University of Alberta

Knowing-How

by

Victor Kumar 

A thesis submitted to the Faculty of Graduate Studies and Research
in partial fulfillment of the requirements for the degree of

Master of Arts

Philosophy

Edmonton, Alberta
Fall 2007

Your file Votre référence
ISBN: 978-0-494-33136-1
Our file Notre référence
ISBN: 978-0-494-33136-1

Abstract:

The essay is an epistemological investigation of "knowing-how", as distinct from "knowing-that". My approach, though ecumenical, is substantively naturalistic: I think of knowledge as a natural kind, a cluster of epistemic-cum-psychological properties. Although the topic is ostensibly epistemological, I draw on many other philosophical disciplines as well, especially the philosophy of mind and normative theory. Three questions guide my inquiry. What kind of mental state is knowing-how? What is the abstract theoretical relationship between knowing-how and knowing-that? And what sorts of activities do we know how to engage in? I argue that knowing-how and knowing-that are distinct, that some of the activities we know how to engage in are non-voluntary and unconsciously performed, and that the theory of knowing-how subsumes the theory of knowing-that. I also highlight some of the interesting connections between the theory of knowing-that and other philosophical issues. The upshot of my discussion is a positive account of knowing-how as a mental state that has both directions of fit—it is both representative and action-guiding—and that does not always encode information propositionally.

For my parents.

Without their unconditional love and support my philosophical studies would not have been possible, nor worth pursuing.

Acknowledgements

I have developed the ideas in this essay through discussion with many people. I am indebted especially to my advisor, Bruce Hunter, for laying out nearly all the issues and possible approaches with regard to this subject in his graduate seminar, taken in the winter of 2006 at the University of Alberta. Small parts of every chapter of the thesis were first articulated in the paper I wrote for Bruce's seminar. To my regret, I am unable to cite Bruce at every point at which he has contributed to my arguments: the entire essay is undoubtedly better because of his critical feedback. Over the past two years Bruce has also influenced my understanding of epistemology more generally.

Thanks also to Nathan Ballantyne, Kenneth Boyd, Richmond Campbell, Wesley Cooper, Julian Fowke, Richard Greene, Octavian Ion, Adam Morton, Jennifer Nagel, and Rob Wilson who read all or part of the essay and/or participated in extended discussions. One of these individuals, Richmond Campbell, is owed special thanks for introducing me to the naturalistic approach I employ in the essay and for providing a model of philosophical rigour and care. Adam Morton provided important counsel during the early stages of the essay.

It is said that one learns as much from one's peers as one's teachers, and that is no less true in my case. A special thanks then to Julian Fowke, Octavian Ion and Kenneth Boyd among other fellow students at the University of Alberta.

Most of chapter two was presented to an audience at the Canadian Philosophical Association Congress at Saskatoon, SK in May 2007. An earlier draft of chapter two was presented at the Northwest Philosophy Conference in Portland, OR in November 2006. Both drafts were presented to fellow students during the Graduate Student Colloquium Series at the University of Alberta in October 2006 and May 2007. Thanks to all three sets of audiences for helpful questions and comments.

Work on this project was supported by, first, a CGS MA scholarship from the Social Sciences and Humanities Research Council in 2006/2007 and, second, a University of Alberta MA entrance scholarship in 2005/2006, both of which allowed me to forgo assistantship duties.

Finally I wish to thank my family for their stimulation and encouragement. No one could ask for more supportive parents, and it is to them that this work is dedicated. In my sister's sharp mind, generosity and indulgence I found my first teacher. To her I owe more than I can return.

Table of Contents:

Chapter 1: Introduction

The topic of this essay is "knowing-how", as distinct from "knowing-that". The distinction is well known, no less to those unfamiliar with the philosophical literature on the topic. It is the difference between an athlete's knowledge and an instructor's, an artist's knowledge and a critic's. The one involves a certain fittingness for practical engagement with the world, the other a theoretical appreciation of that same world.

Arguably, our know-how is more important than our knowledge of facts: intelligence is manifested primarily in action rather than in judgment. In any case, since it grounds our ability to engage in a wide range of intelligent thought and action, knowing-how denotes an important epistemological category—and is worthy of study in its own right since knowing-how and knowing-that are distinct epistemic kinds.

My aim in this essay is to provide a positive account of knowing-how. My strategy will be to draw on material in mainstream epistemology (the theory of knowing-that) in order to formulate and address questions about knowing-how, questions that are, by and large, no less germane to philosophy of mind than epistemology. Issues that are central to the theory of knowing-how include the kind of mental state knowing-how is, the range of activities we know how to engage in, and the relation between knowing-how and knowing-that. This essay provides a coherent set of positions on these issues and others.

In this, the introductory chapter I have three goals. First, for the uninitiated, I

motivate a philosophical interest in knowing-how (section 1.1). Second, I describe some of the background literature relevant to the essay and locate my approach within the literature (section 1.2). And last, I provide a detailed outline of what is to come in the following chapters (section 1.3).

1.1 Why Knowing-How?

Epistemology owes an account of what knowledge is. Motivation for the project finds one of its sources in the importance we place on knowledge, especially in the sciences, and the practical value it enjoys by contributing to the fulfillment of our ends. But as those like Gilbert Ryle are right to point out, much of our knowledge seems to be realized not as propositional knowledge, but rather in the *skills* or *know-how* we exercise as we confront our surroundings (Ryle 1949: ch. 2). Knowing-how and knowing-that are two species of the same genus. Both constitute ways in which intelligent creatures manifest a "cognitive fit" with their environments. Both are important to us and have practical value.

As Ryle notes, there is a bias among philosophers to privilege propositional knowledge as the primary, and perhaps only, manifestation of intelligence (1949: 26). But this presumption is not borne out by our evaluative practices. Physical and social activities, when performed well, merit epistemic approval. An athlete *cleverly* executes a play or maneuver; a manager *deftly* negotiates conflicts between her employees. Chapter 2 of Ryle's book *The Concept of Mind* is filled with similar examples, designed to illustrate that epistemic appraisal has as one of its primary targets a wide range of human activities, and is thus not limited to the formation and

revision of beliefs. To know how to perform these activities is to garner a positive epistemic status.[1]

Knowing-how matters. Indeed, knowing-how often matters as much or more than knowing-that. In academia, for example, we are typically interested in the intellectual know-how of our peers rather than merely their ability to get at the truth or the number and quality of justified true beliefs they possess (see Ryle 1949: 28). Creativity, speed, clarity, economy in expression, etc., are all species of intellectual know-how that are not truth-linked and that contribute to the value of an inquirer. In other contexts we care almost not at all about propositional knowledge. It does not matter to us a bit if a craftsman is in possession of a correct theory about his craft or if a salesperson boasts an extensive array of true beliefs about social influence; we esteem the former if he knows how to wield his tools and the latter if she knows how to make a sale.

Because knowing-how is so central to human knowledge and because philosophers have for the most part neglected to theorize about it, we are in need of a proper understanding of the subject. The need is even more pressing, however, since the concept of knowing-how is increasingly being invoked outside of epistemology (Hawley 2003). Some philosophers of language believe that semantic competence consists in knowing how to use linguistic expressions (Dummett 1991: ch.4; Putnam 1996); philosophers of mind have argued that our acquaintance with qualia amounts to some sort of know-how (Lewis 1988; Nemirow 1990; Mellor 1993); and moral

---

1 Are physical and social activities really subject to *epistemic* evaluation? These activities, one might argue, do not aim at the traditional epistemic goal of truth. The fact that we use the term 'knows how' suggests that the evaluation is indeed epistemic. In chapter three I offer a *broader* understanding of epistemic evaluation, the objects of which need not be truth-linked, and that is perhaps better captured by the phrase 'intellectual evaluation'.

philosophers have begun to shift their attention from knowledge of moral propositions to moral know-how (Blackburn 1996; Clark 2000). These directions of inquiry will be better informed to the extent that we cultivate a philosophical grasp of know-how.

Knowing-how, I've argued, merits attention from epistemologists for the same reasons that knowing-that does. As well, the concept appears useful in other subject areas and is deployed by philosophers there. Furthermore, it seems as if knowing-how and its cognates are discussed quite extensively in other fields *but under the guise of other terms*. Virtue epistemologists investigate epistemic character traits that appear to be composed of various kinds of intellectual know-how. *'Competence'* seems to be another name for knowing-how: philosophers of mind and action are interested in the conditions that make intelligent action possible and on a natural way of understanding competence it is the same as knowing-how. Scientists possess much *tacit knowledge* that underwrites their success at managing the tools of their trade; some of this tacit knowledge certainly seems to be a kind of know-how (see Polanyi 1962). So, perhaps one reason that knowing-how has received such little attention from epistemologists is that it also falls under the jurisdiction of other philosophers, e.g., philosophers of mind, action and science. One might even worry that if knowing-how is the same as competence, then it may be a proper topic of study only for the various disciplines associated specifically with the mind. I find this possibility unlikely, however, given the rich practice of epistemic appraisal associated with know-how.

If it's true that knowing-how is common to many philosophical subjects, the

results of an epistemological investigation of knowing-how may be significant not only for epistemology but also for other philosophical disciplines—and not merely in contexts where know-how is explicitly invoked. For the same reason, it is imperative that we discuss ideas in subjects other than epistemology if we are to understand the phenomenon.

So far I have tried to vindicate the claim that knowledge-how is a proper subject for epistemological inquiry—although the truth of that claim may be more accurately judged at the end of the essay. Next we'll turn to explicit treatments of know-how by various epistemologists over the last sixty years. This will give us a sense of what the major issues are, along with a preview of my take on these issues.

1.2 Intellectualism and Reflectivism

A preoccupation in epistemology with propositional knowledge has left us unfamiliar with practical or procedural knowledge.[2] Theories of know-how thus remain underdeveloped. To be sure, however, the topic has not been entirely neglected in the philosophical literature. Famous discussions in the 1940s by Ryle (1946, 1949) helped to stimulate interest in the subject and more recent work by contemporary philosophers—most notably Jason Stanley and Timothy Williamson (2001)—have revived discussion of knowledge-how among epistemologists and philosophers of mind and action.

---

2 Now if one sort of know-how is the knowledge of how to acquire true beliefs, then epistemologists are not guilty of ignoring this kind of know-how: discussions of what makes a belief justified or count as knowledge are *ipso facto* discussions of a certain type of know-how. This is a controversial thesis, one that I defend in chapter three; but, in any case, it does not imply that epistemologists have paid sufficient attention to *mental know-how in general* (see my brief discussion above regarding other species of intellectual know-how).

As a first stab at carving up the knowing-how literature, let us distinguish between two opposing camps: the intellectualists and the anti-intellectualists. Intellectualism is the view that all knowing-how consists in the possession of doxastic attitudes.[3] To know how to braid hair, for example, is to have (perhaps implicitly) some theory about how hair can be braided. Anti-intellectualism is simply the denial of this claim. Knowing-how, anti-intellectualists think, is not (or not simply) a matter of having the right sort of beliefs.[4]

Ryle is a classic anti-intellectualist.[5] It will be useful to discuss his view first, since nearly all philosophical treatments of knowledge-how since 1950 take their cue from Ryle. Ryle's target, against which he offers a range of objections, is the so-called "intellectualist legend", the idea "that the intelligent execution of an operation must embody two processes, one of doing and another of theorising" (1949: 32). Intellectualism is false, Ryle thinks, because one may have know-how even when one is incapable of articulating the relevant rules or propositions, or even before anyone is capable of doing so. Moreover, intellectualism seems to entail a vicious regress. If intellectualism is true, then knowing-how consists in the possession of a belief. In

---

3 In this chapter and the rest of the essay I use the term 'doxastic state' rather than just 'belief' in my characterization of intellectualism. The reason is that certain intellectualists do not hold knowledge-that to be a species of belief (e.g., Williamson 2002). Nevertheless they still think of knowing-how as realized in a doxastic state of some sort. The term refers primarily to any propositionally encoded state of mind that is truth-apt.

4 Epistemologists may draw the intellectualism/anti-intellectualism distinction in other ways. For example, it may be claimed that intellectualists think know-how is a mental state while anti-intellectualists think it is equivalent to ability; or that intellectualists think that know-how is at least partly constituted by doxastic states (perhaps wholly), while anti-intellectualists deny that belief is a necessary ingredient. I think my way of drawing the distinction is more consistent with actual divisions in the literature. The first alternative way of drawing this distinction (that I have just now mentioned, i.e., mental state vs. ability) makes the two positions contraries, at best. (Mightn't one hold that knowing-how is a mental state *and* that it is an ability? I hope so; it is a position I endorse in chapter four.) At any rate, whichever side of the line the view I defend falls under, it is one that is or would be contested by intellectualists.

5 Other anti-intellectualists include Hawley (2003) and Noe (2005).

order to exercise our know how we must activate the relevant belief. But activating the relevant belief is itself something that we know how to do; and if *this* know-how is *itself* grounded in belief then the exercise of know-how requires the activation of further beliefs, *ad infinitum*; it thus becomes impossible for intelligent practice to get off the ground. Hence intellectualism must be false.

In chapters two and four we'll spend some time evaluating Ryle's negative arguments. (In spite of my sympathy for anti-intellectualism, I'll argue that the first argument is inconclusive and that the second is unsound.) At present we seek only a casual understanding of Ryle, so let us turn briefly to his positive account. Ryle holds that to know how to X is simply to be able to X, all else being equal. If I am able, say, to make moves on a chessboard that are permitted by the rules of chess, I know how to play chess. The *ceteris paribus* clause is important (Hawley 2003): if my opponent has glued the chess pieces to the board, or if I have been temporarily paralyzed, I still know how to play chess (although I am unable to play).

Intellectualist positions on knowing-how have been put forth by Hintikka (1975) Stanley and Williamson (2001), Brown (1970) and Fodor (1981).[6] All four sets of authors argue that knowing-how is simply a special kind of knowing-that. As Stanley and Williamson put it, to know how to, say, boil an egg is to know *that* w is a way of boiling an egg. (One knows *that* bringing some water to a boil, putting an egg in the water, and so on, is a way of boiling an egg.) Hintikka, Stanley and Williamson and Brown all offer linguistic arguments for their position, essentially that attributions of knowledge-how can be *syntactically and semantically assimilated*

6 See also Carr (1979; 1981), Katzoff (1984), Snowdon (2004) and Bengson and Moffett (forthcoming).

to attributions of knowledge-that. Fodor, on the other hand, draws on the nature of psychological explanation and a computationalist theory of mind. More specific discussion of Stanley and Williamson will have to wait until chapter two. Fodor's intellectualism will resurface briefly in chapter four.

Chomsky's account of linguistic competence is another paradigmatic example of intellectualism, the result of a now distant dissatisfaction with the behaviourism that dominated psychology, of which Ryle's behaviourism is, not coincidentally, a close cousin. Chomsky holds that a speaker's linguistic competence—her knowledge of how to form grammatical sentences—is realized in a "certain structure of rules, principles and representations in the mind" (Chomsky 1980: 91). In Chomsky we find a pared down version of intellectualism. Know-how is understood as structured mental representations. Whether it is a form of *knowledge-that*—i.e., whether justification or warrant is a condition on know-how—is, as Chomsky explains, a further claim to which he is unwilling to commit himself (92-93).

Ryle's anti-intellectualism is motivated in part by his behaviourism, by a reluctance to admit the presence of any "ghost-like" substance underlying our intelligent behaviour. Of course, we do not share with Ryle his skepticism about cognitive states as distinct from behavioural dispositions, and this may seem to count in favour of intellectualism about knowing-how. On the other hand, Ryle is correct in identifying an essential tie between know-how and ability. Combining a psychologistic framework for know-how and this insight of Ryle's, we may tentatively infer the neutral thesis that *knowing-how is a cognitive state that grounds our ability to engage in intelligent action.* (This notion of "grounding" is not

altogether clear. I develop a better grasp of the notion in chapter four.) As I have already noted, to know how does not entail having the relevant ability since ability-possession requires that many other conditions obtain, some internal to the agent and some external. For example, an accomplished violinist has know-how but if a tragic accident deprives her of the use of her hands, she lacks the ability to play the violin. And I know how to ride a skateboard down a railing, even if vigilant security guards rob me of the ability to do so.

The question that contemporary intellectualists and anti-intellectualists are split on is whether or not the cognitive state that grounds our ability is a doxastic mental state. Ryle's views notwithstanding, anti-intellectualists need not be behaviourists. Nor need they hold that knowing-how is *equivalent* to ability, all else being equal (although many have). Allowing that know-how is a cognitive state that grounds ability, anti-intellectualists may hold either that a doxastic state is not necessary for know-how or that it is not sufficient. Some anti-intellectualists may find encouragement in a non-representational account of certain mental states and processes. I hope to offer some insights on this matter in the penultimate chapter of this essay, without embroiling myself too deeply in debates about the nature of propositional attitudes or mental states more generally. The position I offer is anti-intellectualist. I argue that know-how's informational content need not be propositionally encoded, and, taking know-how's tie to action seriously, that it must be both representative and directive.

An answer to the question of what knowing-how is (e.g., whether it is a belief or an ability) is obviously crucial for any thorough understanding of the subject.

Another important question asks what range of activities we know how to engage in. One reason why this an answer to this second question is important too is that it is needed in order to shed light on what the phenomenon is. Knowing-how, as I said at the outset, is a certain fittingness for practical engagement with the world. But what range of activities does know-how account for? Exactly how pervasive is knowledge-how? I'll explore this issue presently by introducing another way of carving up the literature.

Our topic in this section up till now has been intellectualist and anti-intellectualist accounts of knowing-how. It is really too bad that these terms have already been appropriated for a very specific issue, since they nicely describe a broader distinction that divides epistemologists.[7] It is a distinction, admitting of degrees, that divides those, on one extreme, who view know-how as involving only voluntary, elaborate and explicitly rule-guided action, and those, on the other extreme, willing to include more automatic, simple and unconscious behaviour within the scope of know-how. An analogous division is found too in the theory of knowledge-*that*. On the one extreme we have radically internalist epistemologists who hold that knowledge is possessed by only full-fledged persons and that it is reflective, inferential and the outcome of voluntary acts of belief-formation. On the other extreme we find radically externalist epistemologists who believe that knowledge is possessed by animals and small children, and often the outcome of automatic, unconscious cognitive mechanisms.

Within this broader distinction there exists not only a divide among theorists

7 There is an analogous debate in metaethics, which suggests that a more apt label for intellectualism may be 'cognitivism' (about know-how). But I see no need to re-assign terms that have an established usage.

as to *what knowing-how is* (i.e., intellectualists vs. anti-intellectualists), but also disagreement on *the nature of the activities that we know how to engage in.* Let us use the term *"reflectivism"* to mark *this* distinction, since one very important difference between the two sides is the degree to which some sort of reflection figures in intelligent activities. Reflectivists will tend to emphasize the importance of voluntary control; we know how to engage in activities or actions only if they are, can be, or once were voluntarily performed. Each of these possibilities represents a progressively weaker form of voluntarism about knowledge-how. Related to "voluntarism" is the issue of conscious awareness. When exercising our know-how, are we necessarily conscious of the cognitive state that constitutes our know-how? Of our "will"? Of the conditions that elicit our performing the activity? Of the activity itself? Reflectivists will tend to answer yes to all these questions.

We have discovered two questions about knowledge-how to which epistemologists may legitimately direct their attention. What is knowledge-how? And what sorts of activities can one know how to perform? Of course, the possibilities for inquiry are richer than this. Some epistemologists have investigated the phenomenology of knowledge-how (Dreyfus and Dreyfus 1990; Young 2004); others have discussed the syntactical and semantic properties of knowledge-how attributions (Brown 1970; Steel 1974; Hintikka 1975; Stanley and Williamson 2001); both philosophers and psychologists have carried out empirical investigations of our concept of knowledge-how (Bengson, Moffett and Wright 2006; Tardiff et al. 2005); feminist epistemologists have brought feminist norms to bear on how we understand knowledge-how (Dalmiya and Alcoff 1993); at least one author has written on the

etymology of 'knows how' (Craig 1999); and still others have drawn various sorts of connections between knowing-how and knowing-that (Stanley and Williamson 2001; Dreyfus and Dreyfus 1990; Clark 2000; Hartland-Swann 1956).

This last topic—namely, the relation between knowing-how and knowing-that—will inevitably arise in different forms throughout the course of the essay. A guiding principle of this thesis is that the two species of knowledge can be understood in a shared theoretical framework. Nevertheless, the essay is more explicitly driven by a concern with the two questions above. I plan to argue against intellectualism, in particular, the variety that tries to assimilate knowing-how to knowing-that (chapter two). Then I defend a radical version of anti-reflectivism, more specifically, a strong form of anti-voluntarism (in chapter three). Having gained an understanding of what sorts of activities we know how to engage in, I will be in a position to articulate and defend an anti-intellectualist account of knowing-how (in chapter four).

My approach in this essay—though ecumenical—is substantively and methodologically naturalistic. The topic is know-how itself, conceived as a natural epistemic-cum-psychological *kind* (see Kornblith 2002). I am, at least in this essay, not concerned with our *concept* of know-how, except as an indirect means of understanding its referent. I do not think of myself as trying to uncover conceptual relations—indeed I am skeptical about their prevalence. Knowing-how, although it may well lack an essence, involves a unified cluster of psychological properties. I will not therefore be moved by intuitions about a small or isolated range of cases. For example, I do not give much weight to the fact that we sometimes use "knows how" in what are obviously ascriptions of knowledge of some practical proposition. In

accord with this naturalistic approach, one of my main concerns, throughout the essay, will be with the psychology of knowing-how.

1.3 The Essay

Now that I've conveyed a general sense of how I plan to proceed in this essay, let me specify in more detail the arguments in each chapter.

In chapter two I defend the anti-intellectualist idea that knowing-how is distinct from knowing-that by appealing to etiological and psychological differences between the two. Unlike knowledge-that, knowledge-how (1) does not entail justification, (2) is not always or fully realized in belief, and (3) possesses a different 'direction of fit'. The first claim needs to be suitably qualified to accommodate views that hold a certain kind of second-order justification to be a necessary condition for knowledge-that; the second finds support in scientific explanations for intelligent activities and the cognitive states referred to in these explanations; the third claim is discussed briefly and dealt with more fully in the final chapter. While all three arguments serve to distinguish knowing-how from knowing-that, (2) and (3) also count against all forms of intellectualism. To round out the chapter, I engage with Stanley and Williamson's recent attempt to assimilate knowledge-how to knowledge-that.

In chapter three I argue for anti-voluntarism about knowledge-how, the idea that we can know how to engage in activities over which we lack voluntary control. In the process I shall also have something to say about whether or not our consciousness must be engaged in some way as we perform intelligent activities. The

evidence that I marshal in favour of anti-voluntarism consists largely of various examples in which individuals seem to know how to engage in activities that are not subject to voluntary control; though I will also defend voluntarism by showing that it coheres with naturalistic account of norm-correctness, an element of the substantively naturalistic theory of knowledge that guides this essay.

These specific arguments aside, discussion of voluntarism in this chapter is located in the context of a very abstract theoretical framework that encompasses both knowing-how and knowing-that. The framework is structured by what are, I will show, two equivalent epistemological theses. First, epistemic norms that have as their object processes of belief formation and revision are part of a broader group of norms concerning intelligent activity in general. Second, acquiring justified beliefs is one of many things we know how to do.

Chapter three is primarily a defence of anti-voluntarism, remember, and in order to make sense of that position, I explain—without the sustained argument that would be required to support such positions—that we must rethink certain views of norms, agency and intentional behaviour. Epistemic norms may rightly make reference to non-voluntary activities given a realist account of epistemic norms. And we are responsible for activities not because they are voluntary but, rather, insofar as they flow from certain aspects of the self. Agency and intentional behaviour bear at best an indirect link to voluntary control. So, this chapter, besides offering some insight into the connection between know-how and voluntary control, will also serve to convey a sense of how issues in the theory of knowledge-how bear on other philosophical subjects, in particular those in normative theory and the philosophy of

mind.

The final substantive chapter of the thesis is the payoff. Equipped with a defence of the distinction between knowing-how and knowing-that (chapter two), and an anti-reflectivist understanding of what range of activities we can know how to engage in (chapter three), we will finally be in a position to say what knowing-how *is*. By confronting the link between know-how and ability, a positive account of knowing-how will emerge. Much of the chapter is spent presenting and defending a general schema for know-how, within which various intellectualist and anti-intellectualist accounts can be seen to fit. On this schema knowing-how grounds, and is a necessary condition on, the ability to engage in intelligent activity. Several objections to this schema will be discussed and defused.

In the rest of chapter four I will argue more fully for a position first floated in chapter two: that knowing-how is *a hybrid mental state*, with both cognitive and "directive" characteristics, the function of which is *to guide behaviour in accord with certain information.* Another feature of this account is a commitment to the claim that know-how is often not propositionally structured. Connectionist accounts of know-how will provide a useful model for this view, since it is plausible that the prototypes of connectionist theory both have a dual functional role and do not always encode information propositionally. As I will explain, a hybrid account also helps to account for non-voluntary know-how.

Before we set ourselves to the task I should highlight some general difficulties that lie in our way. First, the territory is relatively uncharted. Compared to the wealth of

literature on propositional knowledge, work on knowing-how is fairly sparse. This means that as we proceed we will not have the benefit of being able to rely on well-developed theories or learn from the mistakes of thoroughly refuted alternatives. Instead we will have to draw on a few important works in the theory of knowing-how, as well as ideas and distinctions in the theory of knowing-that and elsewhere.

A second difficulty for us is that the phenomenon of practical knowledge is somewhat more obscure than its propositional counterpart. While we are used to ascribing know-how to others, the underlying psychology of knowing-how is relatively unknown. Mainstream epistemology relies easily on folk psychology: we are confident that beliefs, or something of the sort, realize propositional knowledge; belief-forming processes (e.g., perception, testimony) are, perhaps, not altogether unfamiliar either. With knowing-how, on the other hand, non-epistemic folk psychological concepts are of limited help. As well, we typically lack introspective access to the relevant psychological states. Often, that is, we know how to do something without knowing how it is we do it. The topic we have chosen is not without mystery; our task is to dispel it.

Chapter 2: Knowing-How and Knowing-That

If I am to develop a positive account of knowing-how, I must first ward off a view that obviates such an account, namely, the intellectualist view that knowing-how is just a special kind of knowing-that. The view is a relevant alternative because of compelling attempts in the literature by philosophers to affect just such an assimilation (Hintikka 1975, Brown 1974, Stanley and Williamson 2001; Fodor 1981). The most influential attempt has been made by Jason Stanley and Timothy Williamson. As they say, getting clear on the issue is important not only in itself, but also because, as we have already noted, the distinction between knowing-how and knowing-that is exploited outside of epistemology, in other areas of philosophy (2001: 411). Recall: some philosophers of language think that linguistic competence consists not in having some piece of propositional knowledge but rather in knowing how to use linguistic expressions. Stanley and Williamson think, of course, that this claim and others like it rest on a spurious distinction. To know how to do something is, according to these authors, roughly to know *that* it can be performed in such-and-such a way. Later on in the chapter I will evaluate Stanley and Williamson's arguments in favour what I henceforth call "the assimilation thesis".

It is beyond doubt that propositional knowledge contributes to knowledge-how. We often gain know-how by learning some important fact (although just as often we learn some important fact and still fail to know how). A complete treatment of practical knowledge must explain how propositional knowledge figures in both the acquisition and instantiation of know-how. (I take up the issue briefly in chapter

four.) But the idea that knowledge-how is a *species* of knowledge-that is incompatible with certain epistemological and psychological truths. The difficulty is specific not merely to the assimilation thesis but, rather, *all versions of intellectualism.* Not only is knowing-how distinct from knowing-that, but to know how is not identical with being in a doxastic state.

The arguments I offer run as follows. First, knowledge-how and knowledge-that pull apart in that justification is not a necessary condition for knowledge-how (section 2.1). Second, given fairly uncontroversial explanations of behaviour in psychology, knowledge-how and knowledge-that have different realization bases: knowing-how is not realized (always or fully) in doxastic states (section 2.2). With the aim of diagnosing the intellectualists' error, I argue that an implicit or explicit adherence to epistemological voluntarism explains why some philosophers are drawn to intellectualism (also 2.2). Third, knowledge-how and knowledge-that are mental states with different *directions of fit* (section 2.3). The function or purpose of knowledge-that is just to fit the world, while the function or purpose of knowledge-how is also to make the world fit it. To round out the chapter I return, as promised, to Stanley and Williamson's argument for assimilation (section 2.4). Their argument is unsatisfying, I contend, because it draws on an impoverished set of resources.

The following sections in this chapter constitute a defence of the distinction between knowing-how and knowing-that.[8] Before we begin, however, let us discuss a fairly obvious way of drawing the distinction, that is, from the observation that we often know how to do something without, seemingly, possessing any relevant

8 The arguments here therefore also count against the minority view that knowing-that is a species of knowing-how (see, e.g., Hartland-Swann, 1956).

propositional knowledge (see Ryle 1949: 30). Colloquially, it may be expressed that one knows how to do something, although one does not know how one does it. I know how to type an astonishing 80 words per minute, put spin on a tennis serve, convince my mother to send a little extra money this holiday season, and so on. But, for different reasons in each case, I haven't the foggiest idea how I do these things. I cannot consciously entertain a proposition about how these actions are performed.

Consider further: some people even believe that an activity is performed one way when they in fact perform it in another way altogether. Many competent bike-riders think that if one is falling over on a bike one should turn the handlebar in the opposite direction. But this is false!—doing so will instead accelerate your descent. What you're supposed to do—and the same competent bike-riders know *how* to do this—is turn the handlebar in the *same* direction you are falling. Again, apparent know-how absent knowledge-that.

A more exotic example is the case of someone who suffers from anterograde amnesia. (Think of the main character in the film *Memento*.) These individuals have lost the ability to acquire declarative memories of their past—to acquire knowledge-that. Only events that happened before their accident, only facts that they have learned beforehand, may enter their conscious memory. Nevertheless, they are able to learn new motor skills, while if queried they will deny any such competence. Again, we have, at least apparently, knowledge-how without knowledge-that. This example from neuropsychology suggests another argument for the distinction, *viz.*, in a flatfooted construal of certain psychological concepts. In cognitive psychology, *declarative* memory is distinguished from *procedural* memory; the two kinds of

memory are manifested in different cognitive sub-systems. The former corresponds to knowledge-that, the latter to knowledge-how.

This way of supporting the know-how/know-that distinction is not without merit, and establishes a default presupposition in its favour. Unfortunately, however, it begs some serious questions in cognitive science. The cases of knowledge-how described above may be turn out to be case of knowledge-that, so long as the relevant propositional attitudes are, for whatever reason, *inaccessible*. That I cannot assert p, even that I claim not-p, does not guarantee that I do not believe p, *if p is the object of an unconscious belief*. Thus there is not necessarily any inconsistency between the assimilation thesis and the putative counterexamples adumbrated above. It all depends on whether competence in these cases consists in unconscious beliefs—whether, say, procedural memory, like declarative memory, *also* consists in the possession of doxastic attitudes. Now by itself this intellectualist response to the above argument may seem ad hoc. (Why think that the belief is unconscious rather than absent?) The response may be apt, however, if one is in a position to provide some independent motivation for intellectualism, for example, in a computationalist theory of mind (see Fodor 1981).

Whether intellectualism is true or false thus turns on difficult issues in psychology, linguistics and the philosophy of mind—issues that of course cannot be settled decisively here. Nevertheless, in section 2.2, below, I will offer some reason for thinking that knowing-how in these cases does not consist in the possession of unconscious beliefs. My point at present is that any argument that proposes to show that we have know-how without any relevant beliefs can not, if it is to be truly

compelling, simply appeal to common-sense intuitions about everyday cases; it must be informed by cognitive science.

2.1 Justification

Let us then turn to some more promising ways of distinguishing knowing-how from knowing-that. In this section I make plausible a *non-justificatory* account of knowing-how. (My use of 'justification' is intended to be neutral between internalist and externalist conceptions.) I then explain that even if one holds justification is necessary for knowing-how, knowing-how is a condition on knowledge-that, which entails (as we shall see) that knowledge-that requires a higher-order level of justification than knowledge-how. In any case, then, the two are distinct.

Knowing-that requires justification while knowing-how does not. More fully, it is false that knowing how to perform an action requires being justified or warranted in performing the action in the way that one does. Processes that are not justification-conferring (e.g., blows to the head and stray electrons) cannot endow an agent with knowledge-that, but they can give her knowledge-how. The way to understand this is as follows. Knowing-that supervenes on one's *mental history*; a belief counts as knowledge so long as the processes that produce or sustain it are justification-conferring. Knowledge-how, on the other hand, supervenes only one one's *current mental state*; it obtains so long as one is in a mental state that grounds the ability to engage in some intelligent activity.[9] Of course, know-how is often the outcome of

9 Apparent counter-examples to a non-historical account of knowledge-how will be dealt with in chapter four. In this context they are not very compelling, but they offer more of a threat to the anti-intellectualist account of knowledge-how I present later, and so are better dealt with at that stage of the essay.

learning or some other reliable developmental process, but this history merely supports the acquisition of know-how, as a matter of contingent fact. There is no necessary relationship between knowing-how and the processes that produce it. This then rules out any assimilation of knowledge-how to knowledge-that.

Various examples can be constructed so as to draw out the intuition that knowledge-how does not entail justification. The case of swamp-person is persuasive. For those readers that don't already know, swamp-person is a molecule-by-molecule duplicate of you who has by incredible chance coalesced, just this instant, out of a swamp. At the moment of "birth", swamp-person completely lacks propositional knowledge, since she lacks a history of justification-conferring processes. Assume that she has all the same beliefs as you; some of these are of course true. But since they are not formed by a reliable process, or on the basis of evidence, etc., they do not qualify as knowledge. Nevertheless, swamp-person would seem to possess procedural knowledge, indeed all the same know-how that you possess. She knows how to get out of the swamp, how to call a cab to take her to the city, how to do algebra, and so on. Again, since she lacks a history she lacks justification; she knows how but doesn't know that; therefore knowing-how is not a species of knowing-that.

Now while I think that my characterization of swamp-person's knowledge is correct, the case remains controversial for at least two reasons. First, some will balk at the idea that she has intentional states (because, e.g., she lacks a history of relations with the objects her mental states seem to refer to). If she doesn't, then she doesn't have know-how either. Second, swamp-person is a very bizarre case. Our intuitions,

some might think, are reliable, if they are, only with regard to familiar, everyday sorts of cases. We can't be expected to know, simply on the basis of intuition, what sort of knowledge swamp-person possesses.

I am not entirely unsympathetic to either objection, so let's consider a less controversial case. Suppose someone acquires a set of rules about how to fly a plane, not from a competent instructor or through trial and error, but by randomly guessing the function of each control in the cockpit. He memorizes and internalizes the rules such that he acquires a firm set of dispositions to handle the controls in particular ways—as it turns out, in the right sort of ways. I'm inclined here to think that he knows how to fly a plane. It's accidental, to be sure, but the accidentalness involved does not defeat his know-how. *It's an accident that he knows how to fly.*

I'm not sure whether to expect everyone's intuitions to follow mine here. Perhaps we just don't have any universal, wide-ranging intuitions about whether knowing-how does or does not require justification. Furthermore, it should be noted that the argument above presupposes a *historical* account of justification, on which a mental state's justificatory status depends on how it is has been generated or sustained. While this is the dominant view in epistemology, epistemologists who endorse a non-historical account of justification would for this reason be unmoved by the argument.

A more decisive argument can be constructed—one that does not trade on the above intuitions and that does not presuppose a historical account of propositional knowledge—given that knowledge-that requires knowledge-how. Knowing an arithmetical sum requires knowing how to add; knowing what time of the day it is

requires knowing how to read a clock (or the position of the sun, etc.); knowing many truths in experimental psychology requires knowing how to statistically analyze data. Now, it may seem that the possession of some propositional knowledge does not require exercising any relevant know-how, e.g., if one knows only on the basis of testimony. My view is that even these sorts of cases involve know-how. With regard to testimony-based knowledge, for example, we must know how to interpret others, or how to distinguish good testimony from bad, etc. One might also object that some propositions are self-evident, such that one need not exercise any know-how in knowing these propositions. But the fact that a proposition is self-evident does not entail that one need not exercise know-how in acquiring a justified belief. As I've already pointed out, mathematical truths, for example, cannot be known without exercising some mathematical know-how. In any case, even if a small subset of propositional knowledge does not entail know-how, the forthcoming argument focuses on instances of propositional knowledge that do. For now, let us assume that knowledge-that requires knowledge-how, in particular, that justification involves exercising some piece of know-how. My defence of this view will be deferred until chapter 3.

Assume, for reductio, that knowing-how *does* require justification. Then knowing-*that* requires *more* than being justified. The justificatory condition on *knowing-how* imposes an additional second-order justificatory condition on knowledge-that. Not only must one believe p *because* it's true, but one must also be disposed to form beliefs in the way that one does *because* doing so is a good way of getting, let's say, true beliefs. That is to say, to know that P, not only must one be

disposed to (say) reliably generate true beliefs—for even swamp-person may have this disposition; over and above this, one must be justified or warranted in generating true beliefs, in that one's disposition has been formed by a second-order process that reliably engenders reliable belief-forming processes. Proper functionalists, for example, hold a view of this type (see, e.g., Plantinga 1993).

But if knowing-that has a second order justificatory condition, and knowing-how is a species of knowing-that, then knowing-how has a second-order justificatory condition too. And so, since knowing-that requires exercising some know-how, and since a justificatory condition on knowing-how imposes a higher-order justificatory condition on any knowledge-that it is a condition on, knowing-that would require a *third-order* justificatory condition. We have a kind of regress on our hands. Of course, the proper functionalist, for example, should just deny that knowing-how is knowing-that, and thus that know-how also has a second-order justificatory condition. The problem for *any* epistemologist is that since the justification for a belief is the exercise of some know-how, the justificatory conditions on knowing-how must remain at one level *below* those for knowing-that. If knowing-that requires only a second-order justificatory condition, knowing-how has only a first-order justificatory condition. Hence the two must be different kinds of epistemic states.

2.2 Knowing-how and Belief

Given a non-justificatory account of knowing-how, the assimilation thesis is false. Knowing-how and knowing-that are distinct. Still, however, the threat of intellectualism looms. For all I have said so far, knowing-how may simply be a *true*

belief (sans justification). The following two arguments (covered over the next two sections) will deal with this remaining threat, and *a fortiori* with the assimilation thesis. The first will demonstrate that doxastic attitudes are not necessary for knowledge-how, the second that they are not sufficient.

Propositional knowledge is realized in beliefs. But some know-how is clearly not so realized, not even partially. Evidence for this claim can be found in a range of activities—that we know how to engage in—that are not guided by doxastic states. Instead these activities are guided by more basic psychological mechanisms. The relevant beliefs would either be an extravagance or completely unhelpful. Maintaining balance on a bike, drawing connections between ideas, recognizing colors, and so on, are all things that we know how to do without having beliefs about how to do so. Keeping balance on a bike, for example, is effected through automatic corrections to one's center of balance, guided by subliminal sensory input from the inner ears. Beliefs do not enter into this tight loop between sensation and action. Thus knowing how to keep balance on a bike does not consist in a doxastic relationship between an agent and a proposition. While activities such as keeping balance on a bike are somewhat "primitive", it nevertheless it seems that we know how to engage in them.

The idea that beliefs play no role in explanations of "primitive" activities is not entirely uncontroversial. One might insist on the adequacy of doxastic explanations for the activities in question. What I've assumed is that doxastic states are characterized by a significant degree of internal complexity, one not achieved by the states that guide the relevant activities. But it is important to note that some

deeply entrenched accounts of competence in psychology do not invoke beliefs. It is widely believed, for example, that the cerebellum, a phylogenetically old part of the nervous system, is in charge of fine motor control, including the maintenance of balance. *Doxastic attitudes are not located in the cerebellum.*

Perhaps it is only creatures that possess doxastic states that know how to do things, but this does not entail that every action they know how to do is explained by the exercise of a doxastic state. Though beliefs may often serve as constituents of a wide range of knowledge-how, intelligent creatures engage in much intelligent behavior that is not guided by beliefs. As a result of sub-personal training processes or innate developmental mechanisms, our cognitive systems have the capacity to automatically process information and produce behavior in remarkably sophisticated ways. Knowing how to perform these sorts of intelligent activities simply cannot be a matter of knowing that *because doxastic states don't figure in any good explanation for their performance*. An account that ignores our knowledge of how to engage in so-called primitive activities relinquishes the resources needed to provide a sufficiently thorough understanding of knowledge-how.

Given the argument of this section it may seem that I have no reason to reject intellectualism about a restricted range of know-how. If so, then it may be true that the cognitive state that realizes any knowledge of how to perform a less primitive, more sophisticated action is a belief. As a strict generalization, this more limited thesis deserves to be rejected too insofar as "sophisticated" know-how requires the support of "primitive" know-how. When performing sophisticated actions we rely on knowing how to engage in automatic self-correction (see Pollock 1987). For

example, take riding a bike, an activity that, I will assume, involves the activation of beliefs. To know how to ride a bike one must know how to automatically correct one's behavior in all sorts of ways; one must know how to quickly dodge obstacles, maintain one's balance, lean one way if one encounters a slope, etc. The production of *these* activities is not explained by the activation of beliefs. Thus even the performance of a sophisticated action requires primitive "scaffolding"[10]; beliefs about how to perform the act are insufficient.

Beliefs do not seem to suffice even for relatively more autonomous "sophisticated" actions. A plausible candidate for the cognitive state that realizes knowledge-how must be *action-guiding*. And it seems that a belief about how to perform an action merely *initiates* the performance, while it is *guided* by other sorts of cognitive states. Beliefs are not sufficiently fine-tuned to serve us in this regard. Hence while a belief may contribute, even essentially, to autonomous, "sophisticated" knowledge-how, it does not fully constitute it.

There is an explanation, relevant at this point in the essay, for why some philosophers might be drawn towards the assimilation thesis. Central here is *voluntarism*, the question of whether one can know how to engage in activities over which one lacks voluntary control. Voluntarism is the topic of the next chapter so I will discuss it only briefly here.

Many philosophers writing on knowing-how implicitly presuppose voluntarism about knowledge-how. For example, even those anti-intellectualists who take knowing-how to be equivalent to ability understand knowing how to X as the disposition to X if one *tries*. For many philosophers, then, the assimilation thesis may

10 This term is borrowed from Andy Clark (1997).

be an attractive position. This is because a voluntary action is one that is brought about by a belief and a desire; so if voluntarism is true then knowing-how, just like knowing-that, involves belief. This provides a *prima facie* case for the assimilation thesis.

The idea that some of the activities referred to above are such that we can know how to perform them requires that voluntarism be false. If voluntarism is true, if beliefs are required to perform activities of the sort we know how to do, then some of the above activities are not ones that we can know how to engage in *since their execution does not involve beliefs*. Intuitively, the activities in question *do* seem to be of the sort that we can know how to engage in. Voluntarism is therefore false since it is inconsistent with the existence of certain instances of know-how.

There is no direct logical relationship between voluntarism and intellectualist views like the assimilation thesis. Intellectualism might be true even if we can know how to perform non-voluntary activities—being brought about by a belief is not a sufficient condition for voluntary action. Equally, voluntarism might be true though intellectualism false; while essential to the guidance of intelligent activity, beliefs may not, for some reason, suffice to provide one with knowledge-how. Given either voluntarism or intellectualism, knowing-how entails belief. As we have seen, both views are subject to the same counter-example.

The connection between voluntarism and intellectualism, though weak, nevertheless provides the basis of an explanation for why many epistemologists have been drawn to the latter position. Both positions have currency among philosophers who endorse, so to speak, an "over-intellectualized" view of the mind—who, at the

limit, refuse to see knowledge or intelligent activity as anything less than the culmination of a process of explicit reasoning. That someone is a voluntarist therefore counts as indirect evidence that they subscribe to intellectualism.

But even if voluntarists are likely to be intellectualists, how does belief in the former explain belief in intellectualism? Two ways. First, a voluntarist thinks (like everyone else) that belief is implicated in both propositional knowledge and voluntary action. So, as we've already noted, this is why she might think that knowing-how is a special case of knowing-that. Second, a voluntarist would fail to attend to activities for which beliefs are not a driving force, as she thinks these activities are not of the sort we know how to do. So, her voluntarism shields her from a refutation of intellectualism (from its false implication that we can't, e.g., know how to keep balance on a bike). Since voluntarism supports intellectualism in these two ways, and since the two are likely to be paired with one another, we can explain a philosopher's mistake belief in intellectualism on the basis of the support offered by her commitment to voluntarism: she is an assimilationist *because* she is (explicitly or implicitly) a voluntarist.

2.3 Directions of Fit

In this section I articulate a third objection to intellectualism. Unlike knowing-that, knowing-how is a mental state with both directions of fit. The idea will be introduced only briefly here; more detailed elaboration will come in the final chapter of the thesis, where I present a positive account of know-how.

Intentional states are *representative*, in that they are meant to fit the world, or

directive[11], in that they mean for the world to fit them. In the one sort of case the direction of fit is from mental state to world; in the other from world to mental state. Anscombe's original way of introducing the distinction continues to clearly convey the basic idea (Anscombe 1957). A list of groceries may be an inventory or a shopping list. As an inventory, it is representative; it tells one what was bought; what is on the list should match what has been purchased. As a shopping list, it is directive; it tells one what to buy; what is in the shopping cart should match the list. The two directions of fit thus mark two opposing relations of normative dependence.

Propositional knowledge clearly has a representative direction of fit. At minimum it is supposed to accurately represent the world. If a potential piece of propositional knowledge—a belief—does not match the world, then it is the belief and not the world that should change. Even if we extend the relevant epistemic class to include either non-empirical knowledge or informational states the content of which is not propositionally encoded, the relevant states are supposed to represent the way things are (whatever our metaphysics or ontology). On the other hand, knowledge-that is not directive. If the world does not fit one's belief, there is no implication that the world should change to fit it.

Even *normative knowledge-that* lacks a directive direction of fit. This is an important point, since an assimilationist might hold that know-how is a species of normative knowledge-that. Some metaethicists believe that one's possessing a normative belief—that, say, abortion is wrong—entails that one is disposed to comply with the relevant norm. Does a lack of fit between a normative belief and the world

11 'Representative' and 'directive' are terms that Millikan (1996) uses to denote the two directions of fit.

then imply that the world should fit it? No. The belief that abortion is wrong does not lack fit with the world simply if a norm permitting abortion has social currency. For there to be a lack of fit here abortion must be morally permissible—in which case the normative belief ought to be changed. Insofar as a normative *judgment* is directive, it consists of a non-cognitive mental state, e.g., a desire, rather than, or alongside, a normative belief.[12]

Like knowledge-that, knowledge-how is representative.[13] For example, in order to know how to get around one must have accurate representations of one's environment. But unlike propositional knowledge, knowledge-how is essentially directive.[14] [15] Let us understand know-how as the knowledge of how to bring about some state of affairs. If the state of affairs is not instantiated, then one of two things holds. Either the putative know-how is inadequate for bringing about the state of affairs (a representative failure), or the world fails to conform to one's know-how (a directive failure). The latter possibility may obtain for any number of reasons, but our interest is in cases where one acts in a way that is incongruent with one's know-how, e.g., because of a performance error. These sorts of cases are similar to the grocery list *qua* shopping list. The world is meant to fit one's know-how, but it doesn't.

Knowledge-how is a cognitive state that grounds ability. It is thus this essential tie to action-production that distinguishes it from knowledge-that. So while

---

12 This metaethical position is not entirely uncontroversial, but it is plausible enough for our purposes.

13 I explain in chapter four that the representational content of know-how is "buried" or "implied." Propositionally encoded know-how, for example, is not truth-apt but "contains" content that is truth-apt.

14 Thanks to Rob Wilson for discussion that helped to draw out this point and to Richmond Campbell who develops an account of moral judgment as having both directions of fit (Campbell, forthcoming)

15 Millikan discusses these kinds of hybrid psychological states in her 1996.

knowledge-how may involve representations of objects and their causal relations, it also requires a connection between these representations and the structures in the mind responsible for bringing about behavior. Since knowledge-how is directive in this way (and since knowledge-that is not) the assimilation thesis is false.

In the next chapter I will have much to say about an issue I raised in section 2.2, namely, voluntarism about knowledge-how. The topic of *this* section—i.e., directions of fit—will become central to the positive account of know-how I offer in chapter four. Knowing-how, I argue in more detail, is a mental state that has a dual functional role, that is, both directions of fit.

2.4 Stanley and Williamson

Stanley and Williamson (2001) are two prominent proponents of the assimilation thesis. Although they make use of technical issues in syntactic theory and semantics, their argument for assimilation, subtleties aside, can be briefly summarized. They begin by noting that while attributions of knowledge-that express an epistemic relation between agents and propositions, some philosophers have thought that attributions of knowledge-how express an epistemic relation between agents and *actions* (e.g., Carr, 1979, 1981). So, for example, in 'Joe knows *that* the cat is on the mat', Joe is related to the proposition 'the cat is on the mat', while in 'Joe knows *how* to drive a car', Joe is related to the action clause 'to drive a car' (2001: 416-417).

Stanley and Williamson point out that the foregoing analysis of attributions of knowledge-how ignores the insights of recent syntactic theory. Like many other sorts of epistemic and non-epistemic sentences, attributions of knowledge-how express a

relation between a subject and an *embedded question.* In 'Joe knows how to drive a car', Joe stands in an epistemic relation not to the action clause 'to drive a car' but rather to the embedded question 'how to drive a car'. Deferring to Lauri Karttunen's analysis (1977), Stanley and Williamson explain that "an embedded question denotes the set of its true answers" (2001: 420). Thus to know how to do X is to know the answer to 'how do you do X?'. To know how to play tennis is to *know that* such-and-such is a way of playing tennis. In general terms to know how to do X is *to know that w is a way of doing X.*

Stanley and Williamson's paper has already received critical attention.[16] I have a general worry, not discussed in the literature, about their methodology, namely, their method of using syntactical features of language to support their account of knowledge. Essentially, Stanley and Williamson argue that from the perspective of syntactic theory there is no basis for distinguishing between attributions of knowledge-that and attributions of knowledge-how. Both relate a subject to a propositional object. While the object in 'S knows that P' is the denotation of an assertion, the object in 'S knows how to X' is the denotation of an embedded question.

Stanley and Williamson's elaboration and defence of their position is quite sophisticated. Indeed I cannot hope to do justice to their compelling paper in this short exposition. Nevertheless, the basic argument in favour of assimilation is rather

16 See, for example, Koethe (2002) and Rosefeldt (2004), who discuss Stanley and Williamson's introduction of the mysterious "practical mode of presentation", centrally involved in their account of the intensional dimension of knowledge-how ascriptions. See also Brett (1974) who offers a response to Brown's 1971 paper, a precursor to Stanley and Williamson's view, along with Brown's reply to Brett (1974).

shallow, in that it relies on a fairly limited range of evidence.[17] Allowing, for the sake of argument, that the syntax of knowledge ascriptions provides evidence about the nature of two kinds of knowledge, the approach nevertheless ignores all sorts of other considerations that would normally be canvassed by epistemologists interested in the nature of an epistemological category such as knowing-how, some of which I've already discussed. First, Stanley and Williamson do not attend to the *psychology* of knowledge-how—how it is realized, how it is executed. Second, they ignore *the criteria by which we ascribe knowledge-how*. Historical processes of justification are merely contingent antecedents of an agent's know-how. Finally, they ignore certain *paradigm cases* of things people know how to do, in particular those that involve, or require the support of, "primitive" know-how.

The lesson to be drawn here is that Stanley and Williamson's thesis—that knowledge-how is a species of knowledge-that—is under-motivated. Julian Fowke has suggested to me, in conversation, that Stanley and Williamson have, at best, achieved the basis for a negative conclusion. They have successfully undermined linguistic arguments for the distinction between knowledge-how and knowledge-that. Contrary to popular philosophical belief, a syntactical analysis of attributions of knowledge-how does not show that they relate a subject to an action. Thus, from linguistic theory no basis can be found for drawing a distinction between knowledge-how and knowledge-that. Very well, but as we have seen there are an overwhelming number of reasons *in support* of the distinction.

17 Noe (2005) takes a stance on Stanley and Williamson's article that is similar to mine. The difference is that I think their evidence is *sparse*, while he thinks it is *deficient*.

The purpose of this chapter was to clear the ground in preparation for a positive account of knowing-how as well as introduce ideas that will be discussed in more detail in later chapters. I've argued that knowing-how is distinct from knowing-that. Knowing-how does not require justification, is not always or fully realized by beliefs, and has a different direction of fit than does knowledge-that. In establishing this position I've drawn on a wide range of evidence; conversely, it is in Stanley and Williamson's relatively impoverished methodology that their argument has been found wanting.

Chapter 3: Anti-Voluntarism about Knowing-How

Although knowing-how is distinct from knowing-that, an investigation of knowing-how will benefit from the application of certain issues that have featured in the literature on knowing-that. Both, after all, are species of knowledge. In this chapter I take an important concept in the theory of knowledge-that—epistemic voluntarism—and apply it to the theory of knowledge-how. The move is a natural one, since, as I argue, acquiring justified beliefs is itself an exercise of various specific kinds of know-how. Put differently, the normativity of belief-formation is a species of the normativity of intelligent activity in general.

Few epistemologists ever broach the issue of know-how and voluntariness. Aside from Dreyfus and Dreyfus (1990), none do so explicitly, to my knowledge.[18] This chapter will serve, then, to chart the relevant theoretical territory; I also stake out some of this territory by defending an anti-reflectivist position on knowledge-how, roughly speaking, recall, a position that denies that intelligent activity necessarily involves reflective engagement on the part of the subject. In the introductory chapter I suggested that the theory of knowledge-how is linked to other philosophical subjects; a further aim of this chapter is to pursue some of these links.

In the theory of propositional knowledge it is contested whether only voluntary belief-forming processes are subject to epistemic norms. Our interest is broader: belief-forming processes are merely one of several kinds of activities that we

18 The issue of whether justification-conferring processes are voluntary, and whether this is crucial to their epistemic evaluability, is also not often discussed in mainstream epistemology. Especially compared to the close attention given to a closely related topic, internalism vs. externalism.

know how to engage in. *Our* question, in short, is this: can one can know how to engage in activities over which one lacks voluntary control? The answer, I suggest, is yes. Aside from appeals to particular cases of knowing-how, my defence of anti-voluntarism involves more theory-driven arguments about how to delineate the relevant class of activities in light of appropriate accounts of epistemic responsibility and epistemic norms. Many parts of the chapter are either substantively or methodologically naturalistic.

My task in section 3.1 will be to acquaint the reader with epistemic voluntarism, as the concept is understood in the theory of knowledge-that and as I plan to apply it in the present context. A cursory understanding of voluntary action will also be offered. In section 3.2, I argue for anti-voluntarism about knowledge-how. The case I present draws on instances of know-how from everyday experience as well as those studied by scientific researchers. In section 3.3, I examine voluntarist views in mainstream epistemology and draw out implications that anti-voluntarism has for normative theory and philosophy of mind. It will emerge that, given anti-voluntarism about knowing-how, agency, responsibility and intentional behaviour do not bear any of the necessary links to voluntariness supposed by many normative theorists.

3.1 Epistemic Voluntarism

Our topic is epistemic voluntarism (and its denial). Let us be clear from the start that we are not concerned with what is usually understood by the term 'voluntarism', the idea that one can exercise voluntary control over what one believes. Our topic is not

voluntarism about belief but voluntarism about *epistemic evaluability.* To wit, does epistemic evaluation extend properly to activities over which we lack voluntary control? Epistemic *voluntarists* hold that an activity is of epistemic value or disvalue only if it is voluntary. Epistemic *anti-voluntarists* deny this claim. All activities worthy of epistemic credit are, to be sure, intentional, in the sense that they are goal-directed or rule-following; but, anti-voluntarists say, some of these processes are non-voluntary.[19]

Like the internalism-externalism distinction, the present distinction may be drawn in several different ways. One obvious alternative formulation of voluntarism is as the claim that epistemic evaluation extends only to particular activities that we *in fact* exert voluntary control over—rather than those that we *can* exert voluntary control over. On the formulation I plan to work with, a given activity is epistemically evaluable even if it is not in fact performed voluntarily, *so long as it may be performed voluntarily.* Call the first version, then, *modal* voluntarism and the second *de facto* voluntarism. A brief word with regard to my choice of the former over the latter.

Habitual performance can lead to some activities becoming automatic even if they *can* be voluntarily performed. This is a common development as individuals gain expertise over a given activity, as with driving a car (see, e.g., Dreyfus and Dreyfus 1990). Since voluntarists are unlikely to exempt these sorts of activities from epistemic appraisal, it is better to read them as modal voluntarists. Modal voluntarism (as contrasted with de facto voluntarism) is a weaker position in that it

19 But intentional activities are ipso facto voluntary, are they not? I don't think so; I address the issue in section 3.4.

admits a greater number of activities as potential objects of evaluation. Conversely, then, modal anti-voluntarism (compared to its de facto cousin) is a more difficult position to defend. Nevertheless, I plan to argue in favour of modal *anti*-voluntarism. As we proceed in this chapter, other, weaker varieties of voluntarism will present themselves and I shall offer reasons to reject them too.

Some examples of voluntary and non-voluntary belief-forming processes will help us to get a better handle on the distinction. Engaging in a process of explicit inferential reasoning, examining *a priori* intuitions about the truth of general principles, collecting experimental evidence, subjecting data to statistical analysis, consulting others to double-check your methods—all are examples of activities that can be performed voluntarily. Examples of non-voluntary activities are less familiar but include various perceptual processes, such as the perception of color, shapes, temperature and semantic tokens, as well as more higher-order cognitive processes, such as suddenly seeing connections between ideas in different domains, accessing memories relevant to a particular context, selecting the appropriate cognitive resource for a given task, and so on. In any ordinary context (e.g., absent some futuristic bio-technology) we are unable to exercise voluntary control over the performance of these activities.

In the theory of knowledge-that, voluntarists hold that only voluntary *belief-forming processes* are epistemically evaluable (or that voluntary belief-forming processes are somehow epistemically *more valuable*—an important variation on the voluntarist theme, one we'll return to at the end of the chapter). The view is often cast in terms of responsibility. That is, it is held we are epistemically *responsible* for

some belief-forming activity only if it is performed voluntarily.[20]

What is the difference here, if any, between an activity being evaluable, on the one hand, and being held responsible for it, on the other? First, there may seem to be a latent connection between responsibility and a *deontic* conception of normativity, that is, one that makes reference to the discontinuous evaluative categories of permission, obligation, forbiddenness, etc. While it is true that deontologists (as I am using the term[21]) are more likely than other normative theorists to emphasize responsibility, there is no essential tie between the two concepts. Responsibility may, for example, be invoked by those who endorse a means-end approach to normativity, wherein one has reasons to perform activities to varying degrees, as they tends to promote some favoured set of ends. For example, let's say that I have reason to consider all the relevant evidence, relative to the end of having true beliefs. If I follow this evidential norm, I may be responsible for my success.

A second alleged difference between responsibility and mere evaluability concerns the source of what is being evaluated. One is responsible if the activity issues from *oneself*, if it is an expression of one's *agency*; whereas it may be that activities are evaluable even if they issue from outside the agent proper. For example, "activities" that are performed by my digestive system or autonomic nervous system may be "evaluable", although there is no person to take responsibility for these

---

20 Responsibility for belief, many suppose, is inherited from responsibility for our activities of belief-formation (Kornblith 1983; Zagzebski 1996: 219-231). This claim has not gone unchallenged (see, e.g., Katzoff 2001) Owens argues that this fails to make sense of the fact (is it one?) that we are directly responsible for our beliefs (2000: 85-87).

21 In ethics deontology is defined as the view that the right is independent of the good. What one is obligated, permitted, etc., to do does not depend on the consequences of the actions in questions. My use of 'deontic' is meant to refer to any view that employs the discontinuous evaluative categories of obligation, permission, etc., even if these evaluative properties are determined by the consequences of actions.

activities. This makes sense too of the widely-held belief that praise and blame go hand-in-hand with responsibility; only persons can be praised or blamed. As will become clear presently, an account of knowing-how requires a concern with evaluation of activities that are produced by *agents*, i.e., with responsibility.

In mainstream epistemology, the distinction between voluntarism and anti-voluntarism marks a divisive issue, and is sometimes as fundamental to an epistemologist's approach as is her position in the internalism-externalism debate. The distinction between so-called virtue reliabilists and virtue responsibilists, for example, can plausibly be construed as, in part, a distinction between voluntarists and anti-voluntarists in virtue epistemology (more on this in section 3.3). Surprisingly, perhaps, an increasing number of those who accept that many belief-forming processes are non-voluntary *accept* voluntarism, and hold that the belief-forming processes in question are not subject to genuine normative assessment. For these philosophers knowledge is not a fully normative concept.

Until section 3.3, let us set aside voluntarism about propositional knowledge: the voluntarism issue translates smoothly into discussion of knowing-how. If it is possible for us to know how to engage in an activity *then the activity is epistemically evaluable*. For then not only are there good and bad ways of engaging in the activity but, if *we* know how to do it then it is something we are *responsible* for. If an activity weren't epistemically evaluable—if, for example, it happened *to* us rather than issuing *from* us—then we couldn't properly be said to know how to engage in it. The converse is also true. If X is not the sort of activity we can know how to perform (e.g., it is performed by our physiological system, not by *us*) then it is not

epistemically evaluable—at least in the sense that we can be epistemically responsible for it. Thus *to potentially know how to perform an activity is for it to be epistemically evaluable.* This is why, then, voluntarism about knowledge-how is a broader issue than voluntarism about knowledge-that. The latter is a thesis about what sorts of belief-forming processes are epistemically evaluable. Voluntarism about knowledge-how, on the other hand, is about a wider range of activities (including belief-forming processes); it represents a stance regarding the range of physical, social and mental activities subject to epistemic evaluation.

A general conclusion is being drawn here, one not just specific to the voluntarism issue: acquiring justified beliefs is simply one kind of activity that we know how to engage in, one kind of activity that is subject to epistemic evaluation. Obviously, by 'epistemic evaluation' I have in mind something different than most epistemologists. Standardly, an evaluation is epistemic if and only if truth is involved. An activity has value if it leads to more, or a greater proportion of, true beliefs, disvalue if the result is instead more, or a greater proportion of, false beliefs. On my view, truth is not essential to epistemic evaluation. Physical and social activities are evaluable even if they do not lead to truth. By 'epistemic evaluation' I mean something more like 'intellectual evaluation'. An activity has epistemic value if and only if it is good in the way of being intelligent. Sometimes it is intelligent because it is truth-conducive, other times because it tends to lead to another valued end, promoting it in some efficient or clever manner. Something like truth is still involved here. In intelligent activity, there exists a "cognitive fit" between the agent and her environment.

Voluntarism is thus best understood as the view that all intelligent activities (i.e., of the sort we know how to perform) are voluntary. But so far we have said nothing about what it is for an activity to be voluntary rather than non-voluntary. This gap in our discussion may seem excusable, since all of us have a fairly good intuitive grasp of what counts as a voluntary action. But the contrast that we are most familiar with is that between voluntary and *involuntary* activities. I have not used the term 'involuntary' in my characterization of voluntarism (I have instead used 'non-voluntary') because 'involuntary' suggests merely reflexive, *non-intentional* activities. And I am after an account that allows for the epistemic appraisal of activities that are not voluntary but yet intentional. So let me then say something about voluntariness.

A standard view in philosophy, one that I will not challenge, is that an action is voluntary only if it is brought about (e.g., caused) by a belief and a desire. The two mental states, by sharing content though having different but complimentary functional roles, are able to bring about action. A belief that a jacket will keep me warm and a desire to be warm bring about the voluntarily action of putting on a jacket.

I'll say no more about the belief-desire account of voluntary action, except that it seems not to supply a sufficient condition. Beliefs and desires may unconsciously bring about action, and in that case the action is not voluntary. So, an obvious second condition for voluntary action is some sort of conscious awareness. But of what sort? It seems that neither the belief nor the desire must itself be conscious. I reach for my jacket voluntarily, though I am not conscious that the jacket will keep me warm, nor do I feel a desire to be warm.

Of what are we conscious in voluntary action? Voluntary actions are supposed to be those in which we are exercising our wills. So, another possibility is that we are conscious of our will, or of a decision or intention to perform the action. I think that this too is false to the phenomenology of voluntary action. Perhaps we are necessarily conscious of some of the conditions that elicit a voluntary action. Maybe. But it seems that we may perform a voluntary action without knowing why we've done it.

Our awareness in all voluntary action has at least one object, namely, the action itself. For example, if my action of putting on a jacket is voluntary, I am necessarily aware of my action of putting on the jacket. Otherwise my behaviour is automatic. In cases where we are not, at least to some degree, aware of the action itself we are engaging in non-voluntary activities. Notice that one and the same event can be voluntary under one description and non-voluntary under another. In Davidson's case, someone voluntarily turns on the lights but does not voluntarily alert the prowler (Davidson 1980).

In sum, being brought about by a belief and a desire and consciousness of the performance are necessary conditions for normal cases of voluntary action.[22] This second condition will prove important in the next section as I marshal a case against voluntarism.

We are now ready to argue directly for anti-voluntarism about knowledge-how. A close study of intelligent behaviour shows that of the many things we know how to do, some are clearly not under our voluntary control. In section 3.3, I'll reflect

22 They are probably not sufficient however, since a negative condition regarding the absence of external influence, e.g., coercion, is probably needed.

on the opposition to anti-voluntarism in mainstream epistemology. If it's true, as I've argued, that belief-forming activities—of the sort that confer knowledge or justification—are one type of activity that we know how to engage in, then arguments in support of voluntarism about knowledge-that will be directly relevant to anti-voluntarism about knowledge-how.

3.2 Anti-Voluntarism about Knowledge-How

My strategy in this section will be to discuss various cases in which people seem to know how to engage in non-voluntary activities. Where this kind of argumentative strategy is employed in this essay—of appealing to intuitions about examples—I will look to three broad kinds of know-how. They involve, respectively, purely physical actions, social intercourse, and mental activity. Since these three classes exhaust the entire range of activities one can know-how how to engage in, it can be reasonably supposed that the conclusions I draw using these cases will be entirely general.

It is not an accident that a significant range of knowledge-how is non-voluntary. Typically, one of three conditions explains why a given activity lies outside our voluntary control. In the first two kinds of cases, the know-how is either innate or learned unconsciously. Either (1) the activity is too complex or subtle to be consciously guided, or (2) our performance of the relevant activity need not be altered in response to contextual factors. So for reasons of either difficulty in execution or mental economy, the task is delegated to a non-voluntary part of our cognitive system. In the third kind of case, there is a need for speed. While the know-how may be learned consciously, it must eventually be performed at an instance's notice. The

performance becomes non-voluntary—and indeed *can not* be performed voluntarily—because the mental faculty responsible for voluntary action is too slow to execute the action effectively. Each example of non-voluntary know-how that is described below is an instance of one of these three cases. Where appropriate, I will provide an expanded explanation for the need for non-voluntariness in each case.

Given what it is to perform an action voluntarily, voluntarists are committed to the claim that if an action is either not brought about by a belief and a desire, or if the agent is not conscious of the action, then the action is not epistemically evaluable. In chapter two, I argued that there are cases where know-how is exercised without the help of belief. These cases served as counterexamples to both intellectualism *and* voluntarism. In this chapter my strategy is different. First, I describe activities that seem, on their face, to be non-voluntary. In addition, I attack voluntarism's second corollary. That is, I describe cases in which we know how to engage in actions—they are epistemically evaluable—although we are *unaware* of our performing them.

The case for anti-voluntarism is perhaps strongest when one considers mental know-how. Our minds are far too complicated for us to consciously manage all of its operations. Perceptual processes are a clear case in point. We automatically sort all kinds of complex stimuli in our environment; in doing so we are obviously following some classificatory scheme, but we are not doing so explicitly. Another straightforward example is the phenomenon of suddenly grasping something, e.g., a connection between two ideas, an alternative explanation for an event, the reason so-and-so has some character trait, etc. These mental events occur spontaneously, but they do not merely happen *to us*. There are those who we esteem as brilliant because

they are prone to various sorts of spontaneous insights.

A surprising amount of the know-how we exercise when we engage in social intercourse is outside our voluntary control. Some of us know how to compel agreement, or sympathize with those around us, or impose our dominance over others, etc. Each of these acts are themselves voluntary, but in order to secure this voluntary goal we engage in many "subtasks" that are *non-voluntary.* Let us look closely at a particular case, one described by Malcolm Gladwell in an article from the *New Yorker* (2006). Gladwell's subject is Cesar Millan, star of the TV show *The Dog Whisperer.* Millan is famous for being able to tame "problem dogs"; on his show he visits desperate owners in order to help them gain some control over their pets. The article covers various aspects of Millan and his occupation; but what concerns us is how Millan is able to influence the attitude of the dogs he works with by subtle physical gestures of which Millan himself is apparently unaware. Gladwell asks two "movement analysts" to watch tapes of Millan interacting with dogs (with the sound muted). The analysts point out several seemingly minor physical movements that they believe partly explain Millan's power over the otherwise unruly animals. The pace of his gait, his posture, the position of his hands, the way Millan makes eye contact with dogs, the rhythm and simplicity of his movement—all of these actions conspire to induce calmness and submission in the dogs Millan encounters. These actions are the sort of thing that dogs (and humans) pick up on—in a Gestalt-like mode of apperception—but only professionals who study movement can explicitly identify each of these actions and explain their significance. Although Millan is equipped with an impressive understanding of his talent, clearly he is no more aware

of many of these particular movements than would be any layperson observing him. Millan knows how to influence the attitudes of dogs, but a significant part of this know-how is non-voluntary. He has trained himself to understand and engage with the canine psyche, but what he learned how to do is partly outside his voluntary control.

Someone may object, with regard to cases like this, that the activities are in fact voluntary, despite appearances. Millan's actions, for example, are voluntary; it is just that I have provided descriptions of his behaviour on which they aren't voluntary. Millan voluntarily assumes an intimidating posture towards the dogs on his show, but if this act is described at a finer, physical and physiological level, it will come out as non-voluntary. The objection fails on two points. First, it is not clear subtasks described above are identical with whatever Millan is doing voluntarily. Some of these tasks seem to precede or succeed whatever action it is that Millan has voluntary control over, and thus must be numerically different. Moreover, however, even if the actions are identical, the level of description at which it is voluntary is too vague. If we remain at this level it is a complete mystery what it is that Millan is doing well. The really intelligent stuff is happening beneath his consciousness. Take another example. Say, in forming a belief I am aware that I am basing it on "evidence" but I am completely unaware of *any* more specific reason for holding the belief. In this case I am aware of *something*, but not of the reasons for my belief. What confers justification lies beneath conscious awareness.

Let's get back to discussing know-how of the sort that Millan possesses. Many of the sorts of subtle physical movements that are integral to any finely

developed skill are of the sort that were *never* performed voluntary. Instead they arise through innate developmental mechanisms or are learned unconsciously. As I mentioned in chapter two, it is typically our cerebellum rather than our voluntary motor system that is, often from the very beginning, in charge of this kind of behaviour. This provides a reason then to reject another version of voluntarism that has yet to mentioned. Someone might accept that we know how to perform activities that are not voluntarily performed (thus rejecting de facto voluntarism) and those that can not be voluntarily performed (thus rejecting modal voluntarism). But she may be a voluntarist in the sense that she holds we must at least at one time have been able to perform these activities voluntarily. Call this *ontogenetic* voluntarism, if you like. As Millan's very typical, finely developed skills demonstrate, however, this version of voluntarism is false too.

Turning to purely physical know-how, we also find cases where individuals have been trained to do things although the performance is not voluntary. Sports are a good example. An athlete learns the right way to do something—she knows because, e.g., her instructor provides approbation, or she notices herself moving more quickly, or she sees the ball go over the net, etc.—and continues to act in that way although she is not completely aware what she is doing. Take the game of tennis. Like many merely competent tennis players, I taught myself how to play the game and then later on took lessons from a tennis instructor. On my own, I learned how to place "top-spin" on my forehand-side groundstroke. If asked I might tell you that I did this by changing slightly the angle of my racquet and rolling my wrist as I made contact with the ball. Later on, an instructor told me that when I hit a "top-spin" shot

I was starting my stroke with the racquet head much lower to the ground. I was surprised on hearing this news, since I did not know I was doing this. Nevertheless my hitting a "top-spin" shot involved, among other things, knowing how to swing my racquet from a low position to a high position. In his book *Blink*, Gladwell talks briefly about tennis, and interviews the tennis pro, Vic Braden. Braden tells him that although most tennis players think that they are rolling their wrist as they make contact with the tennis ball, in fact digitized computer imaging shows that they do not turn their wrist even an eighth of a degree during contact; the wrist turns only well *after* they make contact (2005: 67-68). So I think that hitting a "top-spin" shot involves turning my wrist—I *try* to do this—but what I'm actually doing is keeping my wrist straight. What I know how to do here is precisely not what I voluntarily attempt to do.

Gladwell's book *Blink*, while lacking the argumentative rigour of philosophical literature, is a good resource for externalist epistemologists. It is a book, as the subtitle reads, about "the power of thinking without thinking." To a lesser extent it is also useful for epistemological anti-voluntarists. It discusses a number of cases where we know how to do something without being aware that we are doing it. Early on in the book, Gladwell discusses an experiment in which University of Iowa researchers present subjects with a gambling game in which they are asked to turn over cards from either a blue deck or a red deck (2005: 8-10). From each card that is turned over the subject will either win or lose money. One of the decks has better payoffs—smaller rewards but fewer losses. It turns out that the subjects "figured the game out before they realized they had figured the game out:

they began making the necessary adjustments long before they were consciously aware of what adjustments they were supposed to be making" (2005: 9-10). The subjects' gambling know-how is thus, at an early stage, outside their voluntary control.

Voluntarism about knowledge-how does not fit what we know about the production of intelligent behaviour. In the next section, I turn to the literature on voluntarism about belief-forming processes, which, as you'll recall, is just a special case of voluntarism about knowledge-how. Forming true beliefs is just one thing it is possible for us to know how to do. Responding to the arguments of voluntarists in this literature will serve to consolidate my anti-voluntarist position. As well, we'll find that anti-voluntarism pushes one towards certain views on the nature of epistemic norms and epistemic agency.

3.3 Know-How, Responsibility and Agency

Few epistemic voluntarists hold that belief-forming processes must be voluntary in order to confer *justification* or *knowledge*. It is agreed by nearly all parties that we simply don't have that sort of control over cognition. (Think again of visual perception.) Epistemic voluntarism, then, as applied to *knowledge-that,* is the claim that only voluntary *belief-forming processes* are subject to normative assessment. To the extent that belief-forming processes are non-voluntary they are *not* of the sort that are subject to epistemic norms. More like digesting than like thinking, one might say. Typically, as I've said, a more specific claim is made: that we are epistemically *responsible* for our activities, or the beliefs that we thereby produce, only if these

activities are voluntary.

Now, to be clear, unlike many voluntarists, I do not wish to understand epistemic normativity within a deontic framework. Nevertheless, a significant argument for voluntarism that I wish to discuss can be understood more generally, so as to apply to any sort of approach to normativity involving responsibility. Essentially, the argument can be understood as making an appeal to some generalized version of the ought-(or should)-implies-can-principle. One is responsible (epistemically, morally, etc.) for an activity only if one has control over whether or not she performs it. Hence the activity must be voluntary.

In a moment, we'll take a look at how this argument plays out in virtue epistemology, in the hands of "virtue responsibilists". In general, virtue responsibilism contrasts with virtue reliabilism in two ways: the former counts as epistemic virtues only those character traits (1) that we voluntarily exercise and (2) that have "internal" identity conditions. That is, with regard to (2), whether a trait is a virtue or a vice depends on the internal features of the trait, e.g., whether if it involves a desire for truth. Virtue reliabilists, on the other hand, identify virtues as those that, first, need not be exercised voluntarily and, second, that have "external" identity conditions, e.g., those that are in fact truth-conducive. So paradigmatic responsibilist virtues include open-mindedness, diligence, intellectual care, etc., while reliabilist virtues range over such cognitive faculties as visual acuity, certain kinds of memory, spontaneous creativity, etc.

Reliabilists, including virtue reliabilists, tend not to discuss the notion of epistemic responsibility in their work; it would seem that they view it as not being

implicated in a wide range of belief-forming processes. The view that I am pushing for in this chapter is that since forming true beliefs is something *we* know how to do, reliabilists *should* think of engaging in non-voluntary and non-reflective activities as something that we can take responsibility for. As we proceed in this section, we'll see what this view entails.

Arguments in favour of voluntarism pop up in the work of many virtue responsibilists, such as Montmarquet and Zagzebski. Consider Montmarquet's argument (2000: 135-6). Reliabilists, he says, are not motivated by any concern for holding agents responsible for their epistemic traits. For clearly "one is not to be blamed for poor visual acuity" (136). But does this imply that responsibility does not accrue to an agent for her perceptual dispositions? Clearly it would be inappropriate to blame or punish Alice simply because she was born nearsighted. But equally, being a thorough researcher (or, more generally, having any sort of motive towards the truth) is also something that is thrust upon us: a dimwitted person can hardly choose to be more thorough in his research. Montmarquet's point just shows that *outward* praise and blame are often not appropriate in cases where someone is epistemically "vicious" or has formed beliefs in an inappropriate manner, and not just when the relevant disposition can not be exercised voluntarily. But this is plausibly because we don't think the relevant values here are significant enough to warrant outward praise and blame. To the extent that there is an argument here, Montmarquet seems to be making the familiar argument that we should not be held responsible for traits the *possession* of which is not under our control. But then this argument rules out responsibilist traits as well. We are, therefore, without an argument for the idea

that an epistemic trait is a locus of responsibility insofar as whether one exercises it is under one's own control.

Taking our cue from well-known discussions of free will and responsibility (see esp. Frankfurt 1971; Watson 1975; Wolf 1990), it is plausible that one is epistemically responsible for some activity to the extent that it partly constitutes oneself *qua epistemic agent*. The question then is whether non-voluntary dispositions (i.e., dispositions to perform non-voluntary actions) help to constitute epistemic agents. An affirmative answer to this question will be offered below. First, however, let us turn to Zagzebski's more nuanced voluntarist view, since for Zagzebski there remains a slightly more indirect, but essential, connection between being an epistemic agent and voluntariness.

Zagzebski accepts that non-voluntary belief-forming processes can be epistemically evaluable—but only if they are subject to some sort of second-order voluntary control. Ah, *another* form of voluntarism! Call this one *second-order* voluntarism. To illustrate, Zagzebski admits that many perceptual processes are non-voluntary, but insists that "[i]f an intellectually virtuous agent had indications that her perceptual ability or her perceptual situation was in some way deviant, she would withhold or withdraw perceptual judgment until she could investigate" (2001: 152). Is this true? Can an agent exert second-order control over her non-voluntary belief-forming processes in this way? Hume didn't think so, at least not if the control is supposed to be systematic. Before I answer this question, however, I should point out that Zagzebski's example is one in which an agent exerts second-order voluntary control late in the perceptual process. We can not help classifying stimuli in various

ways for example, although we may be able to suspend our *judgment*. The other way of construing the example is as one in which an agent corrects her non-voluntary belief-forming process after the fact; this sort of correction is clearly not a second-order process and so must not be what Zagzebski has in mind.

So, do we have the necessary second-order control over mental processes? I don't think so. Many of our mental activities are spontaneous; no amount of effort can prevent their occurrence. Try to directly prevent yourself from noticing some analogy! Often the relevant processes are too fast for us to intervene.[23] Now it is true that intensive training can change otherwise ingrained mental tendencies. But the likelihood that training may interrupt unconsciously learned abilities is low, and that it will interrupt innate abilities perhaps close to nil. Zagzebski's view is especially implausible when it comes to know-how, that is, if it is to apply to physical and social activities as well. The only sort of control we have over fine physical movements is to suppress the larger action patterns during which the former are expressed. I can exert voluntary control over nearly any of my activities if I choose to keep perfectly still or ingest a paralytic; but obviously this is not the sort of second-order control that is required by second-order voluntarism.

I conclude that there is no version of voluntarism about knowing-how in the literature worth defending. Zagzebski's view is interesting however since it represents an attempt to improve upon modal voluntarism; her view is also interesting for me since, like her, I am interested in epistemic responsibility and agency. In order

23 This counter-argument may not go through if Zagzebski's view is that we have second-order control not over our what we believe, but rather over what we accept. Acceptance, like judgment, is a reflectively endorsed, perhaps meta-cognitive state. The view that knowledge is acceptance rather than belief is, however, likely subject to it's own difficulties.

for someone to know how to do X, they must be responsible for doing X. In order to be responsible for some activity it must flow from a disposition that is partly constitutive of oneself. If the issue is *epistemic* responsibility then the dispositions must partly constitutive of oneself *qua* epistemic agent. A common view is that an agent is constituted by dispositions that express voluntary actions. Zagzebski thinks instead that these dispositions must be regulated by voluntary second-order dispositions. If my arguments in this chapter are sound, both of these views are wrong. We can be responsible for, say, belief-forming processes even if they are non-voluntary. But then which dispositions constitute oneself as an epistemic agent?

An answer to the question is far beyond the scope of this essay. Even David Owens, in his book-length treatment of epistemic normativity and control, fails to offer a substantive defence of his answer to this question (2000: ch.8). He suggests that responsibility requires being responsive to reasons (see Dennett 2003 and Wolf 1990 for similar views). An obvious response to this position is to point out that people who are irrational and unresponsive to reasons may still be responsible. I do not say, though, that this account of responsibility is wrong or not worth pursuing. Owen's book presents a compelling case against the "juridical" model of responsibility. The courts link culpability and control, but this may apply only to a restricted range of cases, e.g., those in which punishment is appropriate.

So I won't be offering a theory of agency, responsibility or intentional action, but I will mention one or two theories that remain live options. Perhaps various dispositions towards the world contribute to one's epistemic agency insofar as they are functionally related to one another (thus ruling out mere reflexes). Yet another

possibility is that there is a still more distant relation to voluntariness. For example, perhaps a disposition is partly constitutive of one's agency so long as it is functionally related to voluntary action, either "laterally" or as a "sub-task". I would not necessarily be averse to the latter possibility. My point in this chapter is simply that considerations about know-how push one away from standard views of epistemic agency. It is worth noting, however, that an increasing number of philosophers have come to understand agency and identity as something that lies beneath or beyond our conscious experience.

If I'm right, anti-voluntarism about know-how has certain controversial implications for other philosophical subjects. While I am unable to pursue these issues in depth, I should at least defend the theory itself from objections. One apparent problem for anti-voluntarists is the difficulty of explaining how the norms governing non-voluntary activities can be correct or incorrect. On two popular accounts of what makes epistemic norms correct or appropriate—both of them constructivist—this is simply not possible. First, some epistemologists think that a norm is the sort of thing that an agent can consciously follow; he or she can appeal to them in order to justify her beliefs (e.g., in the face of skeptical threats). What makes a norm correct is that it is justified in a process of reflective equilibrium. Second, other epistemologists think that epistemic norms are correct in virtue of social conventions; they are rules or principles that we agree to, either explicitly, or implicitly in virtue of our dialectical practices. But if there are norms with regard to non-voluntary activities, they are not of the sort that are justified either under reflection or socially, since we are typically unaware of them. Hence anti-

voluntarism is incompatible with two of the leading accounts of epistemic norms.

Another account of norms is needed. Fortunately, there is one and it is eminently plausible. The account belongs to naturalistic, realist normative theorists. On this view, a norm is a statement that affirms the appropriateness (or inappropriateness) of engaging in some activity. The norm is true if and only if the activity is appropriate (or inappropriate).[24] What makes an activity appropriate is typically a matter of its instrumental value, and may depend ultimately on the motivational states of the individual or individuals involved. So, for example, an activity may be appropriate because it promotes ends that one values highly, or because it promotes the having of true beliefs, something which itself promotes one's valued ends (see Kornblith 2002: ch.5). (This reliance on motivational states does not preclude the account from being a realist one.) Here we have a relatively tractable account of what makes for correct epistemic behaviour; although its truth does not constitute a reason to accept anti-voluntarism, it provides a way of making sense of the latter view. It also helps to locate know-how within a fully naturalistic framework, in which knowledge is understood as a natural kind.

Earlier, I mentioned (parenthetically) a view that is a close cousin of voluntarism. One may accept that non-voluntary activities are epistemically evaluable, yet hold that they are epistemically *less valuable*: voluntary action is how we truly express ourselves as sophisticated cognitive agents. I'm skeptical of this view too. In the face of the arguments of this chapter, it is hard to maintain this

24 Norms are not always statements.; they are formulated sometimes as rules. In that case, one should unpack the account in terms of, say, something like correctness-conditions rather than truth-conditions. I choose to focus on one kind of normative utterance (a cognitivist reading of them) in order to simplify.

position without simply relying on the traditional assumption that voluntary and conscious behaviour is central to agency, personhood and intentional behaviour, and thus begging the question. If, as I've suggested, non-voluntary dispositions help to constitute oneself as an epistemic agent, then why should the activities they produce be any less valuable? Many non-voluntary activities (e.g., in perception) are far more sophisticated and complex than voluntary activities. Others are of great importance for the satisfaction of our epistemic goals—they have a greater instrumental value. So how does an activity's voluntariness make it more valuable? I do not say that the "voluntarists" are wrong here, but a substantive argument is owed; the conclusion is neither trivial nor obvious.

In this chapter I've introduced the notion of voluntarism about knowing-how (actually, several related notions) as a way of marking various divisions in the still-new field of the theory of knowing-how. I've argued that several varieties of voluntarism are false, that voluntarism about knowledge-that is a special case of voluntarism about knowledge-how (since to acquire justification is to exercise some know-how), and that knowing-how to perform an activity entails being responsible for it. If anti-voluntarism is true, we are driven towards rethinking standard views of epistemic responsibility and epistemic agency on which they bear a close tie to voluntariness.

Chapter 4: A Hybrid Account of Knowing-How

In chapter two I argued that intellectualism provides an inadequate account of knowing-how. Knowing-how is therefore distinct from knowing-that and merits an account of it's own. The assimilation thesis—the thesis that knowing-how is a species of knowing-that—has, in a way, been turned on it's head: the theory of knowing-that is embedded *within* the theory of knowing-how. Acquiring justified beliefs is just one thing we know how to do.

In chapter three I argued that the activities we know how to engage in are sometimes non-voluntary and unconsciously performed. I suspect that some will remain unconvinced. It would therefore be worthwhile to defend my anti-reflectivist account of know-how by developing a theory that explains the scope of knowing-how.

More pressing, however, is the need for a well-developed alternative to intellectualism. Instead, then, this chapter will consist in a positive account of knowing-how, one that takes seriously its link with action. Knowing-how is a mental state, I argue, that has *both directions of fit* and that is often *not propositionally structured.*

First, I outline a schema or framework for know-how. While there is intense disagreement on what propositional knowledge is, there is widespread agreement on a general schema for knowledge-that. To know that P is to have a justified true belief that P. Plus maybe some fourth condition, such as that there are no undefeated defeaters for one's belief. But there is no established JTB+-like schema for knowing-

how, no schema within which a general understanding of know-how can be gained and disagreements can be couched. I offer in this chapter a candidate schema: knowing-how is a mental state that grounds (in a normative sense yet to be explained) the ability to engage in intelligent activity. Several objections to the schema must be addressed: an entire section is reserved for discussing the application of Ryle's famous regress argument against intellectualism.

After presenting and defending a schema for knowing-how in sections 4.1 and 4.2, I proceed to give a more theoretically-loaded positive account in section 4.3 and 4.4. In section 4.3 I explain more clearly my hybrid account of know-how. As I have it, knowing-how is a mental state that is both representative and directive—although it is not a motivational state. With regard to its directions of fit, knowing-how is analogous to informationally-loaded imperatives of the following sort: 'do X by doing Y'. Hence, although know-how is not truth-apt, there is a truth-apt content "buried" in propositionally encoded know-how of the general form 'X can or should be done by doing Y'. In section 4.4 I make explicit the evidence in favour of the account of know-how. I discuss too a connectionist account of know-how, which explains how know-how can lack propositional structure and further illustrates its dual functional role.

4.1 Schema for a Positive Account of Know-How

Ryle, remember, thought that knowing-how amounts to the physical ability to do something, to follow some set of rules. If Ryle is interpreted in this simplistic way, his view is vulnerable to an obvious class of counter-examples. Often one knows

how to do something without having the physical ability to do it. For example, adverse environmental or social conditions may undermine one's ability. Or while the spirit is willing, the flesh may be weak.

Perhaps we are not being sufficiently charitable to Ryle. Hawley (2003) offers a Rylean, anti-intellectualist account of know-how that avoids this objection. To know how to X, she argues, is to have the ability to X in a certain class of potentially counterfactual situations. (More accurately, for Hawley, it is to have success at doing X in those situations, if one tries.) One knows how to drive if and only if one is able to drive when road conditions are normal, when one is in control of one's limbs, etc.

I think a natural reaction to Hawley's account of know-how is to feel some dissatisfaction. Not because it is false. I think she's right, more or less. Nevertheless, because Hawley's account of know-how is given in terms of a subjunctive conditional, it fails to shed any light on what know-how *is*, on what categorical state of affairs is associated with know-how. The reason why a purely Rylean approach is misguided is that it is *apsychologistic*. Many of the most illuminating epistemological theories provide psychological conditions for knowledge. It is a certain kind of belief, formed in such-and-such a way, etc. Similarly, a full account of know-how must be couched in psychological terminology (although this is not to say that all insights into know-how must make reference to psychological states or properties).

I have argued that intellectualist accounts of know-how fail. But they do fare well in meeting the psychologistic requirement, at least potentially. Stanley and

Williamson, for example, hold that knowing-how is a species of knowing-that. If that were true, knowing-how would be in a position to inherit any number of psychological characterizations that apply to knowing-that.

A major contention in this essay is that know-how has it's own peculiar psychology, different from propositional knowledge. The schema for knowing-how that I wish to outline in this section, however, is, at least on its face, neutral between anti-intellectualist and intellectualist views of knowing-how (including the assimilation thesis).

There is, I think, a general characterization of know-how that is quite plausible. While knowing-how to X does not entail having the ability to X, knowing-how bears an essential tie to ability. Knowing-how is a *necessary* condition on having the ability to engage in intelligent action (see below for a defence of this claim). But as the standard critique of Ryle illustrates, many conditions must obtain for one to have the ability to do something. Of these conditions, which constitute one's know-how? The normative issues discussed in chapter three are helpful here. Our knowing how to X is the basis for our intellectual responsibility for X; it is the condition on ability that is partly constitutive of our intellectual agency. So, knowing-how is a mental state that grounds in this normative sense our ability to engage in intelligent activity.

This schema for knowing-how has the effect of reorienting our view of the formal relationship between knowing-how and knowing-that. It is natural to suppose that the two are analogues: knowing-how is, in the practical domain, the equivalent of knowing-that in the theoretical domain. But this is not quite right. Knowing-how is

more like justified belief than knowledge. *For knowing-how is to ability as justified belief is to knowing-that.* What turns a justified belief into knowledge is truth. Certain conditions that are external to the belief must obtain, namely the state of affairs to which the belief corresponds. Similarly, what turns an instance of know-how into ability are certain conditions that are external to the know-how itself (and external to the agent proper). In general, they are either related to the physical body of the knower or her environment. A marksman's knowledge of how to shoot gives her an ability to shoot so long as, among other things, her limbs are working properly (bodily condition), the air is clear and calm (environmental condition), etc.

Again, knowing-how is a mental state that normatively grounds the ability to engage in intelligent activity—in the sense that it is the condition on ability that lies within the agent, that is the basis for her responsibility in acting the way she does. Ability entails know-how, on this framework. This entailment, while perhaps intuitive, has been challenged by several philosophers (e.g., Carr 1981; Hawley 2003; Snowdon 2003). David Carr argues that "an agent may perform a task of considerable complexity without knowing how [to do it]" (1981: 53). A novice may perform a difficult stunt by luck or chance. His ability to do so, however, does not count as an exercise of know-how, while the same stunt performed by an expert does. "Since the novice actually performed the feat one can hardly deny that he was able to do it (in the sense of possessing the physical power)" (53).

Carr's parenthetical qualification is telling. The sense of 'ability' on which it entails knowing-how is not that of merely possessing the physical power. In that sense I have the ability to serve fifty aces in a row. Clearly this is not how Ryle and

Hawley understand ability. By 'ability' I have in mind the disposition to *reliably* perform a task. A novice lacks ability in this sense.

Carr's objection is not so much refuted as dodged. In my schema for knowing-how, to be able to do X is not to have the mere physical power to do X, but to have some degree of reliability in doing X. [25] Now I will not be moved by the objection that Carr's sense of ability is the more natural one. I am not interested in which sense of ability corresponds to the way the word is more commonly used. If, in my use, the concept of ability is being refined or altered, so be it. Theory development often calls for the refinement or alteration of concepts in order to more precisely classify and relate the objects of study. Whether this construal of ability is warranted depends not on ordinary language intuitions, but on the explanatory power of the theory offered in this essay.

Let us turn now to an objection that at least purports to be concerned with the reliability sense of know-how. Different versions of the objection are rehearsed by Carr (1979), Hawley (2003) and Snowdon (2003). In Carr's example, a dancer is able to perform a dance routine that happens to be identical to another that she has never heard of. She is able to perform both of them but knows how to perform only the first. Hawley's example involves someone, Susie, who annoys her friend, Joe, by tapping on her cigarette box, although she thinks that it is her smoking that annoys him. According to Hawley, Susie is able to annoy Joe although she does not know how to annoy him. In both cases, we seem to have ability without know-how. Hence

25 With regard to reliability, there are probably some contextual dimensions to the possession of ability. Perhaps, for example, for me to have the ability to use a web-browser I must have near 100% reliability at doing so, while a master archer has the ability to hit a bull's-eye at 100 yards even if she is accurate only 25% of the time.

ability does not entail know-how.

Carr argues that ascriptions of know-how relate an agent to an action, and furthermore that the identity conditions of the actions referred to in these ascriptions include the aim or purpose of the action (408). This explains why the dancer does not know how to perform a dance she can not aim to perform, and why Susie does not know how to annoy Joe: it is not the case the she aims to annoy him by tapping on her cigarette box.

It is curious that Carr distinguishes know-how and ability in this way. For it would seem that abilities are *also* related to actions identified in part with their aim or purpose. Apparently, Carr thinks of the actions an agent has the ability to perform as identical with purely non-intentional, physical events. (In Davidson's famous example, then, one knows only how to turn on the lights, but has the ability to alert the prowler as well.) I can't see why this should be so. For if the goal or purpose of an action, X, is not essential to X, then the ability to perform X hardly seems intelligent. Remember it is only intelligent actions that we know how perform. *Pace* Carr, actions one has the ability to perform are identified in part with their aim or purpose. So, a dancer lacks the ability to perform a dance she can not aim to perform; Susie lacks the ability to annoy Joe by tapping on her cigarette box. Of course, both have the physical power to do so. Susie actually does manage to annoy Joe, but since it is not her aim to annoy him by tapping on her cigarette box—since, that is, she can not intentionally do so—she neither knows how nor has the ability to do so.

Even if ability entails knowing-how, one might object to the schema offered here on the grounds that knowing-how is not sufficiently distinct from ability. For

example, it is argued that sometimes knowing-how entails having the relevant ability. If one knows how to do arithmetic, for example, then necessarily one is able to do arithmetic (see, e.g., Bengson and Moffett, forthcoming). I'm not convinced this entailment ever holds. Persistent performance errors may plague even mathematical competence. Philosophers who advance this position are, I believe, guilty of what I call "the analytic fallacy". The fallacy is to identify an analytic connection where there is really only a very tight contingent relationship. Inferring that someone is able to do arithmetic from the claim that they know how is perfectly legitimate, but only because it is intuitively very likely to be true, not because the latter follows deductively from the former.[26] Even if these philosophers are right, however, there is no obvious difficulty for our schema. Just as it may be that some beliefs that, necessarily, when justified are true, thus counting as knowledge (i.e., certain kinds of a priori truths), there may be instances of know-how that are necessarily instances of abilities.

4.2 Ryle's Regress

Knowing-how, I've argued, is a mental state that normatively grounds (and is a necessary condition on) the ability to engage in intelligent activity. Different accounts of knowing-how—both intellectualist and anti-intellectualist—can be seen to fit within this framework. Before I present my own account of know-how within this framework, I must first deal with an argument of Ryle's. The argument purports

26 Bruce Hunter has suggested in conversation that there may be some confusion here between the semantics and pragmatics of know-how. That someone has the ability to X, in these cases, is not entailed by their knowing how to X, although it may be conventionally implied (in Grice's sense). I find the suggestion interesting, but I am more inclined to understand the belief as a mistake about what knowing-how is rather than as a mistake about the meaning of know-how.

to offer a challenge for intellectualism, but in fact threatens *any* account of know-how as a psychological state.

Ryle alleges that intellectualists face a regress.[27] He assumes, first, for reductio, that knowing-how is a species of knowing-that. Second, that all intelligent activities are exercises of know-how.[28] Therefore all intelligent activities are exercises of knowledge-that. Third, Ryle assumes that all exercises of knowledge-that involve entertaining or contemplating certain propositions, for example, propositions describing how a given activity can or ought to be performed. Now "contemplating a proposition" connotes a conscious activity, and if it is understood in this way, Ryle's third assumption is implausible. It is better understood not as a conscious act but, rather, as *the activation of a belief.* Or, in other words, a dispositional belief is made occurrent, where an occurrent belief is not necessarily conscious.

Now activating a belief is itself an intelligent activity, thus an exercise of know-how, and so *its* performance requires that one activate *another* belief, e.g., a belief describing how one's activation of the original belief ought to proceed. But of course this too is an intelligent action, which entails that one must activate another belief in order to carry out this action, and another, and so on, *ad infinitum.* Thus, if intellectualism is true, then one must activate an infinite number of beliefs in order to perform a single intelligent action, and this is obviously absurd.

---

27 The argument I present here may not match exactly the one that Ryle offers in chapter two of *The Concept of Mind.* Think of it as a Rylean argument, adapted to fit the current debate on know-how.

28 In conversation, Bruce Hunter has questioned whether it makes sense to say that one exercises knowledge. Only dispositions, and not mental states, can be exercised. To say that a mental state is exercised is to make a category mistake. I'm not convinced by the objection: I think that some dispositions *are* mental states, and so both can be exercised. In any case none of my points turn on this. If "exercising a mental state" sounds odd, think instead of "activating or using a mental state."

Ryle's regress argument poses a challenge not just for the intellectualist, but for any philosopher who holds that knowing-how is a mental state. To exercise one's know-how one must activate *whatever* mental state realizes this know-how. And if this activation is itself an intelligent activity, itself an exercise of know-how, then it requires the activation of further mental states. No matter what kind of mental state know-how is, then, a regress develops and intelligent activity fails to get off the ground.

Stanley and Williamson's official statement of their intellectualist position on know-how is prefaced by a response to Ryle's argument. The two attack the premise that activating a mental state is necessarily an exercise of know-how (2001: 414-416). I think Stanley and Williamson are right to reject this premise though, as I'll explain, their reasons for rejecting it are unpersuasive.

First, Stanley and Williamson point out that it is possible for one to know how to X only if X is an intentional action.[29] For example, we digest our food but we do not know how to digest our food. Second, they say that activating a belief—or any mental state, presumably—is often a non-intentional action. Thus activating a mental state is not necessarily an exercise of know-how, and therefore does not require the activation of further mental states. But what do Stanley and Williamson offer in the way of an argument for the claim that activating a mental state is not an intentional action? Citing Ginet (1975) they say that exercising knowledge in action clearly does not require any awareness of the relevant mental state. But the claim that intentional actions must be consciously performed rings implausibly. If my arguments in chapter

29 I have implicitly endorsed a stronger view: one knows how to X if and only if X is an intelligent (and not merely intentional) action.

three are sound, intentional actions need not be conscious—this is a corollary of anti-voluntarism about knowing-how. Indeed, Stanley and Williamson's claim is implausible even without assuming anti-voluntarism. Remember that as a result of habit many intelligent actions are performed unconsciously. One may expertly drive car without thinking what one is doing.[30] So then Ginet's point does not help counter Ryle's argument.[31] Furthermore, Ryle argues that activating a mental state (in his words "contemplating a proposition") *is* intentional. It is rule-governed, the sort of thing that we can do well or poorly: "[i]ntelligently reflecting how to act is, among other things, considering what is pertinent and disregarding what is inappropriate" (1949: 61). So activating a mental state *is*, it would seem, an exercise of know-how. Stanley and Williamson are therefore unable to avoid Ryle's regress.

A proper response to Ryle's argument does indeed involve denying his claim that activating a mental state is necessarily an exercise of know-how. In general the activation of some instance of, say, physical know-how is guided by another instance of, this time (typically) unconscious know-how about the activation of one's physical know-how. The second mental state may serve, as Ryle suggests, to ensure that the relevant physical know-how, and not some other mental state, is activated. But the activation of this *second* mental state is in general not itself an exercise of know-how. Instead of issuing *from* us it is something that happens *to* us. Here's what happens. Some external or internal stimuli induces activity in our cognitive system, eliciting a

30 Stanley and Williamson seem to commit themselves to *de facto* voluntarism, a position I dismissed in chapter three as extremely implausible.

31 Stanley and Williamson's error may be in continuing to think of the activity as one in which we contemplate a proposition, a conscious act. Ryle, incidentally, did not appear to think of contemplating a proposition in this way: "the required consideration is often very swift and may go quite unmarked by the agent" (1949: 29).

sequence of mental events that culminates with different instances of know-how being activated. Once this happens, *we* begin to act; the events become the execution of intelligent actions. So Stanley and Williamson are right that activating a mental state is not necessarily an exercise of know-how, not, however, because it is unconscious or non-voluntary, but because the source of the activation lies outside the agent. Even if requiring further development, this is at least a plausible way of addressing Ryle's regress argument.

4.3 Two Directions of Fit

The mental state that normatively grounds the ability to engage in intelligent action is a hybrid state, having both directions of fit. It's content, furthermore, need not be propositionally encoded. I'll spend most of the remaining space in this chapter discussing this view on know-how. Having argued for it already (in sections 2.2 and 2.3 of chapter two), I'll take the liberty of assuming that it already has some plausibility: in the section I justify the *possibility* of this sort of mental state and explain exactly what it is for a mental state to have the two directions of fit I have in mind.

According to Anscombe's definition, remember, an intentional item is representative if it should fit the world, directive if the world should fit it. It is not only possible but commonplace for a mental state or a linguistic token to have both directions of fit. Ruth Millikan offers several examples of intentional items with this feature (Millikan 1996, see also Searle 1979: 18-20 and Jacobsen 1997). Before we proceed to examples of the phenomenon we should note that there is no *prima facie*

incoherence in holding that an intentional item has both directions of fit. Given a functionalist account of attitude or illocutionary force, for example, it may be that one and the same state plays two (or more) functional roles. And this may be true whether the content of the two attitudes is the same or different. Let us first discuss examples of mental states with both directions of fit, only then moving on to examples of sentences with the same feature.

Millikan suggests that *intentions* are both representative and directive. If someone intends P they both believe P will happen and desire P. Here the content of both component attitudes is identical.[32] Millikan also speculates that representations of social norms have a dual functional role. By representing features of the world and motivating action they are able to facilitate the coordination of social behaviours. Richmond Campbell has recently developed an account of the content of moral judgments along these lines (Campbell, forthcoming; see also Campbell and Woodrow 2003). According to Campbell, when someone judges that a state of affairs (for example) is morally good they (typically) both believe that the state of affairs is good and desire that it obtain. In this case, the content of the two component attitudes is different. A final and more tendentious example of hybrid mental states emerges from Millikan's examination of the empirical literature on perceptual representations. On a Gibsonian view, perceptual states simultaneously represent the structure of one's environment and indicate possible ways of interacting with it.

One of the topics of Millikan's 1996 paper is on *primitive* representations that have both directions of fit. She believes that these representations are necessary precursors to states that have only one direction of fit. We need not engage with this

32 More accurately, one desires P and has a *predictive* doxastic attitude toward P.

provocative and more controversial thesis, but it will be useful to consider Millikan's examples of primitive hybrid representations. Mechanisms of imitation in infants may employ hybrid representations, both representing what others are doing and directing the child's own behaviour. Many kinds of representations in animals are similar. The food call of a hen to its brood tells its chicks that there is food and directs them to approach their mother. The "waggle-dance" in bees represents the location of nectar and directs other bees to this location.

Natural language perhaps provides the clearest examples of intentional items with both directions of fit. We may, first of all, omit discussion of sentences that express mental states already mentioned, i.e., statements of what one intends and explicitly moral claims. Millikan's examples of other sorts of sentences include directives given in a declarative form 'We don't eat peas with our fingers' and 'You will report to the CO at 6 a.m. sharp'. These sentences both describe how things are done and elicit behaviour in others by, Millikan says, imparting an intention to a hearer. Millikan also suggests that performative utterances have both directions of fit. If someone says 'The meeting is now adjourned' or 'This ship shall be called *The Queen Elizabeth*', they are both describing how things are (e.g., that the meeting is over) and directing others to act so as make this the case.

Now, I hope, we have a sense of what it is for an intentional item to have both directions of fit, and of how common this phenomenon is. We have also observed that there are many varieties of hybrid intentional items. While all of them are generically representative, some are directive in that they are or express desires, or provide guidance for action, or impart an intention to one's audience, etc. What sort

of hybrid state is knowing-how? That is, in what way is it directive? In what follows, I try to clarify the distinction between directions of fit, and describe exactly what directive state knowing-how manifests.

In chapter two I used Anscombe's original, *normative* characterization to draw the distinction between the two directions of fit. A mental state is representative if it should fit the world, directive if the world should fit it. Other philosophers have drawn the distinction in different ways. Michael Smith, for example, tries to draw the distinction in naturalistic terms (1987: 54). As he has it, a mental state with the content P is representative if it tends to go out of existence in the presence of a perception that not-P, directive if it tends to endure in this condition.

I am not interested in debating the relative merits of Anscombe and Smith's definitions. I wish to point out instead that both are potentially misleading. Both suggest—or at least they may be interpreted as suggesting—that direction of fit distinguishes cognitive and motivational states, beliefs and desires. Representative states are indeed cognitive; if propositionally encoded they are those and only those states that are truth-apt. Directive states include desires, hopes, wishes and the like, but (and this is why Anscombe and Smith's definitions are potentially misleading) *they are not limited to motivational states*.

Know-how is a directive state but it is *not* a motivational state (like some of the examples above). The function of know-how is not to produce behaviour but *to guide it*. A good way of understanding what sort of directive state know-how is to take as an analogy a sentence that has the same directions of fit. Know-how is analogous to informationally-loaded imperatives of the following sort: 'Put one foot

in front of the other landing first with the heel and then with the toe'; 'Open the drawer by lifting up the safety latch'; 'Start the program by double-clicking on its icon'. All of these sentences are directive, but they do not express motivational states.

Although the sentences are not truth-apt, there is a truth-apt assertion buried within these sentences. In the second sentence, the assertion is 'It is possible to open the drawer by lifting up the safety latch.' This is why it happens that we acquire know-how by learning some important fact. Our cognitive systems are adapted to quickly and easily, although not unfailingly, embed representations of facts within directive mental states.

Now that I've more clearly described my hybrid view of know-how, I can turn to making explicit the considerations that count in favour of that view and, as well, to the second half of my positive account, namely, know-how's non-propositional structure.

4.4 Action Guidance and Connectionism

As we've seen, intentional items manifest one or both directions of fit. A sentence's direction of fit is determined by its illocutionary force, by its function in typical conversations. It is often straightforward what direction of fit a mental state has if it is expressed by a common sentence-type. Beliefs are expressed by assertions, and so, like assertions, have a representative direction of fit. What determines a mental state's direction of fit is it's role in action and thought, though there is, of course, functional isomorphy between corresponding mental states and sentences in this

respect.

Now for a mental state for which there are not obvious sentence types that express that state—as with know-how—language use will fail to provide any evidence about that mental state's direction of fit. At the end of the last section, I identified a sentence type that has the same specific direction of fit as know-how, what I called informationally-loaded (IL) imperatives. The example was intended to be purely illustrative. While it is clear that IL imperatives have both directions of fit, and are directive in the same way I think knowing-how is, that fact provides no evidence in favour of the hybrid view of know-how, since, absent some argument, it is controversial whether IL imperatives are directionally analogous to know-how.

In order to mount an argument in favour of the hybrid view, what is needed is a psychological characterization of the directive direction of fit. The representative direction is not the problem: it is plain enough what this direction amounts to, and both myself and my opponents accept that know-how is representative (it is covertly representative, on my view). What is it for a mental state to be directive, if it doesn't amount to just being motivational? A good answer, I believe, runs along Millikan-esque, functionalist lines, and our discussion has already touched on it: a mental state is directive if its immediate function is to produce or guide behaviour. Beliefs also guide behaviour, but that is not their *immediate* function (more on this below). Now in cases where the mental state is neither informationally-loaded nor non-motivational, it will not be able to produce behaviour on its own. Nevertheless it is directive in that it directly guides behaviour when, along with other mental states, it conspires to bring about intelligent action. A simple motivational state has the

function of changing the world so as to fit it. An action-guiding state like knowing-how has the function of making the world fit it in very specific ways, provided the subject is motivated to act. In short, a motivational state *pushes*, while a non-motivational directive state *guides*.

In chapter two I argued that knowing-how is directive in Anscombe's sense. Not only should it fit the world in some way, but the world should also fit it. This normative property of directive mental states manifests itself in two different ways, depending on whether or not the state is motivational. The world should fit motivational states in that the intentionality of these states directs the agent who possesses them to act upon the world. Inaction involves a directive failure. With non-motivational states like knowing-how, on the other hand, inaction does not necessarily involve a directive failure. The world should fit non-motivational directive states in that, given some motivation to act, the specific way the agent acts should conform to the directive state. Rather than inaction it is performance errors that constitute directive failures for non-motivational directive states. For example, suppose one knows how to ride a bike but instead of turning the handlebars to the right when one is leaning to the right, one turns the handlebars to the left and falls down. In that case the world should match one's know how but it doesn't.

Instead of motivating action, states like knowing-how are supposed to guide it. But what about beliefs? Don't they guide action too? Why think then that know-how is any different from belief? The answer is that belief and know-how are *normatively* as well as *functionally* distinct. In both respects there is a gap between belief and action, one that does not obtain between know-how and action. First, it is not

necessary that the world should change to fit beliefs. Performance that is incongruent with a belief does not necessarily involve a directive failure (in Anscombe's normative sense). That one believes X can or ought to be performed in a certain way does not commit oneself to performing X in that way.

Second, beliefs do not directly guide behaviour: that is not their immediate function. Beliefs may produce directive states (e.g., intentions) which serve to guide action, but they are not directly tied to action. For example, in some cases beliefs can not guide action at all. In cases like this someone does not simply lack ability, they also lack know-how. An individual may have correct beliefs about to perform an activity, even to the extent of being able to verbally articulate how an it is done, while lacking knowledge of how to perform the activity themselves. For example, take a hockey fan who has never in his life put on a pair of skates. He believes correctly that a forward should protect the puck with his skates, that a defenceman should always "take the body", that a goalie should always protect the "five-hole" on a breakaway, etc. But he does not know how to do these things. Even if he suddenly acquired the ability to skate, he would not—at least not immediately—possess the knowledge of how to protect the puck with his skates. And so on, *mutatis mutandis*, for the other examples. Having this sort of know-how requires does not consist in mere representative accuracy; it requires a control over one's body that the fellow above lacks.

Sometimes we describe a person's knowledge of some practical proposition—e.g., 'this is how you X'—as her knowing how to X. But this is not the sort of "know-how" that is of interest. We are interested in another sort of phenomenon. The old

saying "those who can, do and those who can't, teach" is on the right track. More accurately, those who know how, do (often), and those who don't know how (in the relevant sense), can only teach. The latter phrase seems to get things right.

One might object to a hybrid account of know-how on the grounds that a mental state with two directions of fit is not two states at all. It is one state with the one direction of fit and one state with the other. Now I am not endorsing the view that know-how's representative and directive properties are inseparable. This would be analogous to the metaethical view of moral judgments as "besires", as consisting of a cognitive and non-cognitive state that do not retain their moral content in isolation from one another (Altham 1987; Little 1997). Rather, my view is that know-how is a unified cognitive phenomenon in that its representative and directional properties mutually support one another when instantiated. This is in line with Richard Boyd's non-essentialist account of natural kinds (Boyd 1988). The objection, furthermore, assumes certain *a priori* conditions on what it is to count as a mental state. As is illustrated by Millikan's contention that hybrid states are necessary precursors to states having one direction of fit, the question of whether one mental state can be both representative and directive is an empirical one.

The reason, in brief, why know-how is not simply representative is that it is immediately action-guiding in a way that representative states can not be. Knowing-how is thus a hybrid mental state. But it is also a mental state that is often not propositionally structured. In what is left of this essay, I describe a connectionist account of know that supports the latter claim, coheres with the hybrid view, and helps to explain the existence of non-voluntary know-how.

Connectionism is a theory of mind that departs in radical ways from more traditional theories (see e.g., Churchland 1995; Clark 2001: ch.4). In traditional computationalism, the mind consists of discrete and syntactically structured representations. Beliefs and other propositional attitudes are composed of units that can be combined in various ways to encode information in attitudes having different propositional contents. In connectionism, representations are instead distributed over a neural network; different representations are related to the extent that they share some overlap across the entire system. Though I will have more to say below about what neural networks are, it is not my business here to discuss the fine details of connectionist models, or more sophisticated versions of those models (e.g., those that are temporally-sensitive). Instead my focus will be on how connectionism can be applied to the understanding of know-how.

In chapter two I argued that the performance of certain intelligent activities is not explained by the activation of doxastic states. More basic psychological mechanisms are responsible. For example, the execution of fine motor control—such as keeping balance on a bike—is brought about by the cerebellum rather than by propositional attitudes in the neocortex. Connectionist models of the mind offer a better way of understanding this sort of neural activity. As we'll see though, it is not just so-called "primitive" activities that are amenable to connectionist explanations. The very sophistication of certain activities calls out for guidance by a more nuanced neural network.

Against an account of knowledge-how as realized in beliefs, Churchland argues that much of our know-how is far too complex to be captured in a set of

general rules (Churchland 1995: ch. 6; see also his 1990 and 1998). Moral know-how, for example, involves being sensitive to countless features of one's social environment in such a way as to produce a subtle variety of behaviour. To know how to, say, respond to others who are grieving does not involve merely following the rule 'if someone is sad, comfort them.' Rather, the type of response that is called for is influenced by numerous different factors, such as who is grieving, what stage of grief they are in, what relationship you bear to them, and so on.

Churchland argues that models of the brain in cognitive neuroscience as well as connectionist artificial intelligence research are better able to accommodate the complexity of our know-how. In brief, to have a piece of know-how is to have one or more "prototypes", neurophysiological states that are activated by types of states of affairs and appropriately connected to various motor commands. A prototype is a pattern of neural activation (or, better, a pattern of *potential* neural activation) that is sensitive to indefinite many dimensions along which a state of affairs can be measured. A number of input nodes are stimulated and activate one or more sets of intermediate nodes, which are themselves connected to several output nodes. Which outputs are activated by which inputs is determined by the weights of the various connections in the system; a period of training and/or innate developmental processes set the weights to appropriate levels. Prototypes are realized in these stored weights which serve to classify input signals into various types and produce as output certain motor commands. It is in this sub-system of a neural network, Churchland argues, that knowledge-how is realized.

Prototypes can guide many different sorts of actions, in that the input they

receive originates not only in physical properties of the environment; they may also be tuned to internal physiological states as well as mental states. Sophisticated prototypes, for example, take account of muscle tension, balance sensations, etc., on the one hand, as well as background information, interpretive frameworks, etc., on the other. Thus the model Churchland provides has the potential to explain a wide variety of knowledge-how—albeit not in very familiar terms.

Now unlike, say, beliefs prototypes are understood, at least by some, as not always encoding information propositionally.[33] This claim is the subject of intense controversy, but I think it has enough able defenders such that I can rely on it. Also unlike beliefs, *prototypes are plausibly understood as having both directions of fit.* They are poised between stimulation and action in such a way as to be in a position to deploy information in the direct guidance of behaviour.

A final important feature of prototypes is their role in non-voluntary action. As was discussed in chapter three, non-voluntary know-how commonly manifests itself as unconscious sub-tasks in voluntary activity. We voluntarily initiate an action and then many intelligent sub-tasks are executed automatically. For example, we resolve to approach a social situation in a certain way, and many subtle, socially significant actions are automatically carried out. This is the sort of thing that neural networks are well-suited for. They are able to be tuned to recognizing and responding to countless details of which we are consciously unaware.

33 Although see Clark (2000) for discussion of how language can enforce syntactic structure onto neural networks.

We've come to the end of the essay. The upshot is an account of knowing-how on which it is distinct from knowing-that. Schematically, knowing-how is a mental state that is a responsibility-grounding condition on the ability to act intelligently. This is a view that many can accept without also endorsing the more specific account of know-how I've developed. On that account, know-how is a mental state that is action-guiding and covertly representational; that is, it functions to both represent the world as it is and direct action in very specific ways. Knowing-how often plays this functional role without encoding information propositionally.

Chapter three was somewhat of a departure from the main topic of this essay, namely, what knowing-how is. The topic there was what range of activities it is possible to know how to engage in; as I've said the answer to this questions is important if we are to understand what phenomenon we are studying. I argued that we know how to engage in non-voluntary activities, often without any awareness that we are doing so. The view is controversial and arguably doesn't bear any necessary connections with views presented elsewhere in the rest of the essay. It also has some even more controversial implications, as I noted, driving a wedge between voluntariness, on the one hand, and agency and intentional behaviour, on the other.

One of the more interesting features of know-how, explored too in chapter three, is its theoretical connection with other philosophical subjects. I won't rehearse all of them here, but one seems worthy of further attention in future work. In chapter three I chose to frame the discussion of voluntarism in terms familiar to normative theory. But this is not the only connection to normative theory. The account of know-how I offer in the final chapter could benefit from a more careful comparison

with a very large debate in metaethics. Metaethicists argue over what direction of fit moral judgments have. Cognitivists think that they are representative, expressivists that they are directive. Along with others, Campbell, remember, argues that moral judgments have both directions of fit. A theory of knowing-how that takes this distinction seriously (such as mine) might benefit from looking more closely at the cognitivism/non-cognitivism debate, especially at arguments that purport to count in favour or against either side in that debate. Conversely, metaethics may have something to gain from knowing-how theorists. For example, they might benefit from a more expanded taxonomy of directive states, of the sort I've begun to explore in this essay.

At the beginning of the essay I noted that procedural knowledge is a relative mystery compared to it's theoretical cousin. I hope to have demystified the phenomenon somewhat. In particular, I hope to have provided a tractable way of investigating specific kinds of know-how, to be used in concert with approaches in, for example, the philosophy of mind and virtue epistemology. Given our theoretical goal of understanding how it is we manage to achieve a cognitive fit with the world, I believe that the project pursued here can be just as illuminating as investigations in mainstream epistemology.

Bibliography

Altham, J. E. J. "The Legacy of Emotivism" in G. MacDonald and C. Wright, eds., *Fact, Science, and Morality* (Oxford: Blackwell, 1987).

Bengson, John and Marc Moffett. "Know-how and Concept-Possession" *Philosophical Studies*, forthcoming.

Bengson, John, Marc Moffett and Jennifer Cole Wright. "The Folk are Intellectualists" Unpublished Manuscript, 2006.

Blackburn, Simon. "Securing the Nots: Moral Epistemology for the Quasi-Realist" in Walter Sinnott-Armstrong and Mark Timmons, Eds., *Moral Knowledge?* (Oxford: Oxford University Press, 1996), 82-200.

Boyd, Richard. "How to be a Moral Realist" in *Essays in Moral Realism*, ed. Geoffrey Sayre-McCord (Cornell University Press, 1988).

Brett, Nathan. "Knowing How, What and That" *Canadian Journal of Philosophy*, 1974, 4, 2: 293-300.

Brown, D. G. "Knowing How and Knowing That, What" in Oscar P. Wood and George Pitcher, Eds., *Ryle: A Collection of Critical Essays* (Doubleday 1970).

Brown, D. G. "Reply to Brett" *Canadian Journal of Philosophy*, 1974, 4, 2: 301-303.

Campbell, Richmond. "What is Moral Judgment", *Journal of Philosophy*, forthcoming.

Campbell, Richmond and Woodrow, Jenna. "Why Moore's Open Question is Open: The Evolution of Moral Supervenience" *The Journal of Value Inquiry* 2003, 37:

353-372.

Chomsky, Noam. *Rules and Representations* (New York: Columbia University Press, 1980).

Churchland, Paul. "Moral Facts and Moral Knowledge" in *A Neurocomputational Perspective* (Cambridge: MIT Press, 1990).

Churchland, Paul. *The Engine of Reason, the Seat of the Soul: A Philosophical Journey into the Brain* (Cambridge: MIT Press, 1995).

Churchland, Paul. "Toward a Cognitive Neurobiology of the Moral Virtues" *Topoi*, 1998, 17: 83-96.

Clark, Andy. "Word and Action: Reconciling Rules and Know-how in Moral Cognition" in Richmond Campbell and Bruce Hunter, Eds., *Moral Epistemology Naturalized* (Calgary, AB: University of Calgary Press, 2000), 267-290.

Clark, Andy. *Mindware: An Introduction to the Philosophy of Cognitive Science* (New York: Oxford University Press, 2001).

Craig, Edward. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis* (Oxford University Press, 1999).

Dalmiya, Vrinda and Alcoff, Linda. "Are 'Old Wives' Tales' Justified?" in Elizabeth Potter and Linda Alcoff, Eds., *Feminist Epistemologies* (New York: Routledge, 1993).

Davidson, Donald. "Actions, Reasons and Causes" in his *Essays on Actions and Events* (Oxford: Clarendon, 1980).

Dennett, Daniel C. *Freedom Evolves* (New York, NY: Viking Press, 2003).

Dreyfus, Hubert and Dreyfus, Stuart. "What is Morality? A Phenomenological Account of Moral Expertise" in David Rasmussen, Ed., *Universalism vs Communitarianism* (MIT Press, 1990).

Dummett, Michael. *The Logical Basis of Metaphysics* (London: Duckworth, 1991).

Fodor, Jerry. "The Appeal to Tacit Knowledge in Psychological Explanations" in his *Representations: Philosophical Essays on the Foundations of Cognitive Science* (Cambridge, MA: MIT Press, 1981), 63-78.

Frankfurt, Harry G. "Freedom of the Will and Concept of a Person" *Journal of Philosophy*, 1971, 68/1: 5-20.

Gladwell, Malcolm. *Blink* (New York: Little, Brown and Company, 2005).

Gladwell, Malcolm. "What the Dog Saw" *New Yorker* 2006, 82, 14: 48-57.

Hartland-Swann, John. "The Logical Status of 'Knowing-That'", *Analysis*, 1956, 16, 5: 111-115.

Hawley, Katherine. "Success and Knowledge-How" *American Philosophical Quarterly*, 2003, 40/1: 19-31.

Hintikka, Jaako. "Different Constructions in Terms of the Basic Epistemological Verbs" in his *The Intentions of Intentionality and Other New Models for Modalities* (Boston: Reidel Press, 1975), 1-25.

Jacobsen, Rockney. Semantic Character and Expressive Content" *Philosophical Papers*, 1997, 26/2: 129-146.

Katzoff, Charlotte. "Knowing How" *Southern Journal of Philosophy*, 1984, 22: 61-69.

Katzoff, Charlotte. "Epistemic Virtue and Epistemic Responsibility" *Dialectica*,

2001, 55, 2: 105-118.

Kornblith, Hilary. "Justified Belief and Epistemically Responsible Action" *The Philosophical Review*, 1983, 82, 1: 33-48.

Kornblith, Hilary. *Knowledge and Its Place in Nature* (New York: Oxford University Press, 2002).

Lewis, David. "What Experience Teaches" *Proceedings of the Russellian Society*, 1988, 13: 29-57.

Little, Margaret Olivia. "Virtue as Knowledge: Objections from the Philosophy of Mind" Nous, 1997, 31: 59-79.

Mellor, H. "Nothing Like Experience" *Proceedings of the Aristotelian Society*, 1993, 90: 1-16.

Millikan, Ruth Garrett. "Pushmi-pullyu Representations" in Larry May, Marilyn Friedman and Andy Clark, Eds., *Minds and Morals* (MIT Press, 1996).

Montmarquet, James. "An Internalist Conception of Epistemic Virtue" in Guy Axtell, Ed., *Knowledge, Belief and Character: Readings in Virtue Epistemology* (Lanham, Md.: Rowman and Littlefield, 2000), 135-147.

Nemirow, Lawrence. "Physicalism and the Cognitive Role of Acquaintance" in William G. Lycan, Ed., *Mind and Cognition* (Oxford: Blackwell, 1990).

Noe, Alva. "Against Intellectualism" *Analysis*, 2005, 65, 4: 278-290.

Owens, David. *Reason without Freedom: The Problem of Epistemic Normativity* (New York, NY: Routledge, 2000).

Plantinga, Alvin. *Warrant and Proper Function* (New York: Oxford University Press,

1993).

Polanyi, Michael. *Personal Knowledge: Towards a Post-Critical Philosophy* (University of Chicago Press, 1962).

Putnam, Hilary. Introduction to Andrew Pessin and Sanford Goldberg, Eds., *The Twin Earth Chronicles* (Armonk, NY: Sharpe, 1996), xv-xxii.

Ryle, Gilbert. "Knowing How and Knowing That" *Proceedings of the Aristotelian Society*, 1946, 46. Reprinted in *Collected Essays vol. II* (New York: Barnes and Noble, 1946).

Ryle, Gilbert. *The Concept of Mind* (Chicago: The University of Chicago Press, 1949.)

Searle, John R. "A Taxonomy of Illocutionary Acts" *Expression and Meaning* (New York: Cambridge University Press, 1979).

Smith, Michael. "The Humean Theory of Motivation" *Mind*, 96: 119-127.

Snowdon, Paul. "Knowing How and Knowing That: A Distinction Reconsidered" *Proceedings of the Aristotelian Society*, 2003, 104: 1-29.

Stanley, Jason and Williamson, Timothy. "Knowing How" *The Journal of Philosophy*, 2001, 98: 411-444.

Steel, Thomas J. "A Puzzle about Knowing How" *Philosophical Studies*, 1974, 25: 43-50.

Tardif, Twila, Henry Wellman, David Liu, Kitty Yau Fong Fung and Fuxi Fang. "Preschoolers' Understanding of Knowing-That and Knowing-How in the United States and Hong Kong", *Developmental Psychology*, 2005, 41, 3: pp. 562-573.

Watson, Gary. "Free Agency" *Journal of Philosophy* 1975, 72/8: 205-220.

Williamson, Timothy. *Knowledge and its Limits* (New York: Oxford University Press 2002).

Wolf, Susan. "Freedom within Reason" (New York: Oxford, 1990).

Young, Garry. "Bodily Knowing: Re-thinking our Understanding of Procedural Knowledge" *Philosophical Explorations*, 2004, 7/1: 37-54.

Zagzebski, Linda. *Virtues of the Mind* (Cambridge: Cambridge University Press, 1996).

Zagzebski, Linda. "Must Knowers be Agents?" in Abrol Fariweather and Linda Zagzebski, Eds., *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility* (New York, NY: Oxford University Press, 2001), 142-157.