the fact that actual ability is necessary for know-how by their very own lights – namely, to make sense of their proposals to understand the practical import of know-how. And again in view of their own concerns about the gradability of know-how and the problem of novel situations, § 8.5 has shown that intellectualists should endorse that the competence to assess the relevant activity is necessary for know-how. Thus, Rylean responsibilism is an attractive alternative for intellectualists. Ryle's Regress constitutes a further argument for this conclusion.

As a final push in this direction, it is interesting to consider the way in which Jason Stanley briefly appeals to some of the pertinent literature in metaethics, particularly to the work of Nomy Arpaly and Peter Railton (cf. Arpaly 2003; Railton 2004; Railton 2006; Railton 2009), in what he takes to be an intellectualist answer to Ryle's Regress. He writes:

Railton has exploited essentially Ryle's regress argument in favor of an account of normative guidance that the intellectualist can straightforwardly adopt. [...] [M]uch of the point of Railton's work is to show that behavior that is "automatized" can nevertheless be "norm-guided" and done for a reason[.] [...] What Railton describes as the "enormous amount of apparent fact" that we take on trust from our perceptual experience and memory is, on my view, simply the body of propositional knowledge that a well-functioning agent possesses. Railton's point is that we act on this knowledge quite automatically. (Stanley 2011b, 19–20).

Here, Stanley appeals to Railton's notion of "default trust" (Railton 2004, 186) in order to explain how competent actors rely on their reasons, their propositional knowledge, in an entirely automatic way. But Stanley does not discuss the way in which Railton explains this notion further, namely as what he labels default, *defeasible* trust – that is, as a way of automatically relying on reasons which is nevertheless *open to rational revision* (cf. Railton 2004, 186–189). As such, this idea of automatically acting in the light of propositional knowledge is only superficially congenial to intellectualism. In the end, it straightforwardly exhibits the very feature of intelligence which I have been stressing earlier. It is not a blind automatism, but an automatic routine which is intentionally cultivated. Thus, Stanley's reference to Railton only serves to add weight to my discussion in §9.4.

This point is also revealed in the account of normative guidance which Stanley quotes from Railton's work as allegedly congenial with intellectualism (cf. Stanley 2011b, 21). The quoted passage is this:

Agent A's conduct C is guided by norm N only if C is a manifestation of A's disposition to act in a way conducive to compliance with N, such that N plays a regulative role in A's C-ing, where this involves some disposition on A's part to notice failures to comply with N, to feel discomfort when this occurs, and to exert effort to establish conformity with N even when the departure from N is unsanctioned and non-consequential. (Railton 2006, 13)

While I cannot discuss Railton's views here in detail, I take it that this description of normative guidance is very close to the Sellarsian groundwork I proposed in §3.4 and to my positive account of intellectual guidance as responsible control in §4.4. Railton explicitly says that the guiding role of a norm must involve noticing when the norm is met and when it is not met and actualizing a disposition to correct and improve. Analogously, my conception of know-how involves the competence to assess oneself and to guide oneself in the light of these assessments. If this is the view Stanley wants to defend, then he is already on the brink of abandoning intellectualism and endorsing something closer to the Rylean responsibilism proposed here.