

DAVID LÖWENSTEIN

Know-how as Competence

A Rylean Responsibilist Account

Studies in
Theoretical
Philosophy

vol. 4

KLOSTERMANN

STUDIES IN THEORETICAL PHILOSOPHY

Herausgegeben von Tobias Rosefeldt
und Benjamin Schnieder

in Zusammenarbeit mit

Elke Brendel (Bonn)
Tim Henning (Stuttgart)
Max Kölbel (Barcelona)
Hannes Leitgeb (München)
Martine Nida-Rümelin (Fribourg)
Christian Nimtz (Bielefeld)
Thomas Sattig (Tübingen)
Jason Stanley (New Brunswick)
Barbara Vetter (Berlin)
Marcel Weber (Genf)

vol. 4



VITTORIO KLOSTERMANN

DAVID LÖWENSTEIN

Know-how as Competence

A Rylean Responsibilist Account





VITTORIO KLOSTERMANN

Gedruckt mit Unterstützung des Förderungs-
und Beihilfefonds Wissenschaft der VG Wort

Bibliographische Information der Deutschen Nationalbibliothek
Die Deutsche Nationalbibliothek verzeichnet diese Publikation in der Deutschen
Nationalbibliographie; detaillierte bibliographische Daten sind im Internet über
<http://dnb.dnb.de> abrufbar.

© Vittorio Klostermann GmbH Frankfurt am Main 2017
Alle Rechte vorbehalten, insbesondere die des Nachdrucks und der Übersetzung.
Ohne Genehmigung des Verlages ist es nicht gestattet, dieses Werk oder Teile in einem
photomechanischen oder sonstigen Reproduktionsverfahren oder unter Verwendung
elektronischer Systeme zu verarbeiten, zu vervielfältigen und zu verbreiten.

Gedruckt auf Alster Werkdruck der Firma Geese, Hamburg,
alterungsbeständig  ISO 9706 und PEFC-zertifiziert .

Druck: docupoint GmbH, Magdeburg

Printed in Germany

ISSN 2199-5214

ISBN 978-3-465-03880-1

for Leonhard
the one who made me sing

Contents

Preface	xi
Introduction	1
Part One. Rylean Responsibilism	9
1 Ryle on Know-how and Intelligence	13
§ 1.1 Intelligence and Normative Practice	13
§ 1.2 Clarifying the Normativity of Activities	18
§ 1.3 Living Up to What it Takes	21
§ 1.4 Reliable Ability and Normal Situations	25
§ 1.5 Intelligence as Normative Guidance	29
§ 1.6 Normative Guidance and Rule-Following	35
§ 1.7 Clarifying Loose Boundaries	39
2 Ryle on Intelligence and Intellect	47
§ 2.1 The Manifestation of Know-how	47
§ 2.2 Teaching and Assessing	52
§ 2.3 Learning and the Intellect	58
§ 2.4 Acquisition and Improvement	62
§ 2.5 A Place for Propositional Knowledge	65
§ 2.6 Ryle on Knowledge and the Mind	70
3 Exercising Know-how	77
§ 3.1 Automatic Exercises of Competences	78
§ 3.2 How Know-how Explains Action	84
§ 3.3 Why Know-how Needs Intentional Action	88
§ 3.4 What Ought to Happen	92
§ 3.5 Consciousness and Knowledge of Action	98
§ 3.6 The Phenomenological Fallacy	102

4	Intellectual Guidance	107
§ 4.1	Understanding as Capacity to Assess	107
§ 4.2	Assessment Capacities as Conceptual	111
§ 4.3	Assessments as Propositional Knowledge	115
§ 4.4	Guidance as Responsible Control	122
§ 4.5	Guidance without Regress	129
§ 4.6	Understanding without Regress	133
§ 4.7	Self-Reflexivity without Circularity	138
	Part Two. Cases, Language, and Intellectualism	143
5	The Practical Nature of Know-how	147
§ 5.1	Ability as Necessary for Know-how	148
§ 5.2	On Impossible Exercises of Know-how	152
§ 5.3	Ability as Insufficient for Know-how	158
§ 5.4	Practical Luck	162
§ 5.5	Masks, Finks, and Mimics	166
§ 5.6	Is Know-how a Disposition?	171
6	The Cognitive Nature of Know-how	177
§ 6.1	The Opacity of Know-how Ascriptions	177
§ 6.2	Is there <i>de re</i> Know-how?	182
§ 6.3	The Problem of Epistemic Luck	187
§ 6.4	Accounting for Epistemic Luck	191
§ 6.5	Procedural Knowledge	196
§ 6.6	The Cases of D.F. and H.M.	199
§ 6.7	Machines, Animals, and Infants	204
7	On ‘Knows how to’	209
§ 7.1	‘Knows how to’ and Ability	210
§ 7.2	Pragmatically Salient Alternatives	213
§ 7.3	The Argument from Translation	219
§ 7.4	Ambiguity Tests	223
§ 7.5	The Polysemy of ‘knows how to’	229
8	From Language to Intellectualism	235
§ 8.1	The Standard Linguistic Account	236
§ 8.2	Uniformity and Universality	240
§ 8.3	The Limits of the Linguistic Approach	245
§ 8.4	Common Ground with Intellectualism	250
§ 8.5	Gradability and Novelty	256
§ 8.6	The Gambit of Detaching Skill	261

9 Intellectualism in Practice	267
§9.1 Action-Guiding Understanding	268
§9.2 Practical Modes of Presentation	270
§9.3 Ryle’s Regress	276
§9.4 Selecting and Applying	281
§9.5 Avoiding Ryle’s Regress	284
§9.6 Guidance for Intellectualists	289
Conclusion	293
Bibliography	297
Indices	317
Index of Cases	317
Index of Persons	319
Index of Subjects	325

Preface

This book was long in the making. I started to work on the problem of know-how in 2010, when I first read the seminal article “Knowing How” by Jason Stanley and Timothy Williamson (2001). Like many others, I found myself fascinated, outraged, and perplexed at the same time, and I maintained this wonderful combination of attitudes for some time because of the intriguing mixture of explanatory aims and methodological approaches in the ensuing literature. Eventually, I started to focus on the origins of this recent and current debate, the texts of Gilbert Ryle (1945a; 1949), where I finally found an approach to know-how, and the beginnings of an account of this notion, which pointed me to a firm and promising path into the thicket of the discussion. In any case, this is the path I chose to take.

The first version of this book was completed in 2014 and defended in the same year as my doctoral dissertation, which was supervised by Holm Tetens and Tobias Rosefeldt. Since then, I have continued to discuss this material in many ways and to work on the text off and on, eventually turning the thesis into this book. Time and distance have allowed me to improve the book significantly, whatever its remaining flaws may be. I fear that I was unable to give due credit to all of the fascinating work which has been published in the meantime. But then again, what is a book but a temporary snapshot of continuously developing thought.

In working on this project, I had intellectual homes with wonderful colleagues at Freie Universität Berlin and at Westfälische Wilhelms-Universität Münster, and I was able to discuss my ideas at several conferences and in a number of reading groups and departmental colloquia at Freie Universität Berlin, at Humboldt-Universität zu Berlin, at Technische Universität Berlin, at Ruhr-Universität Bochum, at the University of Cologne, at the University of Duisburg-Essen, at Technische Universität Dresden, at Georg-August-Universität Göttingen, in Kirchberg am Wechsel, at the University of Konstanz, at Universität Leipzig, at Université du Luxembourg, at West-

fälische Wilhelms-Universität Münster, at the University of Osnabrück, at the University of Riga, at the University of St Andrews, and at the University of Vienna.

At these occasions, and in uncounted further conversations, I was immeasurably fortunate to be able to discuss my work with more and with more wonderful friends and colleagues than I could ever have wished. For their criticism, care, and charity, I am deeply, deeply grateful. This book still has many shortcomings. Without this help and support, there would have been many, many more.

Some have accompanied my work on this book from beginning to end, some have touched on it only slightly, and some may even be ignorant of how much they helped me. While I am unable to give due credit to these contributions individually, the least I can do is try, and most certainly fail, to list the names of those who would have deserved this. These are Adam Westra, Alexander Dinges, Andrea Kern, Andreas Müller, Anna Wehofsits, Ansgar Seide, Barbara Vetter, Beate Sachsenweger, Bolesław Czarnecki, Christian Kietzmann, Christian Nimitz, Christian Quast, David Lauer, David Ludwig, Dirk Kindermann, Dirk Koppelberg, Ellen Fridland, Eliot Michaelson, Emanuel Viebahn, Erik Rietveld, Ernest Sosa, Eugen Pissarskoi, Eva-Maria Jung, Fabian Börchers, Greg Sax, Gregor Betz, Hans-Johann Glock, Hannes Worthmann, Helen Bohse, Henry Jackman, Holm Tetens, Jacob Langeloh, James Andow, Jan Constantin, Jan Janzen, Jan Slaby, Jason Stanley, Jason Streitfeld, Jessica Brown, Joshua Habgood-Coote, Julia Zakkou, Kilu von Prince, Lars Dänzer, Luz Christopher Seiberth, Martin Grajner, Martin Kusch, Martin Weicholt, Matthias Haase, Miguel Hoeltje, Nadja El Kassar, Patrick Shirreff, Pedro Schmechtig, Pirmin Stekeler-Weithofer, Raphael Van Riel, Robert Brandom, Romy Jaster, Sebastian Rödl, Stefan Tolksdorf, Thomas Jussuf Spiegel, Tim Kraft, and Tobias Rosefeldt.

I would like to express my gratitude to the editors and the editorial board of *Studies in Theoretical Philosophy* for including my book in this series, and to Anastasia Urban at Vittorio Klostermann Verlag for her guidance and support during the publication process. I also gratefully acknowledge that the publication of this printed book and e-book have been generously supported by the Verwertungsgesellschaft Wort and by the School of Humanities at Freie Universität Berlin.

Finally, my greatest debt and gratitude is to Kilu, for sharing this journey with me, and for so much more.

Introduction

The concept of know-how is at the core of our self-understanding as creatures both theoretical and practical, as beings who care about getting things right both in thinking and in doing. This book is an investigation of this central concept, an attempt to explain what it is to know how to do something.

The concept of know-how is intertwined with at least three other notions: knowledge, understanding, and ability. In the first step of this Introduction, I will present these interconnections and employ them in order to introduce the most prominent views about know-how currently under discussion. My second step will be a discussion of the methodology with which to approach the topic of know-how. I close with a note on the structure of this book and on the different paths one may wish to take in reading it.

The first conceptual connection leads from know-how to *propositional knowledge*, i.e. to knowing that something is the case, knowing a truth or a fact. Intuitively, knowing how to do something is related to the knowledge of facts about doing so or ways or methods to do so. Second, there is the concept of *understanding*. Again, there is an intuitive relation between knowing how to do something and having an understanding or, equivalently, a conception of how to do so or of certain ways or methods to do it. Third, there is the concept of *ability* or, equivalently, of capacity. Knowing how to do something is intuitively connected with having a capacity, with being able to do what one knows how to do.¹

Can these intuitions be maintained? And what do these relations amount to exactly? These questions are at the core of the problem of know-how.

The most prominent views in the debate about know-how² can be understood as attempts to identify *one* of these three conceptual relations as the *whole* of the explanation of the concept of know-how.

¹ I will continue to use ‘true propositions’ and ‘facts’ interchangeably, similarly for ‘an understanding of x’ and ‘a conception of x’, and for ‘ability’ and ‘capacity’.

² For an overview, see Fantl (2008; 2012), Bengson & Moffett (2011b), Jung (2012), Brown & Gerken (2012b), and Pavese (2016b).

The position which has come to be called *intellectualism* can be understood as the claim that know-how consists in a purely intellectual state largely independent from actual ability. The traditional version of this position, *propositionalist intellectualism*, holds that knowing how to do something is a species of knowing that something is the case, i.e. that know-how is just a special kind of propositional knowledge or propositional knowledge of a special kind of truth. By contrast, *objectualist intellectualism* is the view that knowing how to do something is a species of understanding something, i.e. that know-how is just a special kind of conception or a conception of a special kind of thing. The complementary view has come to be called *anti-intellectualism*. This is the claim that know-how consists in a state of actual ability largely independent from belief or understanding.

The central thesis I shall defend is that all of these positions identify a crucial necessary condition of know-how, but falsely claim that this condition is also sufficient for know-how. Knowing how to do something requires an understanding of how to do so, propositional knowledge about doing so, and the actual ability to do so, but none of these are individually sufficient for know-how. Instead, the concept of know-how requires all of these elements, and it requires them to interact with each other in the right way. Intellectualists and anti-intellectualists are both doubly mistaken – first, because they believe that a mere ingredient of know-how constitutes the whole of the phenomenon, and second, because they believe that their opponents’ accounts about know-how fail to even give necessary conditions. In other words, the central thesis of this book is that we can, and should, maintain all three of the intuitive conceptual connections with which I began.³

In order to make this case, I will discuss the most prominent accounts of know-how which are currently maintained – the propositionalist intellectualism defended by Jason Stanley and Timothy Williamson, the objectualist intellectualism defended by John Bengson and Mark Moffett and, albeit less prominently, the anti-intellectualism defended by Hubert Dreyfus.⁴

While these philosophers are paradigmatic proponents of these positions, the *terms* ‘intellectualism’ and ‘anti-intellectualism’ do not have a completely uniform use in the current debate about know-how.⁵

³ These conceptual connections are also expressed in a remark by Ludwig Wittgenstein which I quote as the motto of Part One of this book on page 11. This is inspired by Eva-Maria Jung who also quotes from this as the motto of her book *Gewusst wie?* (2012), even if she omits the last point about understanding and mastery of techniques.

⁴ I mainly rely on Stanley (2005a; 2011a; 2011b; 2011c) as well as Stanley & Williamson (2001; 2016) and Stanley & Krakauer (2013), on Bengson & Moffett (2007; 2011c), and on Dreyfus (2002; 2005; 2007; 2013) as well as Dreyfus & Dreyfus (1986).

⁵ Of course, these terms are also used elsewhere. But in this book, every use of ‘intellectualism’ refers to intellectualism about know-how and likewise for anti-intellectualism.

For example, intellectualism is often identified with propositionalist intellectualism (cf. e.g. Fantl 2008, 451), simply because objectualist intellectualism is comparatively young. Further, while I restrict the intellectualist position to the view that know-how is something exclusively intellectual, an influential article on the state of play of the debate defines intellectualism as the view that know-how “is or involves” some state of the intellect (cf. Bengson & Moffett 2011b, 7–9). On this definition, the account I will propose is an intellectualist account.

Conversely, while I restrict the anti-intellectualist position to the view that know-how is something exclusively practical, other conceptions of anti-intellectualism hold this position merely to assert the explanatory primacy of practical ability with respect to knowledge and understanding (cf. e.g. Fantl 2011, 128; Dickie 2012, 741). On this definition, the account I will propose is an anti-intellectualist account.

Of course, what counts in the end is not the use of labels but the content of positions. Still, the way I conceive of the position defended here, it is neither intellectualist nor anti-intellectualist, but precisely in the middle ground between these views. In other words, I propose a rapprochement between intellectualism and anti-intellectualism, and I hope to show that both views stand to gain from this.

There are many different methodological strands in the debate about know-how. How can this concept be expressed in ordinary language? Which cases exemplify know-how intuitively? What is the relationship between this concept and terms like ‘practical knowledge’ or ‘procedural knowledge’? And what, precisely, are the *phenomena* which we want to capture with a conception of know-how? What is it we aim to explain? Let me develop the approach pursued in this book by briefly walking through these interconnected questions.

The contrast between know-how and propositional knowledge is often understood in terms of the contrast between what we express as ‘knowing that’ and as ‘knowing how’ in English. But this can be misleading. Know-how is knowledge *how to do* something, i.e. a state expressed paradigmatically with the verb ‘to know’, followed by ‘how’ and infinitive. A sentence like ‘I know how tall Leonhard is’ does not involve the concept of know-how so understood because it involves a finite verb phrase. Instead, it seems to express propositional knowledge. Likewise, propositional knowledge is knowledge *that something is the case*, i.e. a state paradigmatically expressed in English with the verb ‘to know’, followed by the complementizer ‘that’ and an embedded proposition. A sentence like ‘I know that smell’ does

not involve the concept of propositional knowledge so understood because it uses the demonstrative article ‘that’ rather than the homophone complementizer and introduces a noun phrase rather than a proposition. Instead, it seems to express knowledge in the sense of objectual acquaintance.⁶

Given these difficulties, it is sometimes suggested that the contrast between know-how and propositional knowledge can be expressed more clearly as the contrast between ‘theoretical’ and ‘practical’ knowledge.⁷ But what one may know how to do also includes intuitively ‘theoretical’ things like solving mathematical problems, and what one may know to be true also includes intuitively ‘practical’ propositions, i.e. that I am currently running or that it is good for me to take a run regularly. Further, the term ‘practical knowledge’ has come to be used in a number of different senses only one of which is equivalent to know-how. Further concepts discussed under this heading include the ‘practical’ knowledge of what I am currently doing, a special kind of non-observational knowledge of action (cf. Anscombe 1957). The contrast between ‘theoretical’ and ‘practical’ knowledge is therefore less helpful than the contrast between know-how and propositional knowledge.

A further methodological commitment is that a philosophical account of the concept of know-how must be answerable to ordinary intuitions as to when it is appropriate to apply the phrase ‘knows how to’. But it would be a mistake to assume that all of these intuitions will necessarily have to be preserved, let alone shown to be accurate. To take an analogy with propositional knowledge, it is perfectly acceptable to say, of a candidate in a quiz show, that she ‘knew’ the correct answer, even if her true belief was everything but sufficiently justified and therefore does not amount to a case of full-blown propositional knowledge in the sense discussed by many epistemologists.⁸ The same should also be allowed in cases of ‘knows how to’. Loose talk is perfectly fine. It is perfectly acceptable to speak of machines, robots and of all kinds of simple-minded animals as ‘knowing how to’ and ‘knowing that’ such and such. Still, it remains an open question whether they possess genuine know-how or genuine propositional knowledge.

At this point, one may turn to cognitive science where a cognate distinction between ‘declarative’ and ‘procedural’ knowledge plays an important role. It is a natural, and plausible, thought to seek clarity on the philosophical problem of know-how in terms of these scientific concepts. However, this

⁶ These distinctions are discussed in detail in chapter 7.

⁷ This strategy is considered e.g. by Glick (2011), Fantl (2012) and Stanley (2012c).

⁸ Of course, some epistemologists hold that such attributions are actually true and not just loose talk. But these complications can be left open here. We can accept how people use the verb ‘to know’ without *settling* these philosophical questions.

is not the approach I will pursue in this book. The concept of know-how is entangled with scientific discourse, and an account of this notion is answerable to these interconnections. But as philosophers, we should seek the most important criteria for assessing an account of our core concepts in the explanatory roles which they play in the whole of our self-understanding. To employ a famous distinction by Wilfrid Sellars, the manifest image of ourselves and our position in the world is essentially interrelated with the scientific image of ourselves and our position in the world (cf. Sellars 1962). But a philosophical attempt to account for a core concept of the manifest image must stay true to the role of this notion within this framework.

The methodology to be pursued in this book begins with the manifest image. For it is here that we encounter the crucial phenomenon for which the concept of know-how is supposed to provide an account – ‘intelligence’ or ‘intelligent practice’. This name for the explanandum of know-how stems, of course, from Gilbert Ryle, the modern classic of the debate about know-how. And despite many changes and further methodological considerations, the aim to explain what Ryle calls ‘intelligence’ continues to be a core commitment throughout the debate.

In this book, the phenomenon of intelligent practice takes center stage. The core criterion for assessing an account of know-how is its role in explaining what Ryle called ‘intelligence’. All the other questions touched upon in this Introduction are also important, but they are secondary to this central theoretical aim. This methodology will show a clear path through the complex thicket of methodologies and topics in the debate about know-how.

This book is not only Rylean in its methodology, but also in the account of know-how it defends. In his Presidential Address to the Aristotelian Society, entitled “Knowing How and Knowing That” (1945a), and in a chapter with the same title in his celebrated book *The Concept of Mind* (1949), Ryle has made a strong case for the relevance of know-how:

Philosophers have not done justice to the distinction which is quite familiar to all of us between knowing that something is the case and knowing how to do things. In their theories of knowledge they concentrate on the discovery of truths or facts, and they either ignore the discovery of ways and methods of doing things or else they try to reduce it to the discovery of facts. (Ryle 1945a, 5)⁹

The advance of knowledge does not consist only in the accumulation of discovered truths, but also and chiefly in the cumulative mastery of methods. (Ryle 1945a, 15)

⁹ I would like to use the first quotation in this book to make the general typographical remark that all emphases in quoted passages are taken over from the original texts.

Ryle's position has received both support and criticism, and the interpretation of his texts has remained one of the central themes in the debate about know-how. Part of my project in this book consists in a favorable reassessment of Ryle's legacy. Maybe Ryle never intended to give a comprehensive account of the concept of know-how (cf. Hornsby 2011, 81), and maybe this topic even turned into a mere stepping stone of the larger project in *The Concept of Mind*. Still, Ryle's texts already contain at least the firm foundation of a very attractive conception of know-how.¹⁰ To see this, it will prove important to broaden the textual scope of Ryle's works beyond the widely and often exclusively read chapter II of *The Concept of Mind*, particularly to also include chapter V on dispositions and chapter IX on the intellect, as well as Ryle's Presidential Address, where he is occupied more exclusively with know-how.

I will propose an interpretation of these texts intended to deepen some of Ryle's insights and to correct some of his errors, and I will defend such a Rylean account of the concept of know-how against the main contenders in the current debate. Among other things, this will consist in an argument against the widespread view that Ryle is an anti-intellectualist in the sense just distinguished.¹¹ Though a dedicated critic of intellectualism, Ryle does not believe that know-how is *merely* an ability, as the anti-intellectualist position states. Instead, he argues that know-how is a special *kind* of ability – a skill or, equivalently, a competence. However, the explanation of this concept also requires appeal to intellectual states such as understanding and propositional knowledge in order to explain intelligent practice through responsible control of one's acts. At the very least, this is the most faithful and most plausible Rylean position which will emerge from my interpretation – a Rylean responsibilism.

This book is organized in two Parts, but it can be read in many different ways. Part One develops and defends my Rylean responsibilist proposal out of an engagement with Ryle's own texts and a number of independent considerations. This also includes a discussion of anti-intellectualism and its shortcomings. Part Two considers the range of candidate cases and examples of know-how, the linguistic expressions of know-how, and it engages in detail with intellectualist views.

Since every chapter begins with an overview and since dependencies between the material in the individual chapters and sections are mentioned and cross-referenced, readers familiar with the debate about know-how will

¹⁰ This has recently been noted by Kremer (2016) and Elzinga (2016).

¹¹ To name just two examples among many, this view is explicitly held by Stanley & Williamson (2001, 416) and Bengson & Moffett (2007, 45 fn. 25).

hopefully be able to read this book in any way and in any part. However, there is one terminological, or conceptual, thing to bear in mind. In this book, I follow Ryle in speaking of know-how *as* competence. That is, I use ‘know-how’, ‘competence’, and ‘skill’ largely interchangeably. This is highlighted and justified in § 1.1. Readers who are unhappy with this may immediately jump to chapter 7, which deals with the linguistics of ‘knows how to’, or to chapter 5, which discusses alleged counterexamples.

Before presenting a short overview of all chapters, let me also mention that this book does not only end with an index of subjects on page 325, but also with two other indices. Readers who are interested in my take on a specific author or on a particular puzzle case may consult the index of persons on page 319 or the index of cases on page 317.

Part One consists of four chapters. The first pair of these offers the groundwork of a Rylean account of know-how while the second pair adds further independent considerations with which this project can be continued in order to reach the explanatory aim of accounting for intelligent practice.

Chapter 1 introduces the concepts of know-how and of ability and gives an account of the phenomenon of intelligent practice as essentially normative. Further, it provides a sketch of what an explanation of this phenomenon must achieve, namely by understanding know-how, skill, or competence, as an *intelligent* ability rather than a *mere* ability. Chapter 2 goes on to discuss Ryle’s remarks on the intellectual part of an intelligent ability, the role of understanding and propositional knowledge. At times, following Ryle’s declared arguments will require correcting inconsistencies, and I will conclude this chapter by assessing where the Rylean view I propose departs from Ryle’s texts, and where it stands with respect to other important themes in their reception, such as the question of behaviorism.

Chapter 3 discusses the question what it is to exercise know-how, and comments on the interrelations between practice, intentionality, and automaticity. Along the way, I will also discuss and reject some considerations in favor of anti-intellectualism. While these topics are largely independent from more specifically Rylean positions, chapter 4 builds on these results in order to complete the line of thought laid out earlier. Ryle himself failed to provide the crucial final element in the explanation of intelligent practice. But I will offer an account which is congenial to Ryle’s position, a Rylean responsibilist account of normative guidance as responsible control.

Part Two consists of five chapters. The first pair of these discusses the range of examples and puzzle cases which have been proposed as candidate examples of know-how and shows how Rylean responsibilism can account for them. The three final chapters are concerned with the linguistics of ‘knows

how to' and related expressions, with the way in which intellectualists have used linguistic considerations in support of their positions, and with these positions themselves.

Chapter 5 reconsiders and defends the Rylean conception of the relationship between know-how and ability and offers explanations of the relevant counterexamples and puzzle cases. This also includes an analogy with structurally identical cases from the independent debate about dispositions, and a discussion of the question if competences or abilities are themselves dispositions. These considerations are complemented by chapter 6 which addresses the puzzle cases which pertain to the *cognitive* rather than the *practical* nature of know-how. This includes crucial semantic and epistemic properties of know-how, as well as certain themes from cognitive science, including the notion of 'procedural knowledge' and some clinical cases which have been discussed with respect to know-how.

Chapter 7 considers the question how the concept of know-how is expressed in ordinary language, and it defends my Rylean use of the expression 'knows how to' as expressing competence against a number of objections. After both pragmatic and semantic considerations, I will eventually suggest that 'knows how to' is polysemous. Such linguistic considerations are continued in chapter 8, but brought to bear on the position of intellectualism which has prominently been supported with linguistic theory. I discuss and eventually reject this line of thought, but I also present an argument for the claim that there is substantial common ground between intellectualism and Rylean responsibilism, including the question of compatibility with the linguistic data. The final chapter 9 addresses the question if intellectualism can succeed at the explanatory task of accounting for intelligent practice. I argue that the positive intellectualist proposals fail at this task, and I defend Ryle's famous regress objection as an argument which establishes a principal problem for such attempts. In the end, however, I suggest that the way in which intellectualists have reacted to this argument speaks in favor of the rapprochement advertized in this book.

Part One

Rylean Responsibilism

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

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

(Wittgenstein 1953, § 150)

Chapter 1

Ryle on Know-how and Intelligence

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

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

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

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

§ 1.1 Intelligence and Normative Practice

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

*Ryle's Range of Cases*¹

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

§ 1.2 Clarifying the Normativity of Activities

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

§ 1.3 Living Up to What it Takes

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

§ 1.4 Reliable Ability and Normal Situations

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

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

Lottery (Stanley & Williamson 2001, 414–415)

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

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

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

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

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

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

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

Avalanche (Hawley 2003, 27)

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

Lucky Cake (Hawley 2003, 27)

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

§ 1.5 Intelligence as Normative Guidance

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Clocks & Seals (Ryle 1949, 29)

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

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

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

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

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

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

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

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

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

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

Digestion (Stanley & Williamson 2001, 414)

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

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

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

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

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

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

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

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

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

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

§ 1.6 Normative Guidance and Rule-Following

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

§ 1.7 Clarifying Loose Boundaries

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

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

Park Chess

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Driving Contexts (Hawley 2003)

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

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

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

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

Animals & Infants (Ryle 1949, 121)

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

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

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

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

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

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

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

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

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

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

Chapter 2

Ryle on Intelligence and Intellect

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

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

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

§ 2.1 The Manifestation of Know-how

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

But Ryle seems to contradict this view elsewhere:

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

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

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

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

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

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

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

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

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

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

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

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

§ 2.2 Teaching and Assessing

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

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

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

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

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

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

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

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

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

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

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

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

Bela Karoli (Fridland 2012, 9)

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

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

Critics & Writers (Ryle 1949, 49)

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

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

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

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

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

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

Knotting Sailor (Ryle 1949, 55)

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

§ 2.3 Learning and the Intellect

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

§ 2.4 Acquisition and Improvement

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

§ 2.5 A Place for Propositional Knowledge

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Pink Ice Cream (Millikan 2000, 64)

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

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

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

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

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

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

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

§ 2.6 Ryle on Knowledge and the Mind

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Chapter 3

Exercising Know-how

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

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

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

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

§ 3.1 Automatic Exercises of Competences

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

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

Unwelcome Sign

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Similarly, Katherine Hawley holds:

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

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

(1) If one Fs, one employs knowledge how to F.
 [...] Premise (1) is true only when the range of actions is restricted to *intentional actions*. (Stanley & Williamson 2001, 413, 415)

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

§ 3.2 How Know-how Explains Action

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

§ 3.3 Why Know-how Needs Intentional Action

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

§ 3.4 What Ought to Happen

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

§ 3.5 Consciousness and Knowledge of Action

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

§ 3.6 The Phenomenological Fallacy

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

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

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

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

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

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

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

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

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

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

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

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

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

Chuck Knoblauch (Dreyfus 2007, 354)

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

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

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

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

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

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

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

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

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

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

Chapter 4

Intellectual Guidance

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

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

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

§ 4.1 Understanding as Capacity to Assess

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

§ 4.2 Assessment Capacities as Conceptual

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

§ 4.3 Assessments as Propositional Knowledge

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

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

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

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

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

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

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

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

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

This argument establishes the following claim:

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

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

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

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

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

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

Likewise, being competent at swimming or otherwise possessing an understanding of swimming requires knowledge of the form (B2) because it requires the correct assessment of paradigm situations which offer the option to swim and of paradigm situations which do not offer this option at all. To take simple examples, this involves the knowledge that swimming requires a certain minimum amount of water and the absence of other swimmers or objects who are constantly in one's way, and the knowledge that a perfectly calm swimming pool without any people or objects floating around offers a perfect opportunity for swimming. Again, if somebody fails to possess such propositional knowledge, we would, at the very least, have to substantially adjust our interpretation of this person in order to maintain the idea that she is actually still competent to assess what she so blatantly misconceives.

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

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

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

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

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

- (C) Every capacity to assess A-ing involves knowledge of at least some principles of assessment of acts and situations of the forms (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.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

§ 4.4 Guidance as Responsible Control

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

the form (D1), and in the light of one's assessments of the options in a given situation, i.e. of knowledge of the form (D2). These ways rely on the paradigms of fully intentional actions on the one hand and completely automatic performances on the other hand, which I have discussed before, and of which I have already said that they merely mark the clearest ends of the continuum of possible degrees.

For example, I can begin by intentionally assessing my own performances as well as the relevant situation, then go on to consciously judge that such-and-such would be the best thing to do now, determining the best option in my knowledge of the form (D2), and intentionally perform this act.

In the same way, one can also assess one's performances *in action* or *during* their performance. One may realize that one is on the way to complete a very good performance and therefore make sure to keep on track, relying first on one's knowledge about how well one is currently doing of the form (D1). In this way, one may see how one's present course of action could be improved, comparing one's knowledge about one's current options of the form (D2), and see to taking advantage of the best option. Likewise, one may perceive mistakes in one's performances and try to correct them.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

§ 4.5 Guidance without Regress

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

§ 4.6 Understanding without Regress

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

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

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

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

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

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

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

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

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

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

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

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

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

(B 1) Token act *x* in token circumstances *c* is a performance of A-ing of quality *Q*.

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

These verdicts are conceptually connected in the following way.

(E) The self-reflexivity of assessments can be expressed as the conceptual truth (E 1), an instance of which is (E 2).

(E 1) *x* has quality *Q* with respect to *y* just in case *that x has quality Q with respect to y* is an accurate assessment of *x* with respect to *y*.

(E 2) Token act *x* in token circumstances *c* is a performance of A-ing of quality *Q* just in case *that token act x in token circumstances c is a performance of A-ing of quality Q* is, in circumstances *c*, an accurate assessment of *x* as a performance of A-ing.

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

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

claim, ensuring a conceptual bridge between (B1) and (B1'). This shows that assessing performances of an activity and assessing such assessments of that activity are exercises of the same competence – the conceptual capacity to correctly employ the concept of that activity.

A natural source of skepticism about this inference is that, on the face of it, second-order assessments of the form (B1') which refer to first-order assessments of the form (B1) seem to involve a *further* concept over and above those employed in first-order assessments alone – namely, the concept of an *assessment* of A-ing. But I contend that, in fact, (B1) *does* involve the concept of an assessment – not explicitly, of course, but *implicitly*.

I contend that every token act *x* in token circumstances *c* is a performance of A-ing of quality *Q* just in case, *assessed as* a performance of A-ing in circumstances *c*, *x* *has* quality *Q*. This shows that every assessment of the form (B1) implicitly involves the concept of an assessment. For example, for an act to be good *as* a performance of a given activity just is for it to be good *as assessed* as a performance of that activity. The relation between a particular act and a particular activity in virtue of which the act meets the norms of that activity *already is* the assessment relation we were looking for. The competence to make assessments of the form (B1) already entails having the concept of an assessment. Thus, the conceptual truth expressed in (E2) shows that the competence to assess performances of an activity is one with the competence to assess these assessments.

This concludes the first of my three arguments. The second argument also proceeds by way of supporting (E2). Here, I mainly rely on the conceptual nature of the assessment capacities, as discussed in §4.2.

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

- (F) The self-reflexivity of judgment and concept use can be expressed as (F1), an instance of which is (F2), and a further specified version of this, in turn, is (F3).
- (F1) It is conceptually true that *p* just in case *that p* is true.
- (F2) It is conceptually true that 'F' applies to *x* just in case *that 'F' applies to x* is accurate.
- (F3) It is conceptually true that, in circumstances *c*, 'F' applies to *x* just in case *that, in circumstances c, 'F' applies to x* is, in circumstances *c*, an accurate assessment of *x* as a candidate for applying 'F'.

First, (F3) clearly entails (E2). Unlike (E2), (F3) is not concerned with the use of a specific concept of the form 'performance of A-ing of the quality *Q*' and its application to a specific performance *x*, but instead with the use

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

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

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

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

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

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

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

§ 4.7 Self-Reflexivity without Circularity

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

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

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

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

This answer brings out an important simplification with which I have been working so far. Until now, I have assumed that the token circumstances in which the second-order assessment takes place are the same token circumstances in which the first-order assessment takes place. This is why I used the same schematic letter 'c' on both sides in the conceptual principle (E2). However, the circumstances in which a second-level assessment of a first-order assessment is undertaken can, but need not be the same circumstances in which that -level assessment is undertaken.

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

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

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

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

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

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

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

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

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

In § 2.2, I already discussed one example of such an activity when I quoted Ryle's view that the capacity to assess philosophical work requires the capacity to philosophize oneself (cf. Ryle 1949, 53–54). As I argued

there, Ryle falsely generalized from this example and held that every assessment capacity requires the corresponding skill. However, it has remained an open question how this local peculiarity can be explained and why it is that the capacity to assess philosophical work requires the knowledge how to philosophize oneself. Now, I can finally give a straightforward answer to this problem. This answer is simple. Philosophy is concerned with concepts, at least primarily and possibly even exclusively so. Of course, the philosophy of philosophy is itself a controversial field on the intricacies of which I cannot even begin to comment here. I shall nevertheless venture the assumption that the activity of philosophy is, at least largely, engaged with concepts and with their change, analysis, creation, clarification, deconstruction, improvement, and so forth.

As I have just argued, conceptual capacities are the unique kind of competence which is self-reflexive in that the competence to assess its performances is just a meta-level instance of the competence itself. Accordingly, my account directly predicts that every conceptual capacity is such that one cannot merely understand what it takes to exercise it well without also being able to do so oneself. And Philosophizing is a conceptual activity.³¹

A further such example has been proposed by David Carr:

Multiplication (Carr 1981a, 54)

To understand the rules of multiplication is *ipso facto* to be able to multiply and hence there is an air of paradox about the statement—‘he understands multiplication but cannot multiply’.

This is perfectly correct. Understanding the rules of multiplication in the relevant sense and the competence to multiply oneself cannot be separated. The account of know-how I have offered can explain this fact by saying that *performing* acts of multiplication and *assessing* acts of multiplication are two sets of exercises of one and the same conceptual competence – the use of the concept of multiplication.³²

To clarify, the problem brought up by Carr is *not* that one may understand multiplication and fail to know the meanings of expressions like ‘+’ or ‘plus’.³³ One may have no clue whatsoever about these terms, but

³¹ In § 7.1, I will discuss a further case in point by John Bengson and Marc Moffett. This example will be quoted in sentences (5 a–d) on page 212.

³² The view that competence at multiplication is a conceptual affair does not rule out the fact that simple machines like pocket calculators *can* multiply, too. However, their capacities do not constitute genuine competences, but remain on the level of mere dispositions or mere abilities. I will come back to these problems in § 6.7.

³³ A thought experiment along these lines underlies Saul Kripke’s famous interpretation of Ludwig Wittgenstein’s remarks on rule-following (cf. Kripke 1982). On the relationship between these issues and the question of know-how, see § 1.6.

nevertheless understand multiplication perfectly well, maybe on the basis of learning to calculate with non-standard symbols or simply in a different language. In such a case, the person in question would both understand multiplication and have the competence to multiply. She would simply be unable to exercise her competence with the aid of certain tools, which she fails to understand, but which are, in fact, entirely optional. A different sign with the same meaning would do the job as well.

This can be explained very well on the basis of the fact that an understanding of multiplication involves the concept of this activity. As I have argued in § 4.2, there is an important distinction between concepts and the words or symbols which express them. One may therefore possess an understanding of multiplication and thereby possess the concept of this activity, but nevertheless fail to be able to identify a foreign symbol which expresses it. On the other hand, understanding multiplication and thereby having the concept of this activity is already sufficient for knowing how to multiply.

This concludes my discussion of the self-reflexivity of conceptual assessment competences, the final element of my Rylean responsibilist account of know-how. I have argued that Ryle correctly explains know-how as an intelligent ability and I have explained this notion as a reliable ability the exercises of which are guided by the competence to assess performances of this ability and options to do so. This latter competence is itself a piece of know-how, but it can be understood in terms of its self-reflexivity, the fact that its exercises are guided by further exercises of itself.

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 mimics 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 responsibility 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 Bengson, 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 Sytma & 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. (7a), with no third option along the lines of (7b) or (7c).

If it is true that there is a subtle semantic distinction between (7a) and (7b–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. (7a), should be read along the lines of (7b–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 (9 a) 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 (9 a) 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 (9 b) 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 (9 b) back into English would not lead to (8), but instead to ‘Linda can perform a backflip’.

A translation which makes these things clearer than (9 b) is given in (9 c). 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 (9 a) and (9 c) 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 (19 a) and (19 b), 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, (19 c) and (19 d) 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 (13 a) 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, (14 a–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 intellectualist 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 Mastro (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 Mastro (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 responsibility. 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. Mastro 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 languely 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. [(37 a)] is true. Nevertheless, it seems that Hannah does not know how to ride a bicycle, i.e. [(37 b)] 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, (37 a) *entails* (37 b). After all, (37 a) 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, (37 a) 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*

- (M1) If S intelligently ϕ s, then S' ϕ -ing is intelligent because S knows how to ϕ .
- (M2) 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*
- (M3) *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 (M1-2)*
- (M4) 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.
- (M5) *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 (M3-4)*
- (M6) It is false that, if S intelligently ϕ s, then S' ϕ -ing is intelligent because S intelligently performs infinitely many further acts.
- (M7) *Thus*: (M2) is false, i.e. intellectualism is false. *from (M5-6)*

Evidently, the crux of this argument lies with premise (M4). 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 (M4) 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 (M4) 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 (M4) 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.

Conclusion

This book has been an investigation of what it is to know how to do something. I have argued that the concept of know-how is the concept of an *intelligent* ability – of a reliable ability to do well in an activity in virtue of intellectual guidance. Know-how requires the reliable ability to engage in an activity, it requires the understanding of what it takes to do well in this activity, and it requires the propositional knowledge of the quality of what one does and of the options available in a situation. But *pace* both intellectualism and anti-intellectualism, know-how is not identical with any of these elements. It is the complex capacity of their interacting in such a way that competent agents responsibly control their performances in the light of their knowledge and understanding. By seeing how these features are connected with each other, we can fully understand intelligent practice.

In Part One, I have developed this line of thought on the basis of Gilbert Ryle's seminal work on know-how, spelling out the idea of intelligent ability and normative practice in chapter 1 and discussing the crucial role of the intellect for this notion of intelligence in chapter 2. This line of thought has continued with a number of independently motivated considerations about the role of automaticity and intentionality in the exercise of know-how in chapter 3 and culminated in chapter 4 in what I called Rylean responsibilism, an account of what it is to be guided by an understanding of what it takes to do well in an activity.

This line of thought began on the common ground between Rylean responsibilism and anti-intellectualism, with the notion of a reliable ability. But it went on to show why the anti-intellectualist equation of know-how and ability fails to capture the explanatory point of the concept of know-how, the idea of intelligence in the sense of normative guidance. As Ryle already saw, intelligence requires the intellect.

The second larger line of thought pursued in this book was concerned with the intellectualist view that the intellectual elements just shown to be

crucial for know-how are all there is to this concept. In chapters 7 and 8, I have argued that my Rylean use of the expressions ‘knows how to’ and ‘know-how’ is not undermined by the facts and theories about the English language which intellectualists have cited in support of their view. My final assessment of intellectualism in chapter 9 has brought together these considerations with the positive line of thought developed in Part One. As with anti-intellectualism, there is common ground between Rylean responsibility and intellectualism. But I have proposed a number of arguments for the conclusion that intellectualism also fails to give a satisfactory explanation of what the concept of know-how aims to explain, intelligent practice.

In the Introduction, I have already proposed to conceive of these conclusions as a rapprochement between intellectualist and anti-intellectualist views. As I have argued over the course of this book, neither of these positions can be maintained without incorporating some of the materials originally reserved for the opposing view. The anti-intellectualist identification of know-how with ability fails to account for the crucial distinction between *mere* ability and genuine know-how, where the latter also requires intellectual elements like knowledge and understanding. Conversely, intellectualism tries to account for know-how simply as such an intellectual state. But such knowledge or understanding can only be genuine knowledge how to *do* something, as opposed to a *mere* understanding of this activity or mere knowledge *about* it, if abilities are also incorporated into the account.

It seems, then, that the debate about know-how confirms a conjecture which Albert Einstein expressed in 1918 in a letter to Eduard Study – that any two ‘-isms’ become alike when they are cleared of all clutter: “Wenn man zwei beliebige ‚ismusse‘ von allem Unrat säubert, dann werden sie einander gleich.” (cf. Schulmann *et al.* 1998, 890)

At the beginning of Part Two, on page 145, I have quoted Wittgenstein’s remark that some terms of ordinary language require philosophical scrutiny before we can continue to use them without causing confusion. Given the conclusions of this book, I contend that the recent debate about know-how, in the wake of the influential work of Stanley & Williamson (2001), has often been conducted with too much emphasis on language and with too little emphasis on the crucial explanatory point of the concept of know-how – its role in understanding intelligent practice and normative guidance.

Conversely, however, one may also ask why it should even be important, if *that* is the project, to somehow tie it to the expressions ‘knows how to’ and ‘know-how’ in the first place. Would it not be easier to just give away this terminology and straightforwardly talk about skill and competence? This

conceptual policy would certainly be welcomed by intellectualists – indeed it is endorsed in the latest paper by Stanley & Williamson (2016) (cf. § 8.6). On this view, an actor's know-how would not *be* her competence, but only *part* of her competence, and we would lead the discussion in terms of the question how such know-how figures in the intellectual guidance which is *involved* in the exercise of competences.

If this is indeed merely a question of terminology, while all the substantive explanatory issues remain as they are, then I am perfectly happy to endorse this option. This book could then simply be rewritten such that Rylean responsibilism is developed from the intellectualist end, by starting with understanding and propositional knowledge, rather than from the anti-intellectualist end, by starting with ability. Still, the journey would end in the same place, and for the same reasons.

But terminology is seldom *mere* terminology. In fact, I contend that there are positive reasons to maintain the equation of know-how with competence and skill. Competences are epistemic achievements, after all. Just like states of propositional knowledge are instances of getting things right with respect to the facts, competences are instances of getting things right with respect to normative practice. This idea is beautifully captured in the expression 'to know how to do something'. The parallels between competence and propositional knowledge suggest that both are species of the common genus of knowledge. And a natural way to mark this fact is to use 'know-how' to refer to competence.

This perspective can be further substantiated when these parallels between know-how and propositional knowledge are explained in more detail. I contend that they exhibit an analogous structure which is arguably definitive of their nature as epistemic achievements. Ryle expresses this common element in the genus of knowledge as follows:

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

But however *Ryle* thought of this structure of knowledge as the genus common to the species of know-how and propositional knowledge (cf. Kremer 2016), I take it that this idea plausible independently.¹⁹ Just like propo-

¹⁹ Despite a number of differences in detail, arguments along these lines have been variously suggested by philosophers like Carr (1979; 1981a; 1981b), Smythe (1988), Craig (1990), Hawley (2003), J. Williams (2008), Kumar (2011), Tsai (2011b; 2013), and Kremer (2016). On the basis of Robert Brandom's notion of propositional knowledge as what he calls a 'complex hybrid deontic status' (cf. Brandom 1994; 1995; 2000), I have elsewhere sketched an analogous view of know-how (cf. Löwenstein 2011b), generalizing from Brandom's account of perceptual capacity, something he explicitly calls 'another hybrid deontic status' without reconnecting it to propositional knowledge.

sitional knowledge involves but is not exhausted by true belief, know-how involves but is not exhausted by reliable ability. And just like propositional knowledge can be understood as a belief which is true *because* it is justified in the right way, know-how is an ability which is reliable *because* it is intellectually guided in the right way. But of course, this sketch of the parallels between these kinds of knowledge requires much further elaboration.

A further important part of these general epistemological questions concerns the holistic interdependence of know-how and propositional knowledge. As I have argued over the course of this book, not only is it crucial to see that know-how requires certain forms of propositional knowledge, most importantly in the form of correct assessments of the activity in question. Conversely, propositional knowledge also requires know-how in the form of the competence to employ the relevant concepts and in the form of certain more specifically epistemic competences such as the competence to assess an activity. These considerations promise to contribute to a broader picture of both the common nature and the different kinds of our epistemic achievements.²⁰

The arguments presented here raise many further questions and have consequences for a number of important topics, some of which I have already flagged over the course of this book. But it would not be worthy of these further questions to scratch at their surfaces now. Instead, I would like to conclude by stressing that the concept of know-how has indeed turned out to be at the very core of our conception of ourselves. It provides a crucial connection between ourselves as agents and as thinkers, as persons striving for truth and understanding and as persons who try to do well at what they do. Or, to put this point more aptly, Rylean responsibilism reveals that these dimensions have always been interconnected within the concept of know-how.

²⁰ A related topic which I have not been able to discuss is the position Stephen Hetherington calls ‘practicalism’, the view that propositional knowledge is a species of know-how which is held by Hartland-Swann (1956; 1957) and Hetherington (2006; 2008a; 2008b; 2011a; 2011b); for discussion, see Ammerman (1956), Adams (2009), Madison (2012), Fantl (2012), and Kremer (2016). Elsewhere, I have briefly argued that practicalism fails to give a satisfactory account of the manifestation of propositional knowledge in much the same way in which intellectualism fails to give a sufficient account of the manifestation of know-how in intelligent practice (cf. Löwenstein 2013).

Bibliography

- Abbott, Barbara. 2013. Linguistic Solutions to Philosophical Problems: The Case of Knowing How. *Philosophical Perspectives*, **27**, 1–21.
- Adams, Marcus P. 2009. Empirical evidence and the knowledge-that/knowledge-how distinction. *Synthese*, **170**, 97–114.
- Adler, Jonathan. 2012. Epistemological Problems of Testimony. *In*: Zalta (1995 ff.). Spring 2013 Edition. <http://plato.stanford.edu/archives/spr2013/entries/testimony-episprob/>.
- Alexander, Joshua, & Weinberg, Jonathan M. 2007. Analytic Epistemology and Experimental Philosophy. *Philosophy Compass*, **2**, 56–80.
- Allen, Colin, & Trestman, Michael. 2014. Animal Consciousness. *In*: Zalta (1995 ff.). Summer 2014 Edition. <http://plato.stanford.edu/archives/sum2014/entries/consciousness-animal/>.
- Alter, Torin. 2001. Know-How, Ability and the Ability Hypothesis. *Theoria*, **67**, 229–239.
- Ammerman, Robert. 1956. A Note on ‘Knowing That’. *Analysis*, **17**, 30–32.
- Andrews, Kristin. 2012. Animal Cognition. *In*: Zalta (1995 ff.). Winter 2012 Edition. <http://plato.stanford.edu/archives/win2012/entries/cognition-animal/>.
- Annas, Julia. 2011. Practical Expertise. *In*: Bengson & Moffett (2011a). 101–112.
- Anscombe, Gertude Elizabeth Margaret. 1957. *Intention*. Oxford: Basil Blackwell.
- Aristotle. *Physics*. Translation by Robin Waterfield. Introduction and Notes by David Bostock. 2008. Oxford: Oxford University Press.
- Arpaly, Nomy. 2003. *Unprincipled Virtue: An Inquiry into Moral Agency*. Oxford: Oxford University Press.
- Ashwell, Lauren. 2010. Superficial Dispositionalism. *Australasian Journal of Philosophy*, **88**, 635–653.
- Audi, Paul. 2012. Grounding: Toward a Theory of the *In-Virtue-of* Relation. *The Journal of Philosophy*, **109**, 685–711.
- Bach, Kent. 1978. A Representational Theory of Action. *Philosophical Studies*, **34**, 361–379.
- Bach, Kent. 2012. Review of *Know How* by Jason Stanley. *Notre Dame Philosophical Reviews*, March 7, 2012. <http://ndpr.nd.edu/news/29390-know-how-2/>.
- Bargh, John A., & Chartrand, Tanya L. 1999. The Unbearable Automaticity of Being. *American Psychologist*, **54**, 462–479.

- Bartels, Andreas, & May, Mark. 2015a. What a Theory of Knowledge-How Should Explain. A Framework for Practical Knowledge beyond Intellectualism and Anti-Intellectualism. *In: Metzinger & Windt (2015)*. 2(T). 1–20.
- Bartels, Andreas, & May, Mark. 2015b. Preparing the Ground for an Empirical Theory of Knowing-How. A Reply to Ramiro Glauer. *In: Metzinger & Windt (2015)*. 2(R). 1–3.
- Battaly, Heather. 2008. Virtue Epistemology. *Philosophy Compass*, **3**, 639–663.
- Beck, Jacob. 2012. Do Animals Engage in Conceptual Thought? *Philosophy Compass*, **7**, 218–229.
- Beilock, Sian L., Carr, Thomas H., MacMahon, Clare, & Starkes, Janet L. 2002. When Paying Attention Becomes Counterproductive: Impact of Divided Versus Skill-Focused Attention on Novice and Experienced Performance of Sensorimotor Skills. *Journal of Experimental Psychology: Applied*, **8**, 6–16.
- Bekoff, Marc, Allen, Colin, & Burghardt, Gordon M. (eds). 2002. *The Cognitive Animal: Empirical and Theoretical Perspectives on Animal Cognition*. Cambridge: MIT Press.
- Bengson, John, & Moffett, Marc A. 2007. Know-how and concept possession. *Philosophical Studies*, **136**, 31–57.
- Bengson, John, & Moffett, Marc A. (eds). 2011a. *Knowing How. Essays on Knowledge, Mind, and Action*. Oxford & New York: Oxford University Press.
- Bengson, John, & Moffett, Marc A. 2011b. Two Conceptions of Mind and Action: Knowing How and the Philosophical Theory of Intelligence. *In: Bengson & Moffett (2011a)*. 3–55.
- Bengson, John, & Moffett, Marc A. 2011c. Nonpropositional Intellectualism. *In: Bengson & Moffett (2011a)*. 161–195.
- Bengson, John, Moffett, Marc A., & Wright, Jennifer C. 2009. The folk on knowing-how. *Philosophical Studies*, **142**, 387–401.
- Berry, Dianne C., & Broadbent, Donald E. 1984. On the Relationship between Task Performance and Associated Verbalizable Knowledge. *The Quarterly Journal of Experimental Psychology*, **36**, 209–231.
- Bhatt, Rajesh. 2006. *Covert Modality in Non-finite Contexts*. Berlin & New York: Mouton de Gruyter.
- Bird, Alexander. 1998. Dispositions and Antidotes. *The Philosophical Quarterly*, **48**, 227–234.
- Blackburn, Simon. 1984. The Individual Strikes Back. *Synthese*, **58**, 281–301.
- Blome-Tillmann, Michael. 2009. Contextualism, Subject-Sensitive Invariantism, and the Interaction of ‘Knowledge’-Ascriptions with Modal and Temporal Operators. *Philosophy and Phenomenological Research*, **79**, 315–331.
- Blome-Tillmann, Michael. 2013. Conversational Implicatures (and How to Spot Them). *Philosophy Compass*, **8**(2), 170–185.
- Boden, Margaret A. (ed). 1990. *The Philosophy of Artificial Intelligence*. Oxford: Oxford University Press.
- Boghossian, Paul. 1989. The Rule-Following Considerations. *Mind*, **98**, 392–549.
- Boutin, Arnaud, Blandin, Yannick, Massen, Cristina, Heuer, Herbert, & Badets,

- Arnaud. 2014. Conscious awareness of action potentiates sensorimotor learning. *Cognition*, **133**(1), 1–9.
- Brandom, Robert. 1994. *Making It Explicit. Reasoning, Representing, and Discursive Commitment*. Cambridge: Harvard University Press.
- Brandom, Robert. 1995. Knowledge and the Social Articulation of the Space of Reasons. *Philosophy and Phenomenological Research*, **40**, 895–908.
- Brandom, Robert. 2000. *Articulating Reasons. An Introduction to Inferentialism*. Cambridge: Harvard University Press.
- Braun, David. 2006. Now You Know Who Hong Oak Yun Is. *Philosophical Issues*, **16**, 24–42.
- Braun, David. 2011. Knowing How and Knowing Answers. In: Bengson & Moffett (2011a). 244–260.
- Brett, Nathan. 1974. Knowing How, What and That. *Canadian Journal of Philosophy*, **4**, 293–300.
- Brogaard, Berit. 2006. Can Virtue Reliabilism Explain the Value of Knowledge? *Canadian Journal of Philosophy*, **36**, 335–354.
- Brogaard, Berit. 2008a. Attitude Reports: Do You Mind the Gap? *Philosophy Compass*, **3**, 93–118.
- Brogaard, Berit. 2008b. Knowledge-*the* and propositional attitude ascriptions. *Grazer Philosophische Studien*, **77**, 147–190.
- Brogaard, Berit. 2011. Knowledge-How: A Unified Account. In: Bengson & Moffett (2011a). 136–160.
- Brown, David G. 1970. Knowing How and Knowing That, What. In: Wood, George, & Pitcher, Oscar Patrick (eds), *Ryle*. London: MacMillan.
- Brown, David G. 1974. Reply to Brett. *Canadian Journal of Philosophy*, **4**, 301–303.
- Brown, Jessica A. 2013. Knowing-how: linguistics and cognitive science. *Analysis*, published online on February 4, 2013.
- Brown, Jessica A., & Gerken, Mikkel (eds). 2012a. *Knowledge Ascriptions*. Oxford: Oxford University Press.
- Brown, Jessica A., & Gerken, Mikkel. 2012b. Knowledge Ascriptions. Their Semantics, Cognitive Bases, and Social Functions. In: Brown & Gerken (2012a). 1–30.
- Brownstein, Michael. 2014. Rationalizing flow: agency in skilled unreflective action. *Philosophical Studies*, **168**, 545–568.
- Brownstein, Michael, & Michaelson, Eliot. 2016. Doing without believing: Intellectualism, knowledge-how, and belief-attribution. *Synthese*, **193**, 2815–2836.
- Burge, Tyler. 1977. Belief De Re. *The Journal of Philosophy*, **74**, 338–362.
- Buskell, Andrew. 2015. How to be skilful. opportunistic robustness and normative sensitivity. *Synthese*, published online on January 7, 2015.
- Bzdak, David. 2008. On Amnesia and Knowing-How. *Techné: Research in Philosophy and Technology*, **12**.
- Capone, Alessandro. 2011. Knowing how and pragmatic intrusion. *Intercultural Pragmatics*, **8**, 453–470.
- Carr, David. 1979. The Logic of Knowing How and Ability. *Mind*, **88**, 394–409.

- Carr, David. 1981a. Knowledge in Practice. *American Philosophical Quarterly*, **18**, 53–61.
- Carr, David. 1981b. On Mastering a Skill. *Journal of Philosophy of Education*, **15**, 87–96.
- Carter, J. Adam, & Czarnecki, Boleslaw. 2016. Extended Knowledge-How. *Erkenntnis*, **81**, 259–273.
- Carter, J. Adam, & Pritchard, Duncan. 2013. Knowledge-How and Epistemic Luck. *Noûs*, published online on November 21, 2013.
- Carter, J. Adam, & Pritchard, Duncan. 2015. Knowledge-How and Cognitive Achievement. *Noûs*, **91**, 181–199.
- Cath, Yuri. 2009. The Ability Hypothesis and the New Knowledge-how. *Noûs*, **43**, 137–156.
- Cath, Yuri. 2011. Knowing How without Knowing That. In: Bengson & Moffett (2011a). 113–135.
- Cath, Yuri. 2013. Regarding a Regress. *Pacific Philosophical Quarterly*, **94**, 358–388.
- Cath, Yuri. 2014. Revisionary intellectualism and Gettier. *Philosophical Studies*, published online on January 4, 2014.
- Chemero, Anthony. 2003. An Outline of a Theory of Affordances. *Ecological Psychology*, **15**, 181–195.
- Choi, SungHo, & Fara, Michael. 2012. Dispositions. In: Zalta (1995 ff.). Spring 2012 Edition. <http://plato.stanford.edu/archives/spr2012/entries/dispositions/>.
- Christensen, Wayne, Sutton, John, & McIlwain, Doris. 2016. Cognition in Skilled Action. Meshed Control and the Varieties of Skill Experience. *Mind & Language*, **31**, 37–66.
- Clark, Andy, & Chalmers, David. 1998. The Extended Mind. *Analysis*, **58**, 7–19.
- Clarke, Henry. 2016. Wiggins on Practical Knowledge. *Disputatio*, **42**, 113–124.
- Clarke, Randolph. 2009. Dispositions, Abilities to Act, and Free Will: The New Dispositionalism. *Mind*, **118**, 323–351.
- Clarke, Randolph. 2015. Abilities to Act. *Philosophy Compass*, **10**, 893–904.
- Cohen, Neal J., & Squire, Larry R. 1980. Preserved Learning and Retention of Pattern-Analyzing Skill in Amnesia: Dissociation of Knowing How and Knowing That. *Science*, **210**, 207–210.
- Collins, John. 2007. Syntax, More or Less. *Mind*, **116**, 805–850.
- Corkin, Suzanne. 2002. What's new with the amnesiac patient H.M.? *Nature Reviews Neuroscience*, **3**, 153–160.
- Cova, Florian, Dupoux, Emmanuel, & Jacob, Pierre. 2012. On Doing Things Intentionally. *Mind & Language*, **27**, 378–409.
- Craig, Edward. 1990. *Knowledge and the State of Nature. An Essay in Conceptual Synthesis*. Oxford: Clarendon Press.
- Cross, Charles, & Roelofsen, Floris. 2014. Questions. In: Zalta (1995 ff.). Spring 2014 Edition. <http://plato.stanford.edu/archives/spr2014/entries/questions/>.
- Cummings, Sam. 2013. Names. In: Zalta (1995 ff.). Spring 2013 Edition. <http://plato.stanford.edu/archives/spr2013/entries/names/>.

- Damschen, Gregor. 2009. Dispositional Knowledge-How versus Propositional Knowledge-That. In: Damschen, Gregor, Schnepf, Robert, & Stüber, Karsten R. (eds), *Debating Dispositions. Issues in Metaphysics, Epistemology and Philosophy of Mind*. Berlin & New York: Walter de Gruyter. 278–295.
- Davidson, Donald. 1963. Actions, Reasons, and Causes. In: Davidson (1980). 3–19.
- Davidson, Donald. 1971. Agency. In: Davidson (1980). 43–61.
- Davidson, Donald. 1980. *Essays on Actions and Events*. Oxford: Clarendon Press.
- Davidson, Donald. 1987. Knowing One's Own Mind. *Proceedings and Addresses of the American Philosophical Association*, **60**, 441–458.
- Davis, Wayne. 2014. Implicature. In: Zalta (1995 ff.). Fall 2014 Edition. <http://plato.stanford.edu/archives/fall2014/entries/implicature/>.
- Demmerling, Christoph. 2013. Rezension von *Gewusst wie? Eine Analyse praktischen Wissens* von Eva-Maria Jung. *Zeitschrift für philosophische Literatur*, **1**, 1–9.
- Devitt, Michael. 2011. Methodology and the Nature of Knowing How. *The Journal of Philosophy*, **108**, 205–218.
- Di Nucci, Ezio. 2011a. Automatic Actions: Challenging Causalism. *Rationality, Markets and Morals*, **2**, 179–200.
- Di Nucci, Ezio. 2011b. Frankfurt versus Frankfurt: a new anti-causalist dawn. *Philosophical Explorations*, **14**, 117–131.
- Di Nucci, Ezio. 2013. Action, Deviance, and Guidance. *Abstracta*, **7**, 41–59.
- Dickie, Imogen. 2012. Skill Before Knowledge. *Philosophy and Phenomenological Research*, **85**, 737–745.
- Dienes, Zoltan, & Perner, Josef. 1999. A theory of implicit and explicit knowledge. *Behavioral and Brain Sciences*, **22**, 735–808.
- van Dijk, Ludger, & Bongers, Raoul M. 2013. Knowledge and skill: a case for ontological equality. *Frontiers in Human Neuroscience*, **7**(916), 1–3.
- Ditter, Andreas. 2016. Why Intellectualism Still Fails. *The Philosophical Quarterly*, **66**, 500–515.
- Douskos, Christos. 2013. The linguistic argument for intellectualism. *Synthese*, **190**, 2325–2340.
- Drayson, Zoe. 2012. The Uses and Abuses of the Personal/Subpersonal Distinction. *Philosophical Perspectives*, **26**, 1–18.
- Drayson, Zoe. 2014. The Personal/Subpersonal Distinction. *Philosophy Compass*, **9**, 338–346.
- Dreyfus, Hubert L. 2002. Intelligence without representation – Merleau-Ponty's critique of mental representation. The relevance of phenomenology to scientific explanation. *Phenomenology and the Cognitive Sciences*, **1**, 367–383.
- Dreyfus, Hubert L. 2005. Overcoming the Myth of the Mental: How Philosophers can Profit from the Phenomenology of Everyday Expertise. *Proceedings and Addresses of the American Philosophical Association*, **79**, 47–65.
- Dreyfus, Hubert L. 2007. The Return of the Myth of the Mental. *Inquiry*, **50**, 352–365.

- Dreyfus, Hubert L. 2013. The myth of the pervasiveness of the mental. *In*: Schear (2013). 15–40.
- Dreyfus, Hubert L., & Dreyfus, Stuart E. 1986. *Mind over Machine. The Power of Human Intuition and Expertise in the Era of the Computer*. New York: The Free Press.
- Elgin, Catherine Z. 2009. Is Understanding Factive? *In*: Haddock *et al.* (2009). 322–329.
- Elzinga, Benjamin. 2016. Self-Regulation and Knowledge-How. *Episteme*, published online on December 15, 2016.
- Evans, Vyvyan. 2015. A unified account of polysemy within LCCM Theory. *Lingua*, **157**, 100–123.
- Falco, Andrea. 2015. Aristotle on Causality. *In*: Zalta (1995 ff.). Spring 2015 Edition. <http://plato.stanford.edu/archives/spr2015/entries/aristotle-causality/>.
- Falkum, Ingrid Lossius. 2015. The *how* and *why* of polysemy. A pragmatic account. *Lingua*, **157**, 83–99.
- Falkum, Ingrid Lossius, & Vicente, Augustin. 2015. Polysemy. Current perspectives and approaches. *Lingua*, **157**, 1–16.
- Fantl, Jeremy. 2008. Knowing-How and Knowing-That. *Philosophy Compass*, **3**, 451–470.
- Fantl, Jeremy. 2011. Ryle's regress defended. *Philosophical Studies*, **156**, 121–130.
- Fantl, Jeremy. 2012. Knowledge How. *In*: Zalta (1995 ff.). Winter 2012 Edition. <http://plato.stanford.edu/archives/win2012/entries/knowledge-how/>.
- Fara, Michael. 2008. Masked Abilities and Compatibilism. *Mind*, **117**, 843–865.
- Farkas, Katalin. 2016a. Know-wh does not reduce to know-that. *American Philosophical Quarterly*, **53**, 109–122.
- Farkas, Katalin. 2016b. Practical Know-Wh. *Noûs*, published online on June 19, 2016.
- Fine, Kit. 2012. Guide to ground. *In*: Correia, Fabrice, & Schnieder, Benjamin (eds), *Metaphysical Grounding. Understanding the Structure of Reality*. Cambridge: Cambridge University Press. 37–80.
- Fisher, Justin C. 2013. Dispositions, conditionals and auspicious circumstances. *Philosophical Studies*, **164**, 443–464.
- Fitch, Greg, & Nelson, Michael. 2013. Singular Propositions. *In*: Zalta (1995 ff.). Winter 2013 Edition. <http://plato.stanford.edu/archives/win2013/entries/propositions-singular/>.
- Flegal, Kristin E., & Anderson, Michael C. 2008. Overthinking skilled motor performance: Or why those who teach can't do. *Psychonomic Bulletin & Review*, **15**, 927–932.
- Fodor, Jerry. 1983. *The Modularity of Mind*. Cambridge: MIT Press.
- Frankfurt, Harry G. 1978. The Problem of Action. *American Philosophical Quarterly*, **15**, 157–162.
- Frankish, Keith, & Ramsey, William M. (eds). 2014. *The Cambridge Handbook of Artificial Intelligence*. Cambridge: Cambridge University Press.
- Fridland, Ellen. 2010. *Perception and Skill: Theoretical Foundations for a Science of Perception*. New York: The City University of New York, Dissertation.

- Fridland, Ellen. 2012. Knowing-How. Problems and Considerations. *European Journal of Philosophy*, published online on September 19, 2012.
- Fridland, Ellen. 2013. Problems with intellectualism. *Philosophical Studies*, **165**, 879–891.
- Fridland, Ellen. 2014a. Skill Learning and Conceptual Thought. Making a Way through the Wilderness. In: Bashour, Bana, & Muller, Hans D. (eds), *Contemporary philosophical naturalism and its implications*. New York: Routledge. 77–100.
- Fridland, Ellen. 2014b. They've lost control: reflections on skill. *Synthese*, published online on February 28, 2014.
- Fridland, Ellen. 2015. Automatically minded. *Synthese*, published online on January 30, 2015.
- Fridland, Ellen. 2016. Skill and motor control: intelligence all the way down. *Philosophical Studies*, published online on September 8, 2016.
- Frijda, Nico N. 1986. *The Emotions*. Cambridge: Cambridge University Press.
- Gabrieli, John D. E. 1998. Cognitive neuroscience of human memory. *Annual Review of Psychology*, **49**, 87–115.
- Gadamer, Hans-Georg. 1992. *Truth and Method*. Revised translation by Joel Weinsheimer and Donald G. Marshall. New York: Crossroad.
- Gettier, Edmund L. 1963. Is Justified True Belief Knowledge? *Analysis*, **23**, 121–123.
- Gibson, James J. 1977. The Theory of Affordances. In: Shaw, Robert, & Bransford, John (eds), *Perceiving, Acting, and Knowing. Toward an Ecological Psychology*. Hillsdale: Lawrence Erlbaum Associates. 67–82.
- Gibson, James J. 1979. *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin Company.
- Ginet, Carl. 1975. *Knowledge, Perception, and Memory*. Dordrecht: Reidel.
- Ginzburg, Jonathan. 2011. How to Resolve 'How To'. In: Bengson & Moffett (2011a). 215–243.
- Glauer, Ramiro. 2015. The Semantic Reading of Propositionality and Its Relation to Cognitive-Representational Explanations. A Commentary on Andreas Bartels & Mark May. In: Metzinger & Windt (2015). 2(C). 1–12.
- Glick, Ephraim. 2011. Two Methodologies for Evaluating Intellectualism. *Philosophy and Phenomenological Research*, **83**, 398–434.
- Glick, Ephraim. 2012. Abilities and Know-How Attributions. In: Brown & Gerken (2012a). 120–139.
- Glick, Ephraim. 2013. Practical Modes of Presentation. *Noûs*, published online on October 11, 2013.
- Glüer, Kathrin, & Wikforss, Åsa. 2009. The Normativity of Meaning and Content. In: Zalta (1995 ff.). Winter 2010 Edition. <http://plato.stanford.edu/archives/win2010/entries/meaning-normativity/>.
- Goldman, Alvin. 1976. Discrimination and Perceptual Knowledge. *The Journal of Philosophy*, **20**, 771–791.
- Goodale, Melvyn A., & Milner, A. David. 1992. Separate visual pathways for perception and action. *Trends in Neuroscience*, **15**, 20–25.

- Gottlieb, Gabriel. 2011. Unreflective Action and the Argument from Speed. *Pacific Philosophical Quarterly*, **92**, 338–362.
- Greco, John. 2007. The nature of ability and the purpose of knowledge. *Philosophical Issues*, **17**, 57–69.
- Greco, John. 2009. Knowledge and success from ability. *Philosophical Studies*, **142**, 17–26.
- Greco, John, & Turri, John. 2011. Virtue Epistemology. In: Zalta (1995 ff.). Winter 2011 Edition. <http://plato.stanford.edu/archives/win2011/entries/epistemology-virtue/>.
- Grice, Herbert Paul. 1989. *Studies in the Way of Words*. Cambridge: Harvard University Press.
- Grimm, Stephen R. 2006. Is Understanding a Species of Knowledge? *British Journal for the Philosophy of Science*, **57**, 515–535.
- Grimm, Stephen R. 2012. The Value of Understanding. *Philosophy Compass*, **7**, 103–117.
- Groenendijk, Jeroen, & Stokhof, Martin. 1982. The Semantic Analysis of *wh*-Complements. *Linguistics and Philosophy*, **5**, 175–233.
- Groenendijk, Jeroen, & Stokhof, Martin. 1984. *Studies on the Semantics of Questions and the Pragmatics of Answers*. Amsterdam: Universiteit van Amsterdam, Dissertation.
- Grzankowski, Alex. 2014. Attitudes Towards Objects. *Noûs*, published online on 6 June, 2014.
- Habgood-Coote, Joshua. 2017. Knowledge-How: Interrogatives and Free Relatives. *Episteme*, published online on February 27, 2017.
- Haddock, Adrian. 2010. Knowledge of Action. In: Pritchard *et al.* (2010). 241–260.
- Haddock, Adrian, Millar, Alan, & Pritchard, Duncan (eds). 2009. *Epistemic Value*. Oxford: Oxford University Press.
- Hamblin, Charles Leonard. 1958. Questions. *Australasian Journal of Philosophy*, **36**, 159–168.
- Han, ZhaoHong. 2004. Fossilization: five central issues. *International Journal of Applied Linguistics*, **14**, 212–242.
- Han, ZhaoHong. 2012. Fossilization. In: Chapelle, Carol A. (ed), *The Encyclopedia of Applied Linguistics*. Chichester: Wiley-Blackwell. 2205–2210.
- Hartland-Swann, John. 1956. The Logical Status of ‘Knowing That’. *Analysis*, **16**, 111–115.
- Hartland-Swann, John. 1957. ‘Knowing That’: A Reply to Mr. Ammerman. *Analysis*, **17**, 69–71.
- Hawley, Katherine. 2003. Success and Knowledge-How. *American Philosophical Quarterly*, **40**, 19–31.
- Hawley, Katherine. 2010. Testimony and knowing how. *Studies in History and Philosophy of Science*, **41**, 397–404.
- Heinrich, Bernd. 2002. Raven Consciousness. In: Bekoff *et al.* (2002). 47–52.
- Henning, Boris. 2009. The Four Causes. *The Journal of Philosophy*, **106**, 137–160.

- Hermann, Louis S. 2002. Exploring the Cognitive World of the Bottlenosed Dolphin. *In: Bekoff et al.* (2002). 275–83.
- Hetherington, Stephen. 2006. How to Know (that Knowledge-that is Knowledge-how). *In: Hetherington, Stephen* (ed), *Epistemology Futures*. Oxford: Oxford University Press. 71–94.
- Hetherington, Stephen. 2008a. Knowing-That, Knowing-How, and Knowing Philosophically. *Grazer Philosophische Studien*, **77**, 307–324.
- Hetherington, Stephen. 2008b. Not Actually Hume’s Problem. On Induction and Knowing-How. *Philosophy*, **83**, 459–481.
- Hetherington, Stephen. 2011a. *How to Know. A Practicalist Conception of Knowledge*. Chichester: Wiley-Blackwell.
- Hetherington, Stephen. 2011b. Knowledge and Knowing. Ability and Manifestation. *In: Tolksdorf* (2011). 73–99.
- Hintikka, Jaako. 1975. Different Constructions in Terms of the Basic Epistemological Verbs: A Survey of Some Problems and Proposals. *In: The intentions of intentionality and other new models for modalities*. Dordrecht: D. Reidel. 1–25.
- Hintikka, Jaako. 1992. Different constructions in terms of ‘knows’. *In: Dancy, Jonathan, & Sosa, Ernest* (eds), *A Companion to Epistemology*. Oxford: Blackwell. 99–104.
- Hornsby, Jennifer. 2000. Personal and Sub-Personal. A Defence of Dennett’s Early Distinction. *Philosophical Explorations*, **1**, 6–24.
- Hornsby, Jennifer. 2005. Semantic Knowledge and Practical Knowledge. *Proceedings of the Aristotelian Society. Supplementary Volume*, **79**, 107–130.
- Hornsby, Jennifer. 2011. Ryle’s *Knowing-How*, and Knowing How to Act. *In: Bengson & Moffett* (2011a). 80–98.
- Hornsby, Jennifer. 2012. Actions and Activity. *Philosophical Issues*, **22**, 233–245.
- Hornsby, Jennifer. 2016. Intending, knowing how, infinitives. *Canadian Journal of Philosophy*, **46**, 1–17.
- Howell, Robert J. 2011. The Knowledge Argument and the Implications of Phenomenal Knowledge. *Philosophy Compass*, **6**, 459–468.
- Hyman, John. 1999. How Knowledge Works. *The Philosophical Quarterly*, **49**, 433–451.
- Ichikawa, Jonathan Jenkins, & Steup, Matthias. 2012. The Analysis of Knowledge. *In: Zalta* (1995 ff.). Spring 2014 Edition. <http://plato.stanford.edu/archives/spr2014/entries/knowledge-analysis/>.
- Jackson, Frank. 1982. Epiphenomenal Qualia. *The Philosophical Quarterly*, **32**, 127–136.
- Jackson, Frank. 1986. What Mary Didn’t Know. *The Journal of Philosophy*, **83**, 291–295.
- Jaster, Romy. 2016. *Agents’ Abilities*. Berlin: Humboldt-Universität zu Berlin, Dissertation.
- Johnson, Kent. 2006. Externalist Thoughts and the Scope of Linguistics. *Protosociology*, **22**, 19–35.

- Johnston, Mark. 1992. How to Speak of the Colors. *Philosophical Studies*, **68**, 221–263.
- Jung, Eva-Maria. 2012. *Gewusst wie? Eine Analyse praktischen Wissens*. Berlin & Boston: Walter de Gruyter.
- Jung, Eva-Maria, & Newen, Albert. 2010. Knowledge and abilities: The need for a new understanding of knowing-how. *Phenomenology and the Cognitive Sciences*, **9**, 113–131.
- Jung, Eva-Maria, & Newen, Albert. 2011. Understanding Knowledge in a New Framework. Against Intellectualism as a Semantic Analysis and as an Analysis of the Mind. In: Newen *et al.* (2011). 79–105.
- Kallestrup, Jesper. 2009. Knowledge-*wh* and the Problem of Convergent Knowledge. *Philosophy and Phenomenological Research*, **78**, 468–476.
- Karttunen, Lauri. 1977. Syntax and Semantics of Questions. *Linguistics and Philosophy*, **1**, 3–44.
- Katzoff, Charlotte. 1984. Knowing How. *Southern Journal of Philosophy*, **22**, 61–69.
- Kauppinen, Antti. 2007. The Rise and Fall of Experimental Philosophy. *Philosophical Explorations*, **10**, 95–118.
- Keil, Geert. 2003. Über den Homunkulus-Fehlschluß. *Zeitschrift für philosophische Forschung*, **57**, 1–26.
- Kelly, Sean Dorrance. 2002. Merlau-Ponty on the Body. *Ratio*, **15**, 376–391.
- Kenny, Anthony. 1975. *Will, Freedom, and Power*. Oxford: Blackwell.
- Kenny, Anthony. 1991. The Homunculus Fallacy. In: Hyman, John (ed), *Investigating Psychology. Sciences of the Mind after Wittgenstein*. London & New York: Routledge.
- Kenny, Anthony. 2010. Concepts, Brains, and Behaviour. *Grazer Philosophische Studien*, **81**, 105–113.
- King, Jeffrey C. 2011. Structured Propositions. In: Zalta (1995 ff.). Winter 2012 Edition. <http://plato.stanford.edu/archives/win2012/entries/propositions-structured/>.
- Knobe, Joshua. 2007. Experimental Philosophy. *Philosophy Compass*, **2**, 81–92.
- “know”. 2014a. *Merriam-Webster.com Dictionary*. Retrieved May 3, 2014. <http://www.merriam-webster.com/dictionary/know>.
- “know”. 2014b. *Merriam-Webster.com Thesaurus*. Retrieved May 3, 2014. <http://www.merriam-webster.com/thesaurus/know>.
- “know-how”. 2014a. *Merriam-Webster.com Dictionary*. Retrieved May 3, 2014. <http://www.merriam-webster.com/dictionary/know-how>.
- “know-how”. 2014b. *Merriam-Webster.com Thesaurus*. Retrieved May 3, 2014. <http://www.merriam-webster.com/thesaurus/know-how>.
- “knowledge”. 2014a. *Merriam-Webster.com Dictionary*. Retrieved May 3, 2014. <http://www.merriam-webster.com/dictionary/knowledge>.
- “knowledge”. 2014b. *Merriam-Webster.com Thesaurus*. Retrieved May 3, 2014. <http://www.merriam-webster.com/thesaurus/knowledge>.
- Koethe, John. 2002. Stanley and Williamson on Knowing How. *The Journal of Philosophy*, **99**, 325–328.

- Koffka, Kurt. 1935. *Principles of Gestalt Psychology*. London: Lund Humphries.
- Kotzee, Ben. 2016. Learning How. *Journal of Philosophy of Education*, **50**, 218–232.
- Kratzer, Angelika. 1977. What ‘Must’ and ‘Can’ Must and Can Mean. *Linguistics and Philosophy*, **1**, 337–355.
- Kremer, Michael. 2016. A Capacity to Get Things Right: Gilbert Ryle on Knowledge. *European Journal of Philosophy*, published online on October 19, 2016.
- Krifka, Manfred. 2011. Questions. In: von Stechow, Claudia, Heusinger, Klaus, & Portner, Paul (eds), *Semantics. An International Handbook of Natural Language*. Berlin & New York: Mouton de Gruyter. 1785–1804.
- Kripke, Saul. 1977. Speaker’s Reference and Semantic Reference. *Midwest Studies in Philosophy*, **2**, 255–276.
- Kripke, Saul. 1982. *Wittgenstein on Rules and Private Language. An Elementary Exposition*. Oxford: Basil Blackwell.
- Kumar, Victor. 2011. In support of anti-intellectualism. *Philosophical Studies*, **152**, 135–154.
- Kvanvig, Jonathan L. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- Laurence, Stephen, & Margolis, Eric. 1999. Concepts and Cognitive Science. In: Margolis, Eric, & Laurence, Stephen (eds), *Concepts. Core Readings*. Cambridge: MIT Press. 3–81.
- Lavin, Douglas. 2013. Must There Be Basic Action? *Noûs*, **47**, 273–301.
- Lehrer, Keith. 1990. *Theory of Knowledge*. Boulder & San Francisco: Westview Press.
- Levin, Janet. 1986. Could Love Be Like a Heatwave? Physicalism and the Subjective Character of Experience. *Philosophical Studies*, **49**, 245–261.
- Levin, Janet. 2013. Functionalism. In: Zalta (1995 ff.). Fall 2013 Edition. <http://plato.stanford.edu/archives/fall2013/entries/functionalist/>.
- Lewandowska-Tomaszczyk, Barbara. 2007. Polysemy, Prototypes, and Radial Categories. In: Geeraerts, Dirk, & Cuyckens, Hubert (eds), *The Oxford Handbook of Cognitive Linguistics*. Oxford: Oxford University Press. 139–169.
- Lewin, Kurt. 1935. *A Dynamic Theory of Personality: Selected Papers*. New York: McGraw-Hill.
- Lewis, David. 1979. Attitudes De Dicto and De Se. *The Philosophical Review*, **88**, 513–543.
- Lewis, David. 1997. Finkish Dispositions. *The Philosophical Quarterly*, **47**, 143–158.
- Lihoreau, Franck. 2008. Knowledge-How and Ability. *Grazer Philosophische Studien*, **77**, 263–305.
- Luce, Gay Gaer, & Segal, Julius. 1966. *Sleep*. New York: Coward McCann.
- Luntley, Michael. 2009. Understanding Expertise. *Journal of Applied Philosophy*, **26**, 356–370.
- Löwenstein, David. 2011a. Knowledge-How, Linguistic Intellectualism, and Ryle’s Return. In: Tolksdorf (2011). 269–304.

- Löwenstein, David. 2011b. The Social Status of Knowledge-how. Paper presented at the “Second Berlin Conference on Meta–Epistemology. Social Epistemology”, Technische Universität Berlin.
- Löwenstein, David. 2013. Why Know-how and Propositional Knowledge are Mutually Irreducible. In: Hoeltje, Miguel, Spitzley, Thomas, & Spohn, Wolfgang (eds), *Was dürfen wir glauben? Was sollen wir tun? Sektionsbeiträge des achten internationalen Kongresses der Gesellschaft für Analytische Philosophie e.V.* Duisburg & Essen: Universität Duisburg-Essen (DuEPublico). 365–371.
- Löwenstein, David. 2016. A Uniform Account of Regress Problems. *Acta Analytica*, published online on October 24, 2016.
- Madison, Brent James Charles. 2012. Review of *How to Know: A Practicalist Conception of Knowledge* by Stephen Hetherington. *Notre Dame Philosophical Reviews*, February 15, 2012. <http://ndpr.nd.edu/news/28899-how-to-know-a-practicalist-conception-of-knowledge/>.
- Maier, John. 2010. Abilities. In: Zalta (1995 ff.). Fall 2011 Edition. <http://plato.stanford.edu/archives/fall2011/entries/abilities/>.
- Maier, John. 2013. The Agentive Modalities. *Philosophy and Phenomenological Research*, published online on May 6, 2013.
- Malle, Bertram F., & Knobe, Joshua. 1997. The Folk Concept of Intentionality. *Journal of Experimental Social Psychology*, **33**, 101–121.
- Manley, David, & Wasserman, Ryan. 2008. On Linking Dispositions and Conditionals. *Mind*, **117**, 59–84.
- Margolis, Eric, & Laurence, Stephen. 2011. Concepts. In: Zalta (1995 ff.). Fall 2012 Edition. <http://plato.stanford.edu/archives/fall2012/entries/concepts/>.
- Martin, Charles Burton. 1994. Dispositions and Conditionals. *The Philosophical Quarterly*, **44**, 1–8.
- Masto, Meghan. 2010. Questions, Answers, and Knowledge-Wh. *Philosophical Studies*, **147**, 395–413.
- Matsuzawa, Tetsuro. 2002. Chimpanzee Ai and her Son Ayumu: An Episode of Education by Master-Apprenticeship. In: Bekoff *et al.* (2002). 189–198.
- Matthews, Gareth, & Mullin, Amy. 2015. The Philosophy of Childhood. In: Zalta (1995 ff.). Spring 2015 Edition. <http://plato.stanford.edu/archives/spr2015/entries/childhood/>.
- McDowell, John. 1994. *Mind and World*. Cambridge & London: Harvard University Press.
- McDowell, John. 2007. What Myth? *Inquiry*, **50**, 338–351.
- McDowell, John. 2013. The myth of the mind as detached. In: Schear (2013). 41–58.
- McGrath, Matthew. 2012. Propositions. In: Zalta (1995 ff.). Summer 2012 Edition. <http://plato.stanford.edu/archives/sum2012/entries/propositions/>.
- McKay, Thomas, & Nelson, Michael. 2010. Propositional Attitude Reports. In: Zalta (1995 ff.). Winter 2010 Edition. <http://plato.stanford.edu/archives/win2010/entries/prop-attitude-reports/>.
- Mellor, D. Hugh. 1993. Nothing Like Experience. *Proceedings of the Aristotelian Society*, **93**, 1–16.

- Menary, Richard. 2009. Intentionality and Consciousness. *In*: Banks, William P. (ed), *Encyclopedia of Consciousness*. Oxford: Elsevier. 417–429.
- Metzinger, Thomas, & Windt, Jennifer M. (eds). 2015. *Open MIND*. Frankfurt am Main: MIND Group.
- Michaelis, Laura A. 2011. Knowledge Ascription by Grammatical Construction. *In*: Bengson & Moffett (2011a). 261–279.
- Michaels, Claire F. 2003. Affordances: Four Points of Debate. *Ecological Psychology*, **15**, 135–148.
- Miller, Alexander, & Wright, Crispin (eds). 2002. *Rule-Following and Meaning*. Montreal & Kingston: McGill-Queen's University Press.
- Millikan, Ruth Garrett. 1984. *Language, Thought, and other Biological Categories. New Foundations for Realism*. Cambridge: MIT Press.
- Millikan, Ruth Garrett. 2000. *On Clear and Confused Ideas. An Essay about Substance Concepts*. Cambridge: Cambridge University Press.
- Montero, Barbara. 2010. Does bodily awareness interfere with highly skilled movement? *Inquiry*, **53**, 105–122.
- Montero, Barbara. 2013. A Dancer Reflects. *In*: Schear (2013). 303–319.
- Nadelhoffer, Thomas. 2005. Skill, Luck, Control, and Intentional Action. *Philosophical Psychology*, **18**, 341–352.
- Nagel, Jennifer. 2012. Intuitions and Experiments. A Defense of the Case Method in Epistemology. *Philosophy and Phenomenological Research*, **85**, 495–527.
- Nagel, Thomas. 1974. What is it Like to be a Bat? *Philosophical Review*, **83**, 435–450.
- Nemirow, Laurence. 1990. Physicalism and the Cognitive Role of Acquaintance. *In*: Lycan, William G. (ed), *Mind and cognition. A reader*. Cambridge: Blackwell. 490–499.
- Newen, Albert, Bartels, Andreas, & Jung, Eva-Maria (eds). 2011. *Knowledge and Representation*. Stanford & Paderborn: CSLI Publications & Mentis.
- Nida-Rümelin, Martine. 2009. Qualia: The Knowledge Argument. *In*: Zalta (1995 ff.). Summer 2010 Edition. <http://plato.stanford.edu/archives/sum2010/entries/qualia-knowledge/>.
- Ninan, Dilip. 2010. *De Se* Attitudes: Ascription and Communication. *Philosophy Compass*, **5**, 551–567.
- Nolan, Daniel. 2001. What's Wrong With Infinite Regresses? *Metaphilosophy*, **32**, 523–538.
- Noë, Alva. 2005. Against intellectualism. *Analysis*, **65**, 278–290.
- Noë, Alva. 2011. Ideology and the Third Realm (Or, A Short Essay on Knowing How to Philosophize). *In*: Bengson & Moffett (2011a). 196–211.
- Orozco, Joshue. 2011. Epistemic Luck. *Philosophy Compass*, **6**, 11–21.
- Parent, Ted. 2014. Knowing-*Wh* and Embedded Questions. *Philosophy Compass*, **9**, 81–95.
- Park, Shelley M. 1994. Reinterpreting Ryle: A Nonbehavioristic Analysis. *Journal of the History of Philosophy*, **32**, 265–290.
- Parkinson, Brian. 1995. *Ideas and Realities of Emotions*. New York: Routledge.

- Parry, Richard D. 1980. Ryle's Theory of Action in *The Concept of Mind. Philosophy and Phenomenological Research*, **40**, 379–392.
- Pavese, Carlotta. 2015a. Knowing a Rule. *Philosophical Issues*, **25**, 165–188.
- Pavese, Carlotta. 2015b. Practical Senses. *Philosophers' Imprint*, **15**, 1–25.
- Pavese, Carlotta. 2016a. Skill in epistemology I: Skill and knowledge. *Philosophy Compass*, **12**, 642–649.
- Pavese, Carlotta. 2016b. Skill in epistemology II: Skill and know how. *Philosophy Compass*, **12**, 650–660.
- Pavese, Carlotta. 2017. Know-how and Gradability. *Philosophical Review*, forthcoming.
- Perry, John. 1977. Frege on Demonstratives. *The Philosophical Review*, **86**, 474–497.
- Perry, John. 1979. The Problem of the Essential Indexical. *Noûs*, **13**, 3–21.
- Poston, Ted. 2009. Know how to be Gettiered? *Philosophy and Phenomenological Research*, **79**, 743–747.
- Poston, Ted. 2016. Know How to Transmit Knowledge? *Noûs*, **50**, 865–878.
- Pritchard, Duncan. 2005. *Epistemic Luck*. Oxford: Oxford University Press.
- Pritchard, Duncan. 2010. Understanding. In: Pritchard *et al.* (2010). 66–88.
- Pritchard, Duncan. 2012. Anti-Luck Virtue Epistemology. *The Journal of Philosophy*, **109**, 247–279.
- Pritchard, Duncan, Millar, Alan, & Haddock, Adrian (eds). 2010. *The Nature and Value of Knowledge. Three Investigations*. Oxford: Oxford University Press.
- Prochazka, Arthur, Clarac, François, Loeb, Gerald E., Rothwell, John C., & Wolpaw, Jonathan R. 2000. What do *reflex* and *voluntary* mean? Modern views on an ancient debate. *Experimental Brain Research*, **130**, 417–430.
- Railton, Peter. 2004. How to Engage Reason: The Problem of Regress. In: Wallace, R. Jay, Pettit, Philipp, Scheffler, Samuel, & Smith, Michael (eds), *Reason and Value. Themes from the Moral Philosophy of Joseph Raz*. Oxford: Oxford University Press. 176–201.
- Railton, Peter. 2006. Normative Guidance. In: Shafer-Landau, Russ (ed), *Oxford Studies in Metaethics. Volume 1*. Oxford: Oxford University Press. 3–33.
- Railton, Peter. 2009. Practical competence and fluent agency. In: Sobel, David, & Wall, Steven (eds), *Reasons for Action*. Cambridge: Cambridge University Press. 81–115.
- Reber, Paul J., Knowlton, Barbara J., & Squire, Larry R. 1996. Dissociable Properties of Memory Systems: Differences in the Flexibility of Declarative and Nondeclarative Knowledge. *Behavioral Neuroscience*, **110**, 861–871.
- Reed, Nick, McLeod, Peter, & Dienes, Zoltan. 2010. Implicit knowledge and motor skill: What people who know how to catch don't know. *Consciousness and Cognition*, **19**, 63–76.
- Reisberg, Daniel. 1999. Learning. In: Wilson & Keil (1999). 460–461.
- Ren, Huiming. 2012. The Distinction between Knowledge-That and Knowledge-How. *Philosophia*, **40**, 857–875.
- Rey, Georges. 2002. Problems with Dreyfus' dialectic. *Phenomenology and the Cognitive Sciences*, **1**, 403–408.

- Rietveld, Erik. 2008. Situated Normativity: The Normative Aspect of Embodied Cognition in Unreflective Action. *Mind*, **117**, 973–1001.
- Rietveld, Erik. 2010. McDowell and Dreyfus on Unreflective Action. *Inquiry*, **53**, 183–207.
- Riggs, Wayne D. 2009. Understanding, Knowledge, and the Meno Requirement. *In: Haddock et al.* (2009). 331–338.
- Roberts, Craige. 2009. *know-how*: A compositional approach. *In: Hinrichs, Erhard, & Nerbonne, John* (eds), *Theory and Evidence in Semantics*. Stanford: CSLI Press. 183–213.
- Roeber, Blake. 2016. The Pragmatic Encroachment Debate. *Noûs*, published online in June, 2016.
- Rosefeldt, Tobias. 2004. Is Knowing-how Simply a Case of Knowing-that? *Philosophical Investigations*, **27**, 370–379.
- Roth, Martin, & Cummins, Robert. 2011. Intellectualism as Cognitive Science. *In: Newen et al.* (2011). 23–40.
- Rumfitt, Ian. 2003. Savoir Faire. *The Journal of Philosophy*, **100**, 158–166.
- Ryle, Gilbert. 1930. Are There Propositions? *Proceedings of the Aristotelian Society*, **30**, 91–126.
- Ryle, Gilbert. 1945a. Knowing How and Knowing That. *Proceedings of the Aristotelian Society*, **46**, 1–16.
- Ryle, Gilbert. 1945b. *Philosophical Arguments*. Oxford: Clarendon Press.
- Ryle, Gilbert. 1949. *The Concept of Mind*. Edition of 1963. Harmondsworth: Penguin Books.
- Ryle, Gilbert. 1950. “If,” “So,” and “Because”. *In: Black, Max* (ed), *Philosophical Analysis*. Ithaca: Cornell University Press. 323–340.
- Rysiew, Patrick. 2011. Epistemic Contextualism. *In: Zalta* (1995 ff.). Winter 2011 Edition. <http://plato.stanford.edu/archives/win2011/entries/contextualism-epistemology/>.
- Rödl, Sebastian. 2010. The Self-Conscious Power of Sensory Knowledge. *Grazer Philosophische Studien*, **81**, 135–151.
- Sag, Ivan A., Gazdar, Gerald, Wasow, Thomas, & Weisler, Steven. 1985. Coordination and how to distinguish categories. *Natural Language and Linguistic Theory*, **3**, 117–171.
- Sandis, Constantine. 2012. The Objects of Action Explanation. *Ratio*, **25**, 326–344.
- Sanford, David H. 1984. Infinite Regress Arguments. *In: Fetzer, James H.* (ed), *Principles of Philosophical Reasoning*. Totowa: Rowman & Allanheld. 93–117.
- Santorio, Paolo. 2016. Nonfactual Know-How and the Boundaries of Semantics. *Philosophical Review*, **125**, 35–82.
- Sax, Greg. 2010. Having Know-how. Intellect, Action, and Recent Work on Ryle’s Distinction between Knowledge-How and Knowledge-That. *Pacific Philosophical Quarterly*, **91**, 507–530.
- Schaffer, Jonathan. 2007. Knowing the Answer. *Philosophy and Phenomenological Research*, **75**, 383–403.

- Schaffer, Jonathan. 2009a. Knowing the Answer Redux: Replies to Brogaard and Kallestrup. *Philosophy and Phenomenological Research*, **78**, 477–500.
- Schaffer, Jonathan. 2009b. On What Grounds What. In: Chalmers, David J., Manley, David, & Wasserman, Ryan (eds), *Metametaphysics. New Essays on the Foundations of Ontology*. Oxford: Clarendon Press. 347–383.
- Schear, Joseph K. (ed). 2013. *Mind, Reason, and Being-in-the-World. The McDowell-Dreyfus Debate*. London & New York: Routledge.
- Scheffler, Israel. 1968. On Ryle's Theory of Propositional Knowledge. *The Journal of Philosophy*, **65**, 725–732.
- Schiffer, Stephen. 2002. Amazing Knowledge. *The Journal of Philosophy*, **99**, 200–202.
- Schneider, Walter. 1999. Automaticity. In: Wilson & Keil (1999). 63–64.
- Schroeder, Mark. 2012. Showing How to Derive Knowing How. *Philosophy and Phenomenological Research*, **85**, 746–753.
- Schröder, Dirk. 2013. Rezension von *Know How* von Jason Stanley. *Zeitschrift für philosophische Literatur*, **1**, 10–18.
- Schulmann, Robert, Knox, A. J., Jansen, Michel, & Illy, József (eds). 1998. *The collected Papers of Albert Einstein. Volume 8: The Berlin Years: Correspondance, 1914-1918. Part B: 1918*. Princeton: Princeton University Press.
- Schwarzschild, Roger. 2008. The Semantics of Comparatives and Other Degree Constructions. *Language and Linguistics Compass*, **22**, 308–331.
- Schwenkler, John. 2012. Non-Observational Knowledge of Action. *Philosophy Compass*, **7**, 731–740.
- Schwitzgebel, Eric. 2002. A Phenomenal, Dispositional Account of Belief. *Noûs*, **36**, 249–275.
- Schwitzgebel, Eric. 2010. Belief. In: Zalta (1995 ff.). Winter 2011 Edition. <http://plato.stanford.edu/archives/win2011/entries/belief/>.
- Scoville, William Beecher, & Milner, Brenda. 1957. Loss of recent memory after bilateral hippocampal lesions. *The Journal of Neurology, Neurosurgery and Psychiatry*, **20**, 11–21.
- Searle, John Rogers. 1969. *Speech Acts. An Essay in the Philosophy of Language*. Cambridge: Cambridge University Press.
- Searle, John Rogers. 1995. *The Construction of Social Reality*. New York: The Free Press.
- Sellars, Wilfrid. 1962. Philosophy and the Scientific Image of Man. In: Colodny, Robert G. (ed), *Frontiers of Science and Philosophy*. Pittsburgh: University of Pittsburgh Press. 37–78.
- Sellars, Wilfrid. 1968. *Science and Metaphysics. Variations on Kantian Themes*. London: Routledge & Kegan Paul Ltd.
- Sellars, Wilfrid. 1969. Language as Thought and as Communication. *Philosophy and Phenomenological Research*, **29**, 506–527.
- Sennet, Adam. 2016. Ambiguity. In: Zalta (1995 ff.). Spring 2016 Edition. <http://plato.stanford.edu/archives/spr2016/entries/ambiguity/>.
- Setiya, Kieran. 2012. Knowing How. *Proceedings of the Aristotelian Society*, **112**, 285–307.

- Setiya, Kieran. 2014. Intention. *In: Zalta* (1995 ff.). Spring 2014 Edition. <http://plato.stanford.edu/archives/spr2014/entries/intention/>.
- Sgaravatti, Daniele, & Zardini, Elia. 2008. Knowing How to Establish Intellectualism. *Grazer Philosophische Studien*, **77**, 217–261.
- Shope, Robert K. 1978. The Conditional Fallacy in Contemporary Philosophy. *The Journal of Philosophy*, **75**, 397–413.
- Small, Will. 2014. The Transmission of Skill. *Philosophical Topics*, **42**, 85–111.
- Smith, Arthur David. 1977. Dispositional Properties. *Mind*, **86**, 439–445.
- Smythe, William E. 1988. The Conceptual Distinction between Propositional and Skillful Knowing. *In: Baker, William J., Mos, Leendert P., Rappard, Hans V., & Stam, Henderikus J. (eds), Recent Trends in Theoretical Psychology*. New York, Berlin & Heidelberg: Springer.
- Snowdon, Paul F. 2003. Knowing How and Knowing That: A Distinction Reconsidered. *Proceedings of the Aristotelian Society*, **105**, 1–25.
- Snowdon, Paul F. 2011. Rylean Arguments: Ancient and Modern. *In: Bengson & Moffett* (2011a). 59–79.
- Sosa, Ernest. 2007a. *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. Oxford: Oxford University Press.
- Sosa, Ernest. 2007b. Experimental philosophy and philosophical intuition. *Philosophical Studies*, **132**, 99–107.
- Sosa, Ernest. 2009. Knowing Full Well: The Normativity of Beliefs as Performances. *Philosophical Studies*, **142**, 5–15.
- Sosa, Ernest. 2010. How competence matters in epistemology. *Philosophical Perspectives*, **24**, 465–475.
- Sosa, Ernest. 2015. *Judgment and Agency*. Oxford: Oxford University Press.
- Srinivasan, Mahesh, & Rabagliati, Hugh. 2015. How concepts and conventions structure the lexicon: Cross-linguistic evidence from polysemy. *Lingua*, **157**, 124–152.
- Stalnaker, Robert. 1987. *Inquiry*. Cambridge: MIT Press.
- Stalnaker, Robert. 2012. Intellectualism and the Objects of Knowledge. *Philosophy and Phenomenological Research*, **85**, 754–761.
- Stanley, Jason. 2005a. Hornsby on the Phenomenology of Speech. *Proceedings of the Aristotelian Society. Supplementary Volume*, **79**, 131–145.
- Stanley, Jason. 2005b. *Knowledge and Practical Interests*. Oxford: Clarendon Press.
- Stanley, Jason. 2011a. Intellectualism and the Language of Thought. A Reply to Roth and Cummins. *In: Newen et al.* (2011). 41–50.
- Stanley, Jason. 2011b. *Know How*. Oxford: Oxford University Press.
- Stanley, Jason. 2011c. Knowing (How). *Noûs*, **45**, 207–238.
- Stanley, Jason. 2012a. Précis of *Know How*. *Philosophy and Phenomenological Research*, **85**, 733–736.
- Stanley, Jason. 2012b. Replies to Dickie, Schroeder and Stalnaker. *Philosophy and Phenomenological Research*, **85**, 762–778.
- Stanley, Jason. 2012c. The Practical and the Theoretical. *The New York Times*, May 6, 2012.

- Stanley, Jason, & Krakauer, John W. 2013. Motor Skill Depends on Knowledge of Facts. *Frontiers in Human Neuroscience*, **7**(503), 1–11.
- Stanley, Jason, & Williamson, Timothy. 2001. Knowing How. *The Journal of Philosophy*, **98**, 411–444.
- Stanley, Jason, & Williamson, Timothy. 2016. Skill. *Noûs*, published online on May 21, 2016.
- Steel, Thomas J. 1974. A Puzzle about Knowing How. *Philosophical Studies*, **25**, 43–50.
- Stillings, Neil, Weisler, Steven, Chase, Christopher, Feinstein, Mark, Garfield, Jay, & Rissland, Edwina. 1995. *Cognitive Science: An Introduction*. Cambridge: MIT Press.
- Stout, Rowland. 2003. Ryle's Behaviourism. *Revue Internationale de Philosophie*, **223**, 37–49.
- Stout, Rowland. 2010. What You Know When You Know an Answer to a Question. *Noûs*, **44**, 392–402.
- Sutton, John, McIlwain, Doris, Christensen, Wayne, & Geeves, Andrew. 2011. Applying Intelligence to the Reflexes: Embodied Skills and Habits between Dreyfus and Descartes. *Journal of the British Society for Phenomenology*, **42**, 78–103.
- Sytsma, Justin, & Buckwalter, Wesley (eds). 2016. *A Companion to Experimental Philosophy*. Chichester: Wiley-Blackwell.
- Tanney, Julia. 2009. Gilbert Ryle. In: Zalta (1995 ff.). Winter 2009 Edition. <http://plato.stanford.edu/archives/win2009/entries/ryle/>.
- Tanney, Julia. 2011. Ryle's Regress and the Philosophy of Cognitive Science. In: *Rules, Reason, and Self-Knowledge*. Cambridge: Harvard University Press. 249–276.
- Tolksdorf, Stefan (ed). 2011. *Conceptions of Knowledge*. Berlin & New York: Walter de Gruyter.
- Toner, John. 2014. Knowledge of facts mediate “continuous improvement” in elite sport: a comment on Stanley and Krakauer (2013). *Frontiers in Human Neuroscience*, **8**(142), 1–2.
- Toribio, Josefa. 2008. How Do We Know How? *Philosophical Explorations*, **11**, 39–52.
- Traugott, Elizabeth Closs. 1989. On the Rise of Epistemic Meaning in English: an Example of Subjectification in Semantic Change. *Language*, **65**, 31–55.
- Tsai, Cheng-hung. 2011a. Linguistic Know-How: The Limits of Intellectualism. *Theoria*, **77**, 71–86.
- Tsai, Cheng-hung. 2011b. The metaepistemology of knowing-how. *Phenomenology and the Cognitive Sciences*, **10**, 541–556.
- Tsai, Cheng-hung. 2013. The Structure of Practical Expertise. *Philosophia*, published online on December 28, 2013.
- Tulving, Endel. 1999. Episodic vs. Semantic Memory. In: Wilson & Keil (1999). 278–279.
- Tye, Michael. 2011. Knowing What It Is Like. In: Bengson & Moffett (2011a). 300–313.

- Umiltà, Carlo. 2007. Consciousness and Control of Action. *In: Zelazo, Philip David, Moscovitch, Morris, & Thompson, Evan (eds), The Cambridge Handbook of Consciousness.* Cambridge: Cambridge University Press. 327–351.
- Van Gulick, Robert. 2014. Consciousness. *In: Zalta (1995 ff.).* Spring 2014 Edition. <http://plato.stanford.edu/archives/spr2014/entries/consciousness/>.
- Vetter, Barbara. 2013. Multi-Track Dispositions. *The Philosophical Quarterly*, **63**, 330–352.
- Vetter, Barbara. 2014. Dispositions without Conditionals. *Mind*, published online on April 15, 2014.
- Vetter, Barbara. 2015. *Potentiality: From Dispositions to Modality.* Oxford: Oxford University Press.
- Vetter, Barbara. 2016. Are abilities dispositions? *Synthese*, published online on July 12, 2016.
- Viebahn, Emanuel. 2016. Arguing against ambiguity: the zeugma test. Manuscript.
- Viebahn, Emanuel, & Vetter, Barbara. 2016. How many meanings for ‘may’? The case for modal polysemy. *Philosophers’ Imprint*, **10**, 1–26.
- Vivhelin, Kadri. 2004. Free Will Demystified: A Dispositional Account. *Philosophical Topics*, **32**, 427–450.
- Wallis, Charles. 2008. Consciousness, context, and know-how. *Synthese*, **160**, 123–153.
- Watzl, Sebastian. 2011a. The Nature of Attention. *Philosophy Compass*, **6**, 842–853.
- Watzl, Sebastian. 2011b. The Philosophical Significance of Attention. *Philosophy Compass*, **6**, 722–733.
- Weatherson, Brian. 2016. Intellectual Skill and the Rylean Regress. *The Philosophical Quarterly*, published online on July 27, 2016.
- Whittle, Ann. 2010. Dispositional Abilities. *Philosophers’ Imprint*, **10**, 1–23.
- Wieland, Jan Willem. 2012. Regress Argument Reconstruction. *Argumentation*, **26**, 489–503.
- Wieland, Jan Willem. 2013. Infinite Regress Arguments. *Acta Analytica*, **28**, 95–109.
- Wiggins, David. 2009. Knowing How To and Knowing That. *In: Glock, Hans-Johann, & Hyman, John (eds), Wittgenstein and Analytic Philosophy: Essays for P. M. S. Hacker.* Oxford: Oxford University Press. 263–277.
- Wiggins, David. 2012. Practical Knowledge: Knowing How To and Knowing That. *Mind*, **121**, 97–130.
- Williams, John N. 2008. Propositional knowledge and know-how. *Synthese*, **165**, 107–125.
- Williams, Michael. 2008. Responsibility and Reliability. *Philosophical Papers*, **37**, 1–26.
- Wilson, George, & Shpall, Samuel. 2012. Action. *In: Zalta (1995 ff.).* Summer 2012 Edition. <http://plato.stanford.edu/archives/sum2012/entries/action/>.
- Wilson, Robert A., & Keil, Frank C. (eds). 1999. *The MIT Encyclopedia of the Cognitive Sciences.* Cambridge: MIT Press.

- Wittgenstein, Ludwig. 1953. *Philosophische Untersuchungen. Philosophical Investigations*. The German text, with an English translation by G.E.M. Anscombe, P.M.S. Hacker and Joachim Schulte. Revised 4th edition (2009). Chichester: Wiley-Blackwell.
- Wittgenstein, Ludwig. 1977a. *Vermischte Bemerkungen*. von Wright, Georg Henrik (ed). Revised 2nd edition (1978). In: Werkausgabe Band 8 (1984). Frankfurt am Main: Suhrkamp. 445–575.
- Wittgenstein, Ludwig. 1977b. *Culture and Value. A Selection from the Posthumous Remains*. von Wright, Georg Henrik, & Pichler, Alois (eds). Revised 2nd edition with English translation by Peter Winch (1998). Oxford: Blackwell.
- von Wright, Georg Henrik. 1988. An Essay on Door-Knocking. *Rechtstheorie*, **19**, 275–288.
- Wurmbrand, Susanne. 2001. *Infinitives: Restructuring and Clause Structure*. Berlin & New York: Mouton de Gruyter.
- Young, Garry. 2004. Bodily Knowing. Re-thinking our understanding of procedural knowledge. *Philosophical Explorations*, **7**, 37–54.
- Young, Garry. 2009. Case Study Evidence for an Irreducible Form of Knowing How to: An Argument Against a Reductive Epistemology. *Philosophia*, **39**, 341–360.
- Young, Garry. 2011. Irreducible Forms of Knowledge-How in Patients with Visuomotor Pathologies. An Argument against Intellectualism. In: Newen *et al.* (2011). 51–78.
- Zagzebski, Linda. 2001. Recovering Understanding. In: Steup, Matthias (ed), *Knowledge, Truth, and Duty. Essays on Epistemic Justification, Responsibility, and Virtue*. Oxford: Oxford University Press. 235–252.
- Zalta, Edward N. (ed). 1995 ff.. *The Stanford Encyclopedia of Philosophy*. Stanford: CSLI.
- Zardini, Elia. 2013. Knowledge-how, true indexical belief, and action. *Philosophical Studies*, **164**, 341–355.
- Zwicky, Arnold M., & Sadock, Jerrold M. 1975. Ambiguity tests and how to fail them. In: Kimball, John P. (ed), *Syntax and Semantics. Volume 4*. New York: Academic Press. 1–36.

Index of Cases

- Amputee Chef (Snowdon 2003), 155
Amputee Cyclist (Hawley 2003), 155
Amputee Pianist (Stanley & Williamson 2001), 155
Animals & Infants (Ryle 1949), 43
Annoying Smoker (Hawley 2003), 163
Avalanche (Hawley 2003), 26
- Barney (Pritchard 2012), 192
Bela Karoli (Fridland 2012), 54
- Chuck Knoblauch (Dreyfus 2007), 104
Clocks & Seals (Ryle 1949), 32
Comanche Dominance (Kvanvig 2003), 192
Conjurer (Carr 1979), 185
Critics & Writers (Ryle 1949), 54
- D.F. (Goodale & Milner 1992), 199
Digestion (Stanley & Williamson 2001), 34
Distance (Snowdon 2003), 149
Driving Contexts (Hawley 2003), 43
- Ear-Wiggling Frustration (Bengson & Moffett 2007), 151
Exit (Snowdon 2003), 159
- Falcon, 195
Flight Simulator (Stanley & Williamson 2001), 188
- Ginet's Door (Ginet 1975), 281
Gray's Elegy (Carr 1979), 184
- H.M. (Corkin 2002), 201
- Impossible Pudding (Snowdon 2003), 152
Impressive Skating (Hawley 2010), 179
- Klement, 178
Knotting Sailor (Ryle 1949), 55
Kytoon (Bengson & Moffett 2011c), 161
- Light Switch (Hawley 2003), 179
Lottery (Stanley & Williamson 2001), 25
Lucky Cake (Hawley 2003), 26
Lucky Salchow (Bengson & Moffett 2007), 162
- Multiplication (Carr 1981a), 141
- Nolan's Curve (Poston 2009), 188
- Olavi's Finnish Tango (Poston 2009), 189
- Panic (Hawley 2003), 168
Park Chess, 39
Philadelphia (Steel 1974), 180
Pi (Bengson & Moffett 2011c), 154
Pink Ice Cream (Millikan 2000), 69
Press-ups (Snowdon 2003), 158
- Quintuple Salchow (Bengson & Moffett 2007), 150
- Rock (Snowdon 2003), 159
Ryle's Range of Cases, 14

- Sight-Reading (Snowdon 2003), 160
Ski Experts Pat and Albert (Bengson & Moffett 2011c), 150
Ski Instructor (Stanley & Williamson 2001), 150
Speech Impediment (Snowdon 2003), 153
Superman (J. Williams 2008), 178
The Lucky Light Bulb (Cath 2011), 189
Triangles (J. Williams 2008), 186
Unwelcome Sign, 78

Index of Persons

- Abbott, Barbara, 219, 223, 225, 229,
230, 243–246
- Adams, Marcus P., 201, 203, 296
- Adler, Jonathan, 67
- Alexander, Joshua, 217
- Allen, Colin, 205
- Alter, Torin, 248
- Ammerman, Robert, 296
- Anderson, Michael C., 105
- Andow, James, xii
- Andrews, Kristin, 205
- Annas, Julia, 80, 81
- Anscombe, Elizabeth, 4, 102
- Aquinas, *see* Thomas Aquinas
- Aristotle, 28, 71, 85
- Arpaly, Nomy, 291
- Ashwell, Lauren, 170
- Audi, Paul, 248
- Bach, Kent, 89, 243
- Badets, Arnaud, 125
- Bargh, John A., 126
- Bartels, Andreas, 112, 286
- Barz, Wolfgang, 178
- Battaly, Heather, 121
- Beck, Jacob, 205
- Beilock, Sian L., 105
- Bekoff, Marc, 205
- Bengson, John, 1–3, 6, 16, 17, 22, 31,
60, 90, 103, 112, 141, 150, 151,
154, 161–166, 192, 204, 210–218,
223–228, 232, 233, 236, 241–243,
247, 249, 252, 253, 257, 268–270,
274, 285
- Berry, Dianne C., 119
- Betz, Gregor, xii
- Bhatt, Rajesh, 239
- Bird, Alexander, 167
- Blackburn, Simon, 37
- Blanc, Raymond, 155
- Blandin, Yannick, 125
- Blass, Steve, 105
- Blome-Tillmann, Michael, 191, 214
- Boden, Margaret A., 204
- Boghossian, Paul, 37, 113
- Bohse, Helen, xii
- Bongers, Raoul M., 104
- Börchers, Fabian, xii
- Boutin, Arnaud, 125
- Brandom, Robert B., xii, 295
- Braun, David, 236
- Brett, Nathan, 32
- Broadbent, Donald E., 119
- Brogaard, Berit, 67, 191, 193, 219,
222, 223, 237, 243
- Brown, David G., 32, 148, 236
- Brown, Jessica A., xii, 1, 246
- Brownstein, Michael, 100, 102, 119,
187
- Buckwalter, Wesley, 217
- Burge, Tyler, 181, 182
- Burghardt, Gordon M., 205
- Buskell, Andrew, 125
- Bzdak, David, 201
- Capone, Alessandro, 191, 214
- Carr, David, 25, 34, 35, 141, 184,
185, 212, 219, 295

- Carr, Thomas H., 105
 Carter, J. Adam, 32, 190, 191, 242
 Cath, Yuri, 187–191, 193, 248, 281
 Chalmers, David, 34, 242
 Chartrand, Tanya L., 126
 Chase, Christopher, 197
 Chemero, Anthony, 29
 Choi, Sungho, 172, 173
 Christensen, Wayne, 104, 125
 Clarac, François, 99
 Clark, Andy, 242
 Clarke, Henry, 284
 Clarke, Randolph, 28, 172
 von Clausewitz, Carl, 283
 Cohen, Neal J., 197, 201, 203
 Collins, John, 241
 Constantin, Jan, xii
 Corkin, Suzanne, 201
 Cova, Florian, 89
 Craig, Edward, 295
 Cross, Charles, 237, 238
 Cummings, Sam, 186
 Cummins, Robert, 286
 Czarnecki, Boleslaw, xii, 242

 D.F., *see* F., D.
 Damschen, Gregor, 219, 220, 277
 Dänzer, Lars, xii
 Davidson, Donald, 165, 179, 183, 190, 196
 Davis, Wayne, 214
 Demmerling, Christoph, 128
 Descartes, René, 74, 78, 124
 Devitt, Michael, 197, 198, 205, 219, 246
 Di Nucci, Ezio, 89
 Dickie, Imogen, 3, 250–252, 264
 Dienes, Zoltan, 119, 196, 198
 van Dijk, Ludger, 104
 Dinges, Alexander, xii
 Ditter, Andreas, 221, 223
 Douskos, Christos, 244
 Drayson, Zoe, 286
 Dreyfus, Hubert L., 2, 103–106, 117, 125, 285
 Dreyfus, Stuart E., 2, 103
 Dupoux, Emmanuel, 89

 Einstein, Albert, 294
 El Kassar, Nadja, xii
 Elgin, Catherine Z., 110
 Elzinga, Benjamin, 6, 31, 50, 124
 Evans, Vyvyan, 231, 232

 F., D., 199, 200
 Falcon, Andrea, 85
 Falkum, Ingrid Lossius, 230, 232
 Fantl, Jeremy, 1, 3, 4, 30, 207, 277, 296
 Fara, Michael, 172–174
 Farkas, Katalin, 242, 247, 254, 261
 Feinstein, Mark, 197
 Fine, Kit, 248
 Fisher, Justin C., 28
 Fitch, Greg, 67
 Flegal, Kristin E., 105
 Fodor, Jerry, 285
 Frankfurt, Harry, 89, 92
 Frankish, Keith, 204
 Fridland, Ellen, xii, 54, 104, 124–126, 196, 214, 278, 284, 286, 287
 Frijda, Nico, 95

 Gabrieli, John D. E., 197
 Gadamer, Hans-Georg, 104
 Garfield, Jay, 197
 Gazdar, Gerald, 225
 Geeves, Andrew, 104, 125
 Gerken, Mikkel, 1
 Gettier, Edmund, 188
 Gibson, James, 29
 Ginet, Carl, 192, 281, 290
 Ginzburg, Jonathan, 225, 237
 Glauer, Ramiro, 112, 286
 Glick, Ephraim, 4, 24, 229, 247, 248, 250, 272, 273, 276
 Glock, Hans-Johann, xii
 Glüer, Kathrin, 36
 Goldman, Alvin, 192
 Goodale, Melvyn A., 199

- Gottlieb, Gabriel, 104
 Grajner, Martin, xii
 Greco, John, 121, 122, 174, 176
 Grice, Herbert Paul, 41, 219
 Grimm, Stephen R., 108, 193
 Groenendijk, Jeroen, 237, 238
 Grzankowski, Alex, 243

 H.M., *see* M., H.
 Haase, Matthias, xii
 Habgood-Coote, Joshua, xii, 235, 241, 243
 Haddock, Adrian, 102
 Hamblin, Charles L., 237
 Han, ZaoHong, 106
 Hartland-Swann, John, 296
 Hawley, Katherine, 26, 31, 43, 61, 81, 82, 89, 155–157, 163, 168, 169, 173, 179, 187, 295
 Heinrich, Bernd, 205
 Henning, Boris, 85
 Hermann, Louis S., 205
 Hetherington, Stephen, 188, 256, 277, 296
 Heuer, Herbert, 125
 Hintikka, Jaako, 219, 236
 Hoeltje, Miguel, xii
 Hornsby, Jennifer, 6, 15, 30, 102, 282, 286
 Howell, Robert J., 248
 Hyman, John, 68

 Ichikawa, Jonathan Jenkins, 191

 Jackman, Henry, xii
 Jackson, Frank, 248
 Jacob, Pierre, 89
 Janzen, Jan, xii
 Jaster, Romy, xii, 175
 Johnson, Kent, 205, 241
 Johnston, Mark, 166, 167
 Jung, Eva-Maria, xii, 1, 2, 128, 200, 205, 274
 Kallestrup, Jesper, 238
 Karttunen, Lauri, 237
 Katzoff, Charlotte, 34, 219
 Kauppinen, Antti, 217, 219
 Keil, Geert, 286
 Kelly, Sean Dorrance, 199, 200
 Kenny, Anthony, 25, 113, 114, 286
 Kern, Andrea, xii
 Kietzmann, Christian, xii
 Kindermann, Dirk, xii
 King, Jeffrey C., 67, 150
 Knobe, Joshua, 89, 217
 Knoblauch, Chuck, 104, 105
 Knowlton, Barbara J., 203
 Koethe, John, 274, 277
 Koffka, Kurt, 29
 Koppelberg, Dirk, xii
 Kotzee, Ben, 32
 Kraft, Tim, xii
 Krakauer, John, 2, 202, 259, 261, 262, 285–287, 291
 Kratzer, Angelika, 271
 Kremer, Michael, 6, 31, 68, 74, 121, 124, 219, 230, 289, 295, 296
 Krifka, Manfred, 237
 Kripke, Saul, 36, 141, 220
 Kumar, Victor, 295
 Kusch, Martin, xii
 Kvanvig, Jonathan L., 108–110, 192, 193

 Langeloh, Jacob, xii
 Lauer, David, xii
 Laurence, Stephen, 113
 Lavin, Douglas, 35
 Lehrer, Keith, 206
 Levin, Janet, 248, 288
 Lewandowska-Tomaszczyk, Barbara, 224, 230
 Lewin, Kurt, 29
 Lewis, David, 120, 168, 171, 172
 Lihoreau, Franck, 204, 219
 Loeb, Gerald E., 99
 Löwenstein, David, 129, 236, 248, 273, 278, 279, 281, 295, 296
 Luce, Gay Gaer, 100
 Ludwig, David, xii

- Luntley, Michael, 115
- M., H., 201, 202
- MacMahon, Clare, 105
- Madison, Brent James Charles, 296
- Maier, John, 25, 28, 29, 172, 175
- Malle, Bertram F., 89
- Manley, David, 173
- Margolis, Eric, 113
- Martin, Charles Burton, 167, 172
- Massen, Cristina, 125
- Masto, Meghan, 242, 247, 254, 261
- Matsuzawa, Tetsuro, 205
- Matthews, Gareth, 206
- May, Mark, 112, 286
- McDowell, John, 104, 117
- McGrath, Matthew, 67
- McIlwain, Doris, 104, 125
- McKay, Thomas, 67, 178
- McLeod, Peter, 119
- Mellor, D. Hugh, 248
- Menary, Richard, 99
- Michaelis, Laura A., 237
- Michaels, Claire F., 29
- Michaelson, Eliot, xii, 100, 119, 187
- Miller, Alexander, 36
- Millikan, Ruth Garrett, 69, 152, 153, 196
- Milner, A. David, 199
- Milner, Brenda, 201
- Moffett, Marc A., 1–3, 6, 16, 17, 22, 31, 60, 90, 103, 112, 141, 150, 151, 154, 161–166, 192, 204, 210–218, 223–228, 232, 233, 236, 241–243, 247, 249, 252, 253, 257, 268–270, 274, 285
- Montero, Barbara, 106, 125
- Müller, Andreas, xii
- Mullin, Amy, 206
- Nadelhoffer, Thomas, 89
- Nagel, Jennifer, 217
- Nagel, Thomas, 248
- Nelson, Michael, 67, 178
- Nemirow, Laurence, 248
- Newen, Albert, 128, 200, 205, 274
- Nida-Rümelin, Martine, 248
- Nimtz, Christian, xii
- Ninan, Dilip, 120
- Noë, Alva, 34, 155–157, 205, 216–218, 273, 277, 281
- Nolan, Daniel, 279
- Orozco, Joshue, 188
- Parent, Ted, 237, 238
- Park, Shelly M., 75
- Parkinson, Brian, 95
- Parry, Richard D., 124, 277
- Pavese, Carlotta, 1, 74, 257, 261, 273, 275, 276, 287
- Perner, Josef, 196, 198
- Perry, John, 120
- Pissarskoi, Eugen, xii
- Poston, Ted, 187–190
- von Prince, Kilu, xii
- Pritchard, Duncan, 32, 188, 190–195
- Prochazka, Arthur, 99
- Quast, Christian, xii
- Rabagliati, Hugh, 230
- Railton, Peter, 291, 292
- Ramsey, William M., 204
- Reber, Paul J., 203
- Reed, Nick, 119
- Reisberg, Daniel, 197
- Ren, Huiming, 219
- Rey, Georges, 104
- Rietveld, Erik, xii, 29, 38, 104, 127
- Riggs, Wayne D., 110
- Rissland, Edwina, 197
- Roberts, Craige, 225, 237
- Rödl, Sebastian, xii, 28
- Roeber, Blake, 191
- Roelofson, Floris, 237, 238
- Rosefeldt, Tobias, xi, xii, 219, 221, 274
- Roth, Martin, 286
- Rothwell, John C., 99
- Rumfitt, Ian, 219, 223, 244

- Ryle, Gilbert, xi, 5–7, 14–18, 21–27, 30–33, 35–39, 43, 45, 47–75, 77–79, 81, 83–85, 87, 96, 97, 101–103, 114, 115, 119, 123, 124, 129, 130, 140, 171, 173, 204, 205, 212, 221, 230, 236, 256–258, 263, 267, 270, 277, 278, 280, 282–284, 288–291, 293, 295
- Rysview, Patrick, 191
- Sachsenweger, Beate, xii
- Sadock, Jerrold M., 223, 224, 226, 232
- Sag, Ivan A., 225
- Salchow, Ulrich, 150
- Sandis, Constantine, 79, 86
- Sanford, David H., 279
- Santorio, Paulo, 245
- Sax, Greg, xii, 24, 219
- Schaffer, Jonathan, 238, 248
- Schear, Joseph K., 117
- Scheffler, Israel, 289
- Schiffer, Stephen, 273, 277
- Schmechtig, Pedro, xii
- Schneider, Walther, 78, 91
- Schröder, Dirk, 273
- Schroeder, Mark, 238
- Schwarzschild, Roger, 257
- Schwenkler, John, 102
- Schwitzgebel, Eric, 96, 103, 288
- Scoville, William Beecher, 201
- Searle, John Rogers, 40
- Segal, Julius, 100
- Seiberth, Luz Christopher, xii
- Seide, Ansgar, xii
- Sellars, Wilfrid, 5, 93–96
- Sennet, Adam, 224, 230
- Setiya, Kieran, 35, 79, 102
- Sgaravatti, Daniele, 238
- Shirreff, Patrick, xii
- Shope, Robert K., 172
- Shpall, Samuel, 79, 86
- Slaby, Jan, xii
- Small, Will, 63, 187
- Smith, Arthur David, 170
- Smythe, William E., 295
- Snowdon, Paul F., 16, 34, 81, 82, 89, 149, 152, 153, 155, 158–160, 248
- Sosa, Ernest, xii, 30, 174, 176, 207, 217, 256
- Spiegel, Thomas Jussuf, xii
- Squire, Larry R., 197, 201, 203
- Srinivasan, Mahesh, 230
- Stalnaker, Robert, 273, 288
- Stanley, Jason, xi, xii, 2, 4, 6, 25, 34, 75, 82, 83, 89, 92, 133, 150, 155, 165, 166, 188, 190, 191, 196, 198, 200, 202, 204, 210, 211, 216, 222, 223, 225, 235–246, 248, 250–253, 257–265, 270–274, 280, 281, 285–289, 291, 292, 294, 295
- Starkes, Janet L., 105
- Steel, Thomas J., 180–182
- Stekeler-Weithofer, Pirmin, xii
- Steup, Matthias, 191
- Stillings, Neil, 197
- Stokhof, Martin, 237, 238
- Stout, Rowland, 75, 85, 237, 238
- Streitfeld, Jason, xii
- Study, Eduard, 294
- Sutton, John, 104, 125
- Sytsma, Justin, 217
- Tanney, Julia, 39, 75, 286
- Tetens, Holm, xi, xii
- Thomas Aquinas, 28
- Tolksdorf, Stefan, xii
- Toner, John, 125
- Toribio, Josefa, 200, 285
- Traugott, Elizabeth Closs, 232
- Trestman, Michael, 205
- Tsai, Cheng-hung, 277, 295
- Tulving, Endel, 198
- Turri, John, 121
- Tye, Michael, 248
- Umiltà, Carlo, 99
- Urban, Anastasia, xii
- Van Gulick, Robert, 99

- Van Riel, Raphael, xii
Vetter, Barbara, xii, 172–176, 230,
232, 240
Vicente, Augustin, 230
Viebahn, Emanuel, xii, 230, 232, 240
Vivhelin, Kadri, 174

Wallis, Charles, 100, 201, 246
Wasow, Thomas, 225
Wasserman, Ryan, 173
Watzl, Sebastian, 99
Weatherson, Brian, 277, 289
Wehofsits, Anna, xii
Weicholt, Martin, xii
Weinberg, Jonathan M., 217
Weisler, Steven, 197, 225
Westra, Adam, xii
Whittle, Ann, 172
Wieland, Jan Willem, 279
Wiggins, David, 74, 85, 219, 282, 290,
291
Wikforss, Åsa, 36
Williams, John N., 163, 178, 179,
181, 186, 277, 295

Williams, Michael, 123
Williamson, Timothy, xi, 2, 6, 25,
34, 82, 83, 89, 92, 133, 150, 155,
188, 211, 216, 223, 225, 235–237,
239, 240, 248, 252, 257, 260–265,
270–273, 280, 281, 291, 294, 295
Wilson, George, 79, 86
Wittgenstein, Ludwig, 2, 11, 36, 37,
39, 127, 141, 145, 234, 286, 294
Wolpaw, Jonathan R., 99
Worthmann, Hannes, xii
Wright, Crispin, 36
von Wright, Georg Henrik, 79
Wright, Jennifer C., 216–218, 224,
226
Wurmbrand, Susanne, 221

Young, Garry, 128

Zagzebski, Linda, 110
Zakkou, Julia, xii
Zardini, Elia, 191, 238
Zwicky, Arnold M., 223, 224, 226,
232

Index of Subjects

- ability, 1, 22–23, 29, 221
situation-specificity of, *see*
situation-specificity
wide vs. narrow, *see* specific vs.
general
as disposition?, 171–176
as reliable, 25–27, 163
specific vs. general, 28–29, 148–
149, 152–153, 157, 159–162, 214–
215
vs. being able, *see* as reliable
vs. mere possibility, *see* as reliable
- achievement, 21–35, 295–296
- activities
normative account of, 18–19, 41–
43, 57–58, 97–98, 131–132, 185
normative demands of, 19–20
status of norms of, 20–21
vs. actions, 14–15
- affectivity, 128
- affordance, 29, 126–128, *see also* abil-
ity
- agnosia, 199–201
- ambiguity, 219, 223–228
tests, 223–228, 232–233
- amnesia, 197, 201–203
- animal knowledge, 207
- animals, 43–44, 204–206
- answers, *see* knowledge-*wh*
- anti-intellectualism
and Ryle, 29–39, 71–75
common ground with Rylean re-
sponsibilism, 21–29
demarcation of, 2–3
rapprochement with intellectual-
ism, 3, 294
- Argument from Linguistics, 245–248
- Argument from Translation, 220–223,
225–226, 243
- artificial intelligence, 196, 204
- assessment capacity, 107–115, 254–
255, 258–292
and regress, 133–138, 264–265
- assessments, 55–57, 115–122, 254–255
automatic, 126
from testimony, 132–133, 139
of one's acts, 119–120, 122–129,
138–140
of one's options, 119–120, 122–129,
138–140
- attention, 98–100
- automaticity, 15, 78–81, 202–203
and learning, 90–91
and normative guidance, 92–98,
103–104, 281–282, 285–287
degrees of, 80–81
- basic action, 34–35, 44, 151, 160,
175–176, 201
- behaviorism, 74–75
- belief, 96, 103, *see also* propositional
knowledge
- cancellability, *see* know-how ascrip-
tions, and ability ascriptions
- capacity, *see* ability
- capacity epistemic luck, *see* epistemic
luck

- competence, 7, 30, *see also* know-how
 concealed questions, 243, 250
 conception, *see* understanding
 concepts, 112, 114
 conceptual capacity, 112–115, 136–
 137, 141–142, 211–213
 conditionals, 172–173
 conjunction reduction, *see* ambiguity,
 tests
 consciousness, 98–100
 constitution, 247–251, 263–264
 constitutive vs. regulative rules, 40–
 41
 context-dependence, 39–45, 157, 160–
 162, 229
 contradiction test, *see* ambiguity,
 tests
 control, 88–89, 124–125, *see also* re-
 sponsible control
 conversational implicature, *see* impli-
 cature
- de dicto*, 181–187
de re, 181–187
de se, 120, 239
 declarative knowledge, 4, 196–199,
 203
 as propositional knowledge?, 197–
 198
 default trust, 291–292
 demonstrative conception, 70, 114,
 116–118, 183
 demonstrative knowledge, *see* demon-
 strative conception
 directed discontent, 127–128
 dispositions, 87, 171–174
 and abilities, 171–176
 disquotations, 245–246, 249
- ellipsis test, *see* ambiguity, tests
 embedded questions, *see* knowledge-
wh
 environmental luck, *see* epistemic
 luck
- episodic memory, *see* declarative
 knowledge
 epistemic luck, 108, 187–196
 capacity epistemic luck, 194–196
 environmental luck, 190, 192–193
 experimental philosophy, 216–219
 explanation of action, 84–88
 as intelligent, 85–86
 as intentional, 85–87
 explicit knowledge, *see* implicit vs.
 explicit
 extended mind, 242
- fact, 1
 factivity, 109–110
 finkishness, 167–170, 187
 intrinsic, 170
 formal cause, 85
 fossilisation, 106
 free relatives, 242, 249
 functionalism, 288–289
- Gettier case, *see* epistemic luck
 gradability, 27, 108, 256–259
 ground, 247–251, 263–264
- hommunculus, *see* sub-personal ex-
 planation
- image-like representation, 128
 imperatives, 65
 implicature, 213–219, 231–232
 implicit vs. explicit, 38–39, 136, 197–
 198, 255, 283–284
 indexicality, 119–120, 273
 infants, 43–44, 205
 intellect, 58–62
 intellectualism
 and regress, *see* Ryle's Regress
 and the practical import of know-
 how, *see* intelligent practice,
 explanation of
 common ground with Rylean re-
 sponsibilism, 250–256, 289–292
 demarcation of, 2–3, 191

- objectualist, 2, 192, 249–250, 253, 255, 268–270, 273–274
- propositionalist, 2, 190–191, 236, 245–249, 253–276
- rapprochement with anti-intellectualism, 3, 289–292, 294
- revisionary, 191
- ‘intelligence concepts’, 15–18
- ‘Intelligence’ vs. ‘intelligence’, 16–17
- intelligent practice, 5, 15–35
 - explanation of, 5, 29–35, 122–129, 265, 268–276
- intentional action, 15, 78–79, 88–89
- intentional error, 282–283, 286–287
- know-how
 - and knowledge-*wh*, *see* knowledge-*wh*
 - as competence or skill, 23–24, 210–211, 261–265, 294–296
 - as disposition?, 171–176
 - gradability of, *see* gradability
 - learning of, 60–65
 - manifestation of, 47–52, 262–263
 - practical import of, *see* intelligent practice, explanation of
 - teaching of, 52
 - vs. knowledge *about* how, *see* vs. mere understanding
 - vs. mere ability, 32–35, 147, 158–159, 202–203, 221
 - vs. mere *de re* know-how, 183–187
 - vs. mere understanding, 47–52, 140–142, 147, 149–150, 154–155, 157, 211–215, 218–219, 243
- know-how ascriptions, 3–4, 210–219
 - and ability ascriptions, 211–219
 - and embedded questions, *see* knowledge-*wh*
 - and knowledge-*wh*, *see* knowledge-*wh*
 - as ambiguous, *see* ambiguity
 - as gradable, *see* gradability
 - as polysemous, *see* polysemy
 - bifurcational view of, 219–234
 - in Cantonese, 223
 - in French, 220, 244
 - in German, 37, 220–222
 - in Modern Greek, 244
 - in Montana Salish, 223
 - in Russian, 223
 - in Turkish, 223
 - ‘knowing how’, 3–4
 - ‘knowing that’, 3–4
 - knowledge *about* how, *see* know-how, vs. mere understanding
 - knowledge argument against physicalism, 248
 - knowledge what it is like, 248
 - knowledge-how, *see* know-how
 - knowledge-that, *see* propositional knowledge
 - knowledge-*wh*, 190–191, 236–240
 - as capacity to choose answers, 241–242, 247, 254, 261
 - as objectual knowledge, 242–243, 247
 - as propositional knowledge, 237–241
 - contextual relevance of answers, 238
 - knowing answers *as* answers, 238
 - mention-all-reading, 238
 - mention-some-reading, 238
 - propositions vs. sets vs. questions as *relata*, 238, 241
 - strong exhaustivity, 238
 - uniformity of, 241–243, 260–261
 - weak exhaustivity, 238
 - with infinitival *wh*-complement, 238–240
- machines, 204
- manifest image, 5
- masking, 153–154, 166–187
 - intrinsic, 170–171
- McDowell-Dreyfus-debate, 103–105, 117
- memory, *see* declarative knowledge, *see* procedural knowledge

- mimicking, 162–166, 170
- minimal vs. evaluative norms, 40–41
- mirror-reading studies, 201–203
- modality, 172–173, 238–240, 271–273
 abilitative, 21–29, 238–240, 271–273
 deontic, 238–240
- motor acuity, 285
- non-observational knowledge of action, *see* practical knowledge
- normal situations, *see* situation-specificity
- normative guidance, 30–31, 122–129
 and regress, 129–133, 289–292
- normative practice, 17–21, 92–98
- normativity of meaning, 36–39
- norms, *see* normative practice
- novelty, 259–261
- objectual knowledge, 222, 225
- objectualist intellectualism, *see* intellectualism, objectualist
- ontological priority, 247–251, 263–264
- opacity, 177–187
- option, *see* ability, specific vs. general
- ought-to-be, 93–98, 123
- ought-to-do, 93–98
- perception, 123
- perceptual acuity, 285
- personal-level explanation, *see* sub-personal explanation
- phenomenology of expertise, 102–106
- polysemy, 229–234
 and ambiguity tests, 232–233
 and implicature, 231
 Ryle on, 230
- practical knowledge, 4, 101–102
- practical luck, 87, 133, 162–166, 170–171, 190
- practical mode of presentation, 179, 273–276
 computer analogy for, 275–276
- practicalism, 296
- pragmatic encroachment, 190–191
- pragmatically salient alternatives, *see* salient alternatives
- ‘PRO’, *see* unpronounced constituents
- procedural knowledge, 4, 196–199, 202–203
 as know-how?, 197–198
- proper names, 185–186
- propositional knowledge, 1, 65–68, 115–122
 and stakes-sensitivity, *see* pragmatic encroachment
 as disposition?, 68, 96, 103, 288–289, 295–296
- demonstrative, *see* demonstrative conception
- implicit, *see* implicit vs. explicit
- interdependence with know-how, 73–74, 290–291, 296
- Ryle on, 65–68, 71–74, 121, 288–289
- testimonial, *see* testimony
- verbalizability of, *see* demonstrative conception
- vs. (justified) true belief, 70, 121–122, 187–188
- propositionalist intellectualism, *see* intellectualism, propositionalist
- propositions, 1, 66–68, 247
- questions, *see* knowledge-*wh*
- reflective knowledge, 207
- reinforceability, *see* know-how ascriptions, and ability ascriptions
- reliability, *see* ability, as reliable
- responsibilism, 123
- responsible control, 122–129, 292
 and regress, 129–133, 289–291
- revisionary intellectualism, *see* intellectualism, revisionary
- rote, *see* automaticity
- routine, *see* automaticity
- rule-following, 36–39

- rules, *see* normative practice, *see* rule-following
- Ryle's Regress, 129, 276–289
 and applying, 278–280, 283–284
 and automaticity, 281–282, 285–287
 and contemplating, 280
 and functionalism, *see* functionalism
 and kinds of intellectualism, 278
 and selecting, 278–280, 282–284
 and sub-personal explanation, *see* sub-personal explanation
- salient alternatives, 213–219
- self-assessments, *see* assessments
- self-monitoring, *see* assessments
 Ryle on, 124
- self-reflexivity, 135–138
 and circularity, 140
- semantic memory, *see* declarative knowledge
- situated normativity, 126–128
- situation, *see* ability, specific vs. general
- situation-specificity, 25–29, 126–128,
see also assessments, of one's options
- skill, 7, 124–125, 250–251, 261–265,
see also know-how
- sleepwalking, 100–101
- stakes-sensitivity, *see* pragmatic encroachment
- standard, *see* normative practice
- Standard Linguistic Account, 236–245
- Steve Blass Disease, 105
- sub-personal explanation, 285–287
- success, *see* achievement
- Swampman, 196
- testimony, 61, 67–70, 132–133, 187
- theoretical knowledge, 4
- understanding, 1, 31, 68–70, 107–115, 192–194
 propositional vs. objectual, 108, 193
- unpronounced constituents, 239, 244
- vagueness, 39–45, 157, 160–162
- verb phrase deletion, *see* ambiguity, tests
- verbalizability, *see* demonstrative conception
- virtue reliabilism, 121–123, 174–176
- ways of doing something, 240, 252, 284–287
- wh*-complement, *see* knowledge-*wh*
- zeugma, *see* ambiguity, tests

