Practical Knowledge and Linguistic Competence

Natalia Waights Hickman

Brasenose College

Student number: 2765251

Thesis submitted for the DPhil in Philosophy

Word count : 56,134

Practical Knowledge and Linguistic Competence

Abstract

This thesis is comprised of four long independent papers, animated by an interest in certain parallels between the debates over practical "knowledge how", on the one hand, and linguistic competence, on the other. In particular, the *implicitness* of the putative content of both forms of knowledge, and their distinctive *productivity*—the fact that possession of the knowledge equips one to meet certain practical success conditions in the face of indefinite novelty—are two themes that loom large in the dissertation. A third is the connections, relevant to characterising both know-how and linguistic competence, between intentional action, guidance by reasons, and attributions of propositional knowledge.

The first two papers develop and defend a broadly "intellectualist" account of knowledge how to ϕ, according to which this knowledge is propositional in *content*. These aim, *inter alia,* (i) to motivate intellectualism in terms that should be largely congenial even to its opponents; (ii) to incorporate within intellectualism the anti-intellectualists' insights concerning distinctive conditions on *knowing* practically; (iii) to demonstrate that the proper target of anti-intellectualist arguments concerned with novelty is not intellectualism but a naïve representationalist view of intentionality; (iv) in this connection, to highlight parallels between a strand of Ryle's anti-intellectualist regress arguments and Wittgenstein's notorious rule following arguments.

The third and fourth papers defend cognitivist views of linguistic competence, according to which this is (at least partly) constituted by propositional knowledge of syntax and semantics. A pivotal idea exploited in support of cognitivism is that language use is, as Dummett (1996) maintained, distinctively and intrinsically *rational* activity. These papers aim, *inter alia*, (i) to defuse objections based on distinguishing language and linguistic competence as domains of enquiry; (ii) to identify as *rational* the form of explanation putative linguistic knowledge offers of exercises of linguistic competence (iii) to show that anti-cognitivism fails to do justice to one important rational dimension of language use (iv) to show that routine linguistic transactions rely upon rational sensitivity to semantic information, and that, following the "knowledge view of reasons" speakers' semantic competence must therefore be regarded as partly constituted by semantic *knowledge.*

Contents

Introduction

The series of papers comprising this dissertation developed from a special interest in what I think are natural parallels between the debates over knowledge-how on the one hand, and linguistic competence, on the other. Although each paper is intended as an independent and self-standing response to particular points of controversy within these debates, I take them to be mutually supportive, and they reflect a perspective on the two issues that is informed by a single set of commitments and concerns.

The debate over knowledge how, sparked by Stanley and Williamson's (2001) provocative paper, has focused almost exclusively on the question whether knowledge how is a species of propositional knowledge, as argued by the neo-*intellectualists* Stanley and Williamson, among others. A central controversy in the epistemology of language in effect presupposes a negative answer to this question, and is focused on whether linguistic competence is a species of propositional knowledge (as per *cognitivist* views) or, instead, a kind of practical know-how.

One finds a close symmetry in the dialectics of these debates, for some predictable reasons. For example, if knowledge how and linguistic competence are kinds of propositional knowledge, as intellectualists and cognitivists respectively contend, the propositional content is in both cases largely implicit or tacit. That is, a subject purported to possess this knowledge is not typically able to put that knowledge into words or bring it to consciousness, indeed she may lack the conceptual resources to do so. Suspicion of the notion of propositional knowledge that is implicit in this way is thus a central source of

resistance both to cognitivist views of linguistic competence and intellectualist views of knowledge how.

The productive character of both language mastery and many sophisticated forms of knowledge how has also been at the centre of the respective debates. All knowledge is productive, in the sense that it can potentially license indefinitely many novel inferences and motivate indefinitely many novel actions. But linguistic knowledge, and knowledge how to ϕ are both productive in a distinctive way (in my own view owing to the fact that both consist in knowledge of general *rules* and principles). First, they are productive in a *restricted* way, being, as it were, earmarked for deployment in particular practices: linguistic knowledge equips its possessor to compose and correctly interpret indefinitely many novel sentences of her language—and not to do much else, as long as the knowledge is implicit; knowledge how to ϕ equips its possessor to act correctly (in a way constitutive of successful ϕ-ing) in indefinitely many novel ϕ-ing tasks or occasions. Second, and relatedly, the productivity of linguistic competence and knowledge how to ϕ has a normative dimension: both kinds of knowledge equip one to achieve something particular or meet determinate standards in the face of indefinite novelty—typically in cases of knowing how to ϕ, to successfully engage in the activity or perform the action; in the linguistic case, to understand or communicate with other speakers of the language.

Curiously, as I point out in the second paper in this thesis, this common feature has figured as a consideration motivating opposite conclusions in the debates over knowledge how and language mastery: while the productivity of linguistic knowledge is a central motivation for cognitivism, the fact that know-how equips an agent for indefinite novelty has been hailed as ruling decisively *against* intellectualism. This is in part because anti-

intellectualists—mistakenly, as I shall argue—implicitly identify knowledge how to ϕ with general competence in ϕ-ing. They are emphatic that so much *besides* propositional knowledge is integral to an agent's competence to ϕ in the face of indefinite novelty: particular sensibilities, interests, goals, concerns and even character virtues and emotional dispositions will control the use to which the agent's propositional knowledge is put in each new circumstance (Wiggins 2012). The same must surely be said of linguistic knowledge, without returning any verdict as to whether the knowledge itself is propositional: its *use* in communication depends on numerous other attitudes, interests and sensibilities, including those investigated in Gricean pragmatics (Wilson and Sperber 2012). Understanding what others are saying by the use of their words indisputably requires more than knowledge of their language.

Ultimately, I have come to think that the roots of the parallel between the debates over know-how and linguistic understanding are deeper. Suppose that linguistic competence *is* a form of knowledge-how. Then the question whether linguistic competence is a kind of propositional knowledge, is just a special instance of the general question whether states of knowledge how are states of propositional knowledge. It is in this light that I now understand some of the controversy concerning linguistic understanding, and my ultimate aim is to provide accounts of knowledge how in general and knowledge how to speak/use a language in particular that exhibit the appropriate harmony. The view towards which I am building will subsume semantic cognitivism, the view that speakers possess and rely upon implicit semantic knowledge, within an account of knowing how that is broadly intellectualist—i.e. that sees knowledge how as propositional *in content*. What follows is concerned with motivating the position, which integrates three commitments: intellectualism about knowing how, cognitivism about linguistic understanding; and the

view that linguistic understanding is a form of knowing how. The four central papers of the thesis are concerned with the more or less independent defence of linguistic cognitivism and intellectualism. I will take the opportunity here to say a little to motivate the third commitment, and then conclude by summarising the arguments and conclusions of the four core papers.

Linguistic Understanding as Knowledge How

There are several pressures to place linguistic competence in the category of practical knowledge how, in addition to the distinctive nature of their productivity. I shall note four.

First, we make free use of the construction 'knows how' in imputations of linguistic competence, in ascribing both native competence to children ('Sally knows how to speak now') and second language mastery to adults ('John knows how to speak French'). Of course, the ability to *speak* cannot be a necessary condition of linguistic competence, any more than can the ability to read and write, since there are competent users of language who do not exhibit that competence in speech or do not exhibit it in reading and writing. But knowledge how to use the expressions of a language *somehow* is *prima facie* a condition on competence (Dummett 1996). And if the knowledge how ascribed in these cases is not just (some degree or facet of) linguistic competence itself, it is unclear what it could be—'John knows how to speak French' is plausibly just an idiomatic equivalent of 'John knows French'; and if we say that Sally is learning how to speak, what are we saying about her if not that she is on the way to mastering language?

Second, the most salient point of similarity between linguistic competence and other forms of knowledge how is that both are apparently *practical* competences which are typically acquired through immersion and participation, not formal instruction, and which are evaluated on the basis of performances which demonstrate the knowledge by *deploying* it, not by stating it.

Third, linguistic competence and other forms of knowledge how share the feature of inarticulability and inaccessibility to consciousness. An ordinary speaker deploying her linguistic competence in judgements of grammaticality or ambiguity will not typically be able say in what the knowledge she drew upon consists. She knows the judgement is the right one, but cannot say why. Similarly, one can *know* how to swim, how to sing, how to climb, how to make people laugh or how to comfort a friend without being able to *say* how to.

One possibility, of course, is that linguistic competence is a species of know-how, but that varieties of know-how are just complex abilities, with no propositional content which is expressible even in principle (e.g. Hornsby 2005). –This is a conclusion I will argue against in the sequel. The "abilities" here are apparently partly cognitive—their exercise is a demonstration of intelligence. Their exercise also exhibits an extraordinary sophistication which can be described in terms of rules, principles, methods or criteria; and it is also intentional behaviour, which calls for rationalising explanation in terms of the agent's motivating beliefs or knowledge states. Moreover, it is widely recognised that an agent's *knowledge* how to ϕ cannot be equated with her *ability* to ϕ, for example because knowledge how survives disability (witness the arthritic piano teacher who is no longer able to play

the piano herself). These factors point towards knowledge how and linguistic competence being varieties of *implicit* propositional knowledge.

Fourth, both kinds of competence seem not to require the kind of epistemic warrant that is usually held to be a condition on uncontroversial cases of explicit knowledge *that*. Ted Poston (2009) and Yuri Cath (2011), among others, have argued that unlike paradigmatic cases of propositional knowledge, states of knowledge how do not 'gettierize'. If one learns how to ϕ from an unreliable source or from lucky/unsafe testimony, one will nonetheless count as knowing how to ϕ as long as one can deploy what one has learned effectively in action. Similarly, as Dean Pettit (2002) has argued, semantic understanding does not seem vulnerable to gettier-style counter-examples. If Anne is told the meaning of an expression by someone who intended to mislead her, but the testimony is accurate; or told the meaning of an expression by someone who would on any other day have misinformed her, Anne's understanding of the expression will remain intact.

In both the cases of knowing how and semantic understanting, intuitions will vary: it isn't absolutely clear that the relevant knowledge or understanding survives the abrogation of various epistemic warrant conditions. The interpretation of the data is also controversial: invulnerability to gettier-style counterexample has been held to show that the relevant states are not ones of propositional knowledge, rather than that this is propositional knowledge of a distinctive sort. My own view is that, to the extent that the non-gettierizability verdict seems right, this reflects features of these types of knowledge noted above, which have nothing to do with being propositional or otherwise: it reflects the specific function the knowledge has for guaranteeing some standards or success conditions in action are met. Put crudely, where forms of know-how are concerned (of which I believe

linguistic understanding is a variety), the proof is in the pudding. The knowledge in question is *ascribed on the basis of the subjects' achieving what that knowledge is designed to enable one to achieve*, irrespective of *how* the subject came by this enabling knowledge. I do not myself think this impugns in any way the claim that such knowledge has propositional content, but here I want only to stress that non-gettierizability is ostensibly peculiar to cases of knowing how and linguistic understanding, and that this common feature is related to others that they share. This is then one more reason to think that linguistic competence is a species of practical knowledge how.

Summary of Papers

1. Knowing in the "executive way": Knowing How, Rules, Methods, Principles and Criteria

Opponents of intellectualism about knowledge how naturally ally themselves with Ryle, whose seminal essays on the knowing-how/knowing-that opposition are likely responsible for its becoming orthodoxy in the second half of the 20th century. This paper argues for a version of intellectualism, however, that is suggested by Ryle's *positive* discussions of the phenomenon. On this view knowledge how to ϕ consists in knowledge of the rules and/or principles, methods and procedures that characterise ϕ-ing itself; rules, methods or procedures that, if one knows *how*, one typically "knows…in the executive way of being able to apply them" (Ryle 1949, p. 38).

In elaborating this proposal I make clear its roots in Ryle's work, taking as a foil the crude "capacity thesis" widely (mis)attributed to him, according to which knowledge how to ϕ

is just the ability or capacity to ϕ; and I show how the contrasts Ryle drew between typical cases of knowledge-*how* and of knowledge-*that* do not indicate that knowledge how is non-propositional in content, but rather that know*ing* how involves a distinctively practical grasp of what is known.

I explain the affinity of the *quasi*-Rylean proposal concerning know-how with John Hyman's (1999, 2015) view of factual knowledge, and discuss important points of contrast with Stanley and Williamson's (2001) intellectualist account. Drawing partly on David Wiggins' (2012) work, I also propose an approach to understanding practical ways of knowing, which is motivated in favour of Carlotta Pavese's (2015) recent account of practical senses.

2. Novelty and Knowledge (How)

This paper examines and defuses the central philosophical objection to contemporary intellectualism, the view that knowledge how to ϕ is a species of propositional knowledge. This objection begins from the idea that knowledge-how equips an agent for indefinite novelty in the specific circumstances and challenges of ϕ-ing; it is held that no finite number of propositions could budget for all of these eventualities in advance, and thus that knowledge-how cannot reduce to knowledge of propositions. The challenge picks up one thread in Ryle's attempted *reductio* of intellectualism by appeal to vicious regress.

The blueprint for Ryle's regress arguments, I argue, is Wittgenstein's notorious discussion of rule following, which also problematizes indefinite *novelty* in the application of (knowledge of) rules. The target there, however, is not intellectualism about knowing how,

but a naïve representationalism about knowing a rule, suggesting that the anti-intellectualist misidentifies her target. Indeed, I argue that it is precisely because of this conflation of naïve representationalism with intellectualism that the traditional contrast between knowing-*how* and knowing-*that* has been accorded so much significance; and it is likewise for this reason that recent interrogation of the knowing how/that distinction has generated such heated controversy.

Against the anti-intellectualist's contentions, I maintain that uncontroversial cases of propositional knowledge are also indefinitely versatile, and that propositional knowledge alone *does* equip an agent for indefinite novelty—indeed I take this to be one of the upshots of Wittgenstein's regress argument concerning knowledge of a rule (the argument against intermediaries). In the case of knowledge how, it can be seen to consist in propositional knowledge as long as the propositional content is allowed to be general or generic, e.g. knowledge of general rules, ways and methods of ϕ-ing, which cover a wide range of (in certain respects) heterogeneous cases. For Stanley and Williamson's (2001) variety of intellectualism, this means that their "ways to ϕ" must be taken to be properties that are very coarsely individuated.

3. Devitt on Linguistic Competence

Cognitivism is the view that speakers possess implicit propositional knowledge of language and that this knowledge is constitutive of their linguistic competence. In particular, cognitivist views credit speakers with knowledge of a grammar or semantic theory, a system of rules competent to deliver an interpretation of any sentence of the speaker's language.

This paper seeks to examine and defuse a line of objection to cognitivism most recently and comprehensively developed by Michael Devitt (2006, 2011a), according to which cognitivism rests on a failure to distinguish languages and linguistic competence as topics of enquiry, and on the subsequent conflation of the theories and forms of explanation appropriate to each. In particular, Devitt accuses the cognitivist of assuming without argument a view that turns out (on his construal) to be highly implausible: the view that the rules of a language's grammar, *qua* psychologically real, must be isomorphic with "processing rules" that govern the speaker's production and interpretation of sentences of that language.

I argue that the implausibility of the cognitivist view on Devitt's construal is owed to a combination of (a) the tendentious framework Devitt builds for its assessment, (b) misrepresentation of cognitivism as a thesis about the structure of language processing, and (c) the illegitimate appeal to insect, bird and machine behaviour as models for comparison with human language use. Whatever may turn out to be the case with the behaviour of certain insects and birds, a salient and important feature of human language use is that it is intentional, and so *prima facie* subject to *rationalising* explanations. This makes problematic the hypothesis that Devitt favours—that speakers are in fact 'ignorant of language'. Knowledge of language is suited to provide the required rationalising explanations of linguistic behaviour; the ignorance hypothesis leaves language use unrationalized.

4. (Implicit) Knowledge, Reasons and Semantic Understanding

This paper aims to bring recent work in epistemology and action theory to bear on the question whether speakers have implicit knowledge of a compositional semantic theory. In particular, I discuss two proposals concerning the normative and constitutive roles of knowledge in practical rationality: Hawthorne and Stanley's (2008) 'knowledge-action principle,' and the 'knowledge view of reasons', originally advanced by Unger (1975) but developed more recently by Hornsby (2007a, 2007b) and Hyman (2011, 2015). I exploit these proposals in support of semantic cognitivism, the view that speakers have propositional knowledge of meaning and rely upon this knowledge in language production and comprehension.

Hawthorne and Stanley argue that knowledge provides a norm of practical rationality, one constraining reasons on which it is appropriate to act: *ceteris paribus*, one should act on the reason that *p* only if one knows that *p*; acting in ignorance contravenes the norm and is rationally criticisable on those grounds. I argue that language use is intentional and thus *prima facie* reason guided, thus subject to Hawthorne and Stanley's norm. The upshot is that semantic anti-cognitivism not only *fails to rationalise* language use, but has the consequence that it is systematically *irrational* along some dimension, for it systematically violates the knowledge norm.

The knowledge view of reasons is the view that knowledge is a condition on agent's being sensitive or responsive to a fact as reason, or being 'guided by it'. Accordingly, pressure to impute knowledge that *p* to an agent varies with pressure to locate *p* among the reasons to which the agent was sensitive. I argue that speakers routinely exhibit the relevant

sensitivity to semantic reasons and that this sensitivity is non-redundant to successful communication; thus speakers must be credited with propositional knowledge of meaning.

I discuss the problem of distinguishing guidance from conformity where an agent's putative reason is implicit, but argue from cases that speakers' behaviour must sometimes be described as *motivated* and not merely rationalised by semantic facts. I also discuss and reject the view first advanced by Evans (1981), that speaker's "knowledge" of language is not inferentially integrated with other beliefs and desires, and in this sense is not a full or genuine propositional attitude.

Knowing in the "Executive Way": Knowing How, Rules, Methods, Principles and Criteria.

Introduction

Stanley and Williamson (2001) among others have provided compelling arguments against equating knowledge how to ϕ and the ability to ϕ; against the idea that there is a substantive distinction between knowledge *that* such-and-such and knowledge *how to* ϕ; and in favour of the view that knowledge how to ϕ is in fact a variety of propositional knowledge—the thesis now called 'intellectualism'. Opponents who continue to champion the knowing how/that distinction naturally ally themselves with Ryle, whose seminal essays[1] on the target contrast are likely responsible for its becoming orthodoxy in the second half of the 20th century. Yet on re-reading those early essays, the substantive conflict between Ryle's *positive* contentions concerning knowing how and present-day intellectualism is difficult to make out. (The so-called 'capacity thesis', the thesis that knowledge how to ϕ is just the ability of capacity to ϕ, which is often attributed to him, is nowhere to be found in those texts.)

This paper expounds a variety of intellectualism about knowledge how that is in fact extracted from, or suggested by, some of Ryle's work; one that I think is more natural and compelling in details than the view offered by Stanley and Williamson. On this view knowledge how to ϕ is knowledge of the rules, principles, methods or procedures that

[1] Ryle 1945-6 and 1949, ch. 2

characterise ϕ-ing itself; rules, principles, methods or procedures that, if one knows *how*, one typically "knows…in the executive way of being able to apply them" (Ryle 1949, p. 38), but whose explicit formulation plays a significant role in the teaching, moderation and evaluation of ϕ-ing.

It is high time that certain misconceptions concerning Ryle's own position are corrected, but that is not my principal aim here. I will borrow from his work selectively myself, and my attempt to distil from this a general reductive thesis of any sort—let alone an intellectualist thesis—concerning knowledge how to ϕ, is doubtless contrary to the spirit of Ryle's approach. Nevertheless, I hope that what follows will make clear that Ryle's discussions of knowledge how are a good deal richer and more sophisticated than his critics have generally allowed. A further subsidiary purpose of making explicit the roots of this view in Ryle's writings, is to drive home a point made by Jason Stanley, which does not appear to have been assimilated (or else it has been received with scepticism)[2], and that is that subscribing to intellectualism about know-how is entirely independent of one's larger commitments regarding the nature of mind, and in particular, the relationship between cognition and action (Stanley 2011, pp. 13-18). A glance at the history of dissent from Ryle's knowing how/that distinction should suffice to cast doubt on the contrary view: Alan White (1982) and Adrian Moore (1997), both philosophers who were prepared to go along with Ryle in thinking of knowledge-how as a complex disposition or ability that is directly exhibited in action, nevertheless argued against him that knowledge how is knowledge of a proposition[3], one that may be called upon to address a question about

[2] See the Introduction to Bengson and Moffett (2011).

[3] In fact, Moore (1997, ch. 8) argues that only *some* knowledge how is knowledge of a proposition concerning what to do, while some is irreducibly practical (so is some knowledge *that*, on Moore's account).

what to do. Conversely, some who champion Ryle's knowing how/that distinction subscribe to Fodorian representatioanlism (e.g. Devitt 2011). Turning to more recent work in epistemology, I will argue that the profile Ryle gives for knowing-*how* coincides with John Hyman's (1999, 2015) view of propositional knowledge, according to which knowledge that *p* is the ability to do (infer, think, feel, say, etc.) things for the reason that *p*, where *p* is a fact or truth.

The first half of the paper is devoted to an exposition of the intellectualist view I prefer; the second discusses contrasts with Stanley and Williamson's position, and compares Pavese's (2015) account of practical modes with an approach that borrows from Wiggins' (2012) work on expertise. I begin with some preliminary ground clearing, among other things to make vivid the distinction between intellectualism and other theses with which it is often conflated. In §§ 2-4 I discuss Ryle's positive views, using the crude capacity thesis as a foil, and offer a version of intellectualism developed from these. In §5 I register some criticisms of Stanley and Williamson's proposal, and motivate the view extracted from Ryle as an alternative. §6 is concerned with conditions on knowing practically, or under a practical mode of presentation.

1. Contemporary Intellectualism and *The Concept of Mind*

Contemporary intellectualism is the thesis that knowledge how to ϕ is propositional or factual knowledge. It is not a thesis concerning skill, expertise or competence in ϕ-ing, but only *knowledge how* to ϕ. Importantly, it is a thesis concerning what this knowledge is knowledge *of*, i.e. concerning knowledge in the 'object' sense of cognitive attitude verbs, *viz*. that it is propositional or factual. As to what kind of state or potentiality of an agent

knowledge how is, intellectualism maintains only that knowing how is a knowledge relation to a proposition or fact (not to an act- or activity-type), and leaves it open whether there are distinctive conditions associated with *knowing* how, i.e. with that knowledge in the 'act' sense of the cognitive attitude verb.[4] And indeed, according to the best known intellectualist account (Stanley and Williamson 2001), *knowing* how often involves engaging propositional content under a "practical more of presentation", which in turn entails possessing complex dispositions to action. Likewise, on my way of thinking about intellectualism, its core contention is that features we might pre-theoretically take to distinguish knowledge-how and knowledge-that at the level of content or object, in fact reflect differences in the way that content is grasped, differences in a *way* of knowing.

Ryle insisted on the following contrasts: While knowing-*how* is typically acquired and exercised through practice, and often independent of an ability to put that knowledge into words, knowledge-*that* is typically transferred and exhibited by stating what is known, i.e. by making the object or content of knowledge explicit (e.g. Ryle 1949, p. 49). Intellectualism does not conflict with Ryle's contentions here. What intellectualists do not accept is the inference from the practical uptake and expression of knowledge how to the conclusion that it is non-propositional. Why? One reason is that we frequently ascribe propositional knowledge by combining the verb "knows" with an interrogative, as in "John knows *when* the train got in", or "Alice knows *where* to find the keys". Knowledge here is ascribed in terms of a question that the ascribee has the knowledge to answer, and

[4] Philosophers since Aristotle have distinguished act, power and object uses of cognitive attitude terms (see e.g. Aquinas (1953[1475], Q. 17, 1st reply), where the 'act' use is reserved for occurrent states. Following Alvarez (2010) I transfer the 'act' label to what Aquinas and his contemporaries regarded as the power use: to denote an enduring state or disposition. The same distinction is drawn and emphasised in philosophical logic without appeal to the language of acts or powers: what is believed, judged or stated, i.e. a proposition, is contrasted with a believing, judging or saying of that proposition—what might now be called a propositional *attitude*. See Kneale & Kneale ([1962] 1984) and Prior (1971).

it is overwhelmingly natural to think this is how ascriptions of knowledge *how* function also. But this argument for intellectualism does not address what are, *prima facie,* reasons for resisting it. To my mind, the best arguments for dispelling the appearance that inability to *say* what one knows demonstrates ignorance are due to Ryle. Crudely put, it is how we act and and not what we parrot that is the test of whether we have knowledge; and assertion is just one action-type among many. This will be discussed at length below.

Ryle's *The Concept of Mind* is an assault on dualism—principally the dualistic opposition of cognition and action—and the alternative picture of mind it presents is *quasi*-behaviourist. On this view, psychological predicates are not thought to be eliminable or reductively analysable in terms of behaviour; but they *are* thought to be properly characterised in terms of behaviour which provide defeasible "criteria" for the application of those predicates. Every aspect of mind of which the book treats provides Ryle with another vantage point from which to rail against the "Myth of the Ghost in The Machine" (Ryle 1949, ch. 1), and to motivate as well as apply his methodology of elucidating mental phenomena by registering patterns in their typical expressions or manifestations, primarily expressions and manifestations in overt behaviour. Accordingly, Ryle's chapter on knowing how and knowing that involves, on the negative side, a polemic against a pseudo-Cartesian idea of the exercise of practical intelligence as an execution of two tandem operations taking place in distinct realms (one the one hand, contemplating a proposition in the inner realm of consciousness; on the other hand, subsequently or concomitantly acting (*ibid,*pp.29, 31-32; Ryle 1945-6); and on the other hand, a repeated redirection of the reader's attention to the ways in which knowledge-how is revealed and expressed in behaviour—to the public criteria whose satisfaction leads us to credit someone with knowledge how to ϕ.

Two points about that chapter (and the earlier Aristotelian Society paper) are more immediately obvious if one bears in mind its discursive and methodological background. First, Ryle's concentration there on the behavioural expression of knowledge how to ϕ, the paradigm of which is, of course, ϕ-ing itself, is simply the natural recommendation of a *quasi*-behaviourist investigation of what it is *to know* how to ϕ, and continuous with his investigation of emotion, will and perception in the rest of the book. It does not signal a reduction of knowledge how to ϕ to the corresponding ability. Second, this focus is appropriate to Ryle's interest in the conditions under which an agent is said to possess knowledge how—as opposed to an interest in *what*, if anything, such an agent knows. He identifies these conditions on knowing-how with overt features of the agent's practice—in particular, the feature of being *rule or procedure-governed*—which, when recognised by third parties, warrant their crediting the agent with know-how. There is thus a sense in which Ryle's concerns are orthogonal to those of contemporary intellectualists, so that the positive contentions of each side are independent, and compatible. Ryle was interested in what it is to possess know-how; intellectualists in whether this knowledge is *of* a fact or proposition[5]. But as I shall argue below, Ryle's work in fact suggests a form of intellectualism. For he does take the rules or procedures by which an agent's practice is governed to be known by him, only what singles the agent out for an ascription of know-how is the *ability to apply the rules or procedures in practice*, not the ability to rehearse or communicate them linguistically.

2. Knowing How: The Traditional Account

[5] Rumfitt 2003, p.159 and 2011, fn. 2, has also tried to register, in slightly different terms, how Ryle's concerns ran orthogonal to those of his neo-intellectualist critics.

It is widely assumed and often asserted that Ryle's positive view of knowing how was that knowledge how to do something is simply the ability to do it, for example, that knowing how to ride a bicycle is the ability to ride a bicycle. In the section of their paper that deals with Ryle's account, Stanley and Williamson (2001, p. 416) write: "According to Ryle, an ascription of the form '*x* knows how to *F*' merely ascribes to *x* the ability to *F*." A similar summary of Ryle's view is given by Snowdon (2003). The Capacity Thesis, as Snowdon calls it, says that "knowing how to G does in fact consist in being able to G, in having the capacity to G. Knowing how ascriptions ascribe abilities or capacities to do the mentioned action." (Snowdon 2003, p. 2). Snowdon then presents two quotations from chapter II of *The Concept of Mind*, from which, Snowdon asserts, it can be seen that "clearly…Ryle is affirming that knowing how is an ability, that is, he is affirming CT." (*ibid*, p.3).

Contempt for the so-called Capacity Thesis is as widespread as the assumption that it originates with Ryle[6]. Counterexamples abound, highlighting a number of important respects in which ability and knowledge how differ. First, critics point out that a person is often unable to exercise their knowledge how, but this does not imply they have temporarily lost it: a manufacturer's ability to produce plastic moulds is contingent upon the supply of oil, but his knowledge how to is not; an explosives expert may be unable to defuse a bomb he cannot access, but it does not follow that he does not know how to defuse it. Second, physical impairment can deprive me of certain abilities without depriving me of knowledge how; for instance, severe arthritis would not rob me of my knowledge how to play the piano. Third, one may know how to do something and *never* have possessed the corresponding ability. Stanley and Williamson (2001, p. 420.) give the

[6] For the widest range of counterexamples, see Snowdon (2003). Others can be found in Stanley and Williamson (2001) Ginet (1975), White (1982), and Moore (1997).

example of a ski instructor who can teach others more agile than herself how to perform a complex jump, though she lacks the ability to perform it herself. Finally, we can acquire, possess, and improve our abilities to do some things without knowing how to do them, such as the ability to wiggle our ears.

The first point is evidently off target. It is true that an explosives expert does not lose his knowledge how to defuse a bomb whenever a locked door prevents him from exercising it, but neither does he lose his abilities on this account. We say he is not 'able to' or 'cannot' defuse the bomb meaning that he lacks the occasion or opportunity; this is an obviously inappropriate interpretation to give the 'able to' of the Capacity Thesis.[7] A man does not lose the ability to see (i.e. go blind) when he sleeps, or the ability to draw when he forgets his pencil. Opportunity and external impediment affect the exercise of abilities in just the same way they affect exercises of knowledge how, so any (putative) counterexamples which trade on an opportunity sense of 'able to' fail to break the equation between knowledge how and ability.

Perhaps the remaining objections to the Capacity Thesis could be partly countered in a similar way; it might be argued that they fail to address the relevant notion of ability. But such a strategy could only take us so far: taken together, I believe the counterexamples found in the literature show incontrovertibly that the equation of knowing how to ϕ with the ability to ϕ is mistaken. But they entirely fail to show that *Ryle's* account of knowing

[7] Ryle devotes several pages to a discussion of the different uses of 'can' an 'able to' in chapter V of *The Concept of Mind* (1949). For more on the 'can' of opportunity and the 'can' of ability, see also Kenny (1989, chapter V).

how is unsatisfactory. Ryle did not advance the Capacity Thesis, and his own account incorporates all the distinctions just registered between knowing how and ability.

3. Ryle on Knowing How

One point to notice about Ryle's discussions of knowing how, is that for him knowing how was not so much the *explanandum* as the *explanans*. He took the contrast between knowing how to do things and knowing *that* something is the case to be familiar and obvious to his readers, and he thought he could appeal to it in order to upset a conception of *intelligence* as purely intellectual, exercised in contemplating truths (Ryle 1945-6, p.1). His first interest, then, was not in knowing how but in what he called intelligence concepts, e.g. skill, shrewdness, acumen and wit (*ibid*, p.5 and *passim*; Ryle 1949, p.25-27). His negative claim was that these *do not* reduce to knowledge of facts or truths (*ibid*, p.30-31). One can accumulate knowledge of facts and truths by being told them or by looking them up in books. But when one infers on the basis of one fact that a second must obtain, or when one determines, because of some fact or facts one has learned, to take one course of action rather than another, or to undertake the action in a certain way and not another, then one is exercising *intelligence in the use of one's knowledge* (*ibid*, p.28). And that intelligence cannot be acquired merely by being informed of what is the case. As Ryle put it, "stupidity is not the same thing, or the same sort of thing, as ignorance." (*ibid,* p.25), and it cannot be remedied in the way ignorance is remedied. Enlarging the stock of facts on which one can draw does not make one more ingenious in the ways one draws on them. Knowledge can be accumulated but intelligence, ingenuity, skill and judiciousness can only be developed. (*ibid*, p.41).

In these contentions I believe Ryle was right[8], but they do not bear at all on contemporary intellectualism, which is a thesis about *knowing how* and not about varieties of intelligence.

Ryle's positive claim about intelligence was that it is a kind of knowledge how (*ibid*, pp. 28, 32), as against knowledge that anything is the case, and this claim is certainly mistaken. His real aim was to show that intelligence is a practical potentiality exhibited in action as much as thought, and that it is then exhibited *directly*: it is an individual's *acting or thinking* itself that shows her to be either intelligent or stupid, not her possession of information—information on which she putatively reflects before acting or thinking (*ibid*, pp.25-6). But if all this is right, it shows that intelligence does not reduce to knowledge; it does not show that it is a kind of *practical* knowledge rather than *propositional* or *factual* knowledge. Intelligence is not equivalent to knowing how to do anything, or with knowing how to do things in a particular way, other than *intelligently*[9]. It is true that a modicum of intelligence is necessary for knowing how to do anything, but it is just as necessary for knowing that something is the case. And it is true that one may show special intelligence in the rapid acquisition and ingenuous exploitation of knowledge *that*, and that this intelligence is not itself reducible to further propositional knowledge (Ryle 1945-6, p.7, quoted above). But it is not reducible to knowledge how either. One can just as well show intelligence in the rapid acquisition and ingenuous exploitation of knowledge how (e.g. how to play chess or

[8] See Weatherson (forthcoming) for a defence of Ryle's position concerning the irreducibility of intellectual *skill* to propositional knowledge.

[9] One might think that this concedes Ryle's point, for his claim was that intelligence amounts to knowing how *to do things in an intelligent manner*. Thus intelligent action is distinguished not by proceeding from contemplation of propositions but by exhibiting a certain method or *modus operandi* :"what distinguishes sensible from silly operations is not their parentage but their procedure."(Ryle 1949, p. 38). But my point is that Ryle's contention that intelligence is dispositional and manifest directly in action does not entail that intelligence is a kind of practical *knowledge*. Intelligence is no more *knowing how* to do things intelligently than elegance is knowing how to do things elegantly, dexterity is knowing how to do things dexterously or punctuality is knowing how to arrive on time.

tennis), and this intelligence surely is not a further form of knowledge how. Varieties of intelligence or intellectual skill are not reducible to knowledge how to do something, any more than they are reducible to knowledge that something is the case.

So Ryle was, I believe, misguided in his attempt to elucidate forms of intelligence by elucidating forms of knowledge how. But in their own right, when stripped of some of the behaviourist excesses[10], Ryle's positive discussions of knowledge how contain more insight than error, insight that is obscured by the myth of the crude capacity thesis. These are the passages that Snowdon (2003) quotes, which he thinks taken together show that Ryle endorsed the view that knowing how to ϕ is just the ability to ϕ.

> 'Intelligent' cannot be defined in terms of 'intellectual' or 'knowing how' in terms of 'knowing that' (Ryle 1949, p.32)

and

> When a person is described by one or other of the intelligence epithets such as 'shrewd' or 'silly', 'prudent' or 'imprudent', the description imputes to him not the knowledge or ignorance of this or that truth, but the ability, or inability, to do certain sorts of things. (*ibid*, p. 27)

Were this all that Ryle had ever said about knowing how, it might be tempting to think, as Snowdon presumably does, that the first passage equates intelligence with knowing how, while the second proffers a definition of intelligence in terms of abilities. But even if this supposition were acceptable, it wouldn't justify the imputation to Ryle of the Capacity Thesis. It isn't clear from these passages what kind of knowledge how Ryle thinks intelligence involves, and there is nothing to suggest that he thinks knowing how to do something amounts to being able to do *that thing*. And this is, of course, not all that Ryle said about knowing how. Interestingly, the seven pages devoted to a positive discussion of

[10] I have in mind such contentions as that the difference between intelligent, intentional behaviour on the one hand, and non-intentional behaviour on the other, consists in "the absence or presence of certain sorts of testable explanatory-cum-predictive assertions." (Ryle 1949, p.25).

knowing how in the Presidential Address make no mention of abilities at all. Instead the key points are:

> (a) When a person knows how to do things of a certain sort … his knowledge is actualised or exercised in what he does. …;
> (b) When a person knows how to do things of a certain sort…his performance is in some way governed by principles, rules, canons, standards or criteria…it is always possible in principle, if not in practice, to explain why he tends to succeed, that is, to give the reasons for his actions…;
> (c) We certainly can, in respect of many practices…extract principles…but when we try to express these principles we find that they cannot easily be put in the indicative mood. (Ryle 1945-6, pp.8-11).

The second point is the most salient in that earlier essay, and it is mentioned or discussed on virtually every page of the chapter on knowing how in *The Concept of Mind*, yet it has been ignored by critical and sympathetic commentators alike. Exercises of knowledge how, Ryle says, are 'observances or rules or canons, or applications of criteria' (Ryle 1949, p.42); they 'are distinguished by their procedure' (*ibid*, p. 49); they exhibit a 'method or *modus operandi*'(Ryle 1945-6, p. 3). I will say more about these insights later, but it should already be obvious that Ryle's profile for knowing how forbids its assimilation to unlearned or brute capacities, such as the ability to raise one's arm or the ability to digest food. Exercises of these abilities do not have *a modus operandi*; they are not "governed by principles, rules, canons, standards or criteria" (*ibid*, p.8).

An example of how abilities *do* come to be connected with knowledge how in Ryle's work is provided by his discussion of the boy learning to play chess:

> It should be noticed that the boy is not said to know how to play if all that he can do is to recite the rules accurately. He must be able to make the required moves.... His knowledge *how* is exercised primarily in the moves he makes... So long as he can observe the rules, we do not care if he can also formulate them. (Ryle 1949, p.41).

In order to be credited with knowledge how to play, the chess player must be able to make the right moves, avoid and veto proscribed ones, and so on. But this claim is thoroughly

plausible, and quite different from the claim that knowing how to play chess consists in the ability *to play chess*. What it suggests instead is that knowing how to play chess involves the ability *to follow the rules* of chess in practice. Demonstration of this ability is both necessary and sufficient to warrant an attribution of knowledge how to play, while reciting the rules is neither necessary nor sufficient. The same point is made in the 1945-6 essay when Ryle says, in elaboration of point (a) registered above, that a man's observation of rules or principles in exercising knowledge how "must be realised in his performance of his tasks. It need not (though it can) be also advertised in …paying some external or internal lip-service to those rules or principles." (Ryle 1945-6, p. 8-9).

Ryle comes closest to giving a straight definition of knowing how in the following passage:

> Knowing how, then, is a disposition, but not a single track disposition like a reflex or a habit. Its exercises are observances of rules or canons or the applications of criteria, but these are not tandem operations of theoretically avowing maxims and then putting them into practice. (Ryle 1949, p.42)

Knowing how is, on Ryle's account, not a single-track but a *multi-track* disposition, one of the "higher grade dispositions…the exercises of which are indefinitely heterogeneous."(*ibid*, p.44).What does this mean? It means that knowing how to ϕ is not merely exercised in ϕ-ing. It is also exercised in thinking about ϕ-ing, evaluating one's own and others' performances, criticising and instructing:

> You exercise your knowledge how to tie a clove-hitch not only in acts of tying clove-hitches and in correcting your mistakes, but also in imagining tying them correctly, in instructing pupils, in criticising the clumsy or incorrect movements and applauding the correct movements that they make, in inferring from a faulty result to the error which produced it, in predicting the outcomes of observed lapses, and so on indefinitely. (Ryle 1949, p. 55).

This is unsurprising if knowledge how is knowledge of rules, principles or criteria, which is exercised in their practical application or observation. For if one knows the rules or criteria governing a practice, one can draw on this knowledge in multifarious ways. In any

case, it should be clear that what Ryle says about knowing how to tie a clove-hitch he would say about knowing-how to perform a ski jump also. There is simply nothing in Ryle's writing to suggest he would have denied, or must be unwittingly committed to denying, that the skiing instructor Stanley and Williamson discuss knows how to perform the jump, even thought she lacks the agility to perform it herself. The point Ryle wanted to drive home rather is that the ski instructor's knowledge how is not demonstrated by her ability to recite a procedure from a text-book, but by her ability to appeal to that procedure, implicitly or explicitly, in evaluating, criticising and modifying the performance of her students. It is precisely because an injured skier can still train others that we can see her know-how survives, and thus is independent of, disability. It it still what she *does* on its basis that demonstrates her knowledge how.

It may certainly be said against Ryle that this feature of being primarily expressed in action belongs to knowing *that* as well as knowing *how*. For we do not credit parrots with knowledge of the facts that are semantically expressed by sentences they parrot, but we credit human beings with knowledge of those facts when they *act on them*[11].But Ryle's error here lies in over-intellectualising knowledge-*that*, or in failing to do justice to its practical or dispositional aspect; it is not in over-simplifying knowledge-*how*.

A note about point (c) above, i.e. the fact that when we specify the principles and rules governing many practices with respect to which one can have know-how, it is natural to express them using imperatives. Ryle suggested that this also upsets the intellectualist

[11] For example, we judge that someone holds sexist beliefs not because she often articulates them, but because she reaches judgements and decisions, and displays overt behaviour, for which sexist beliefs provide the only intelligible premise.

doctrine, but he does not say why. He suggests later on in the paper that the use of the imperative has to do with the pedagogical use of rule-formulations: "they are imperative because they are disciplinary, they are in the idiom of the mentor." (*ibid*, p.12). Ryle says nothing further about the contrast between declarative and imperative moods or what it is supposed to signal. Some philosophers have urged that knowledge of rules, which often though not always are expressed using imperatives, is not knowledge that is propositional, among other things because rules are not "truth-apt": they may obtain or not, but are not true or false. There is evidence that Ryle too thought there is a substantive rule/proposition distinction[12]. Combining this assumption with the view of knowledge how to ϕ that I extract from Ryle's work (as knowledge of rules, principles, criteria, methods of procedures that characterise ϕ-ing) would have the consequence that *some* knowledge how is non-propositional. But some uncontroversial and explicit forms of knowledge *that* would turn out to be non-propositional too, for example knowledge that the bishop in chess moves diagonally. This is to say that emphasising the use of the imperative in the expression of what is known does not support Ryle's distinction between knowing-*how* and knowing-*that*, or cohere with the terms in which he drew it. For even if there is some interesting difference between knowledge of rules and knowledge of propositions or facts (the former but not the latter sometimes being expressed by means of the imperative), it cuts across the distinction between "practical" knowledge, typically acquired through and expressed in action, and knowledge-that such-and-such is the case, which can be picked up easily by reading or hearing the fact, rule or principle asserted[13].

[12] See the second paragraph of Ryle 1945-6, p. 7.

[13] It is better aligned with the distinction between propositional and practical knowledge drawn by Adrian Moore (1997, ch. 8).

What, then, is Ryle's conception of knowledge how? I do not think I or anyone can claim, by any amount of careful exegesis, to recover a precise thesis to which Ryle was committed. If he did not proffer any slogan theses or definitions of knowledge how, it is because he did not endorse any.[14] But there is a picture that Ryle paints of knowledge how and its uses, one that makes certain features salient, and the one that receives most attention, as I have said, is the idea that exercises of knowledge how "incorporate special procedures" (1945-6, p.6), are governed by rules, principles or criteria, or exhibit a special method or *modus operandi*. Which rules, principles, methods or procedures are applied by someone with knowledge how to ϕ? –Those that characterise or define ϕ-ing itself; those whose observation makes an act or activity count or qualify as ϕ-ing. The student learning how to play chess observes the rules of the game, including those determining the powers of each piece; as she progresses she will discover and apply more complex principles of strategy; eventually she may 'invent' them. The amateur dancer follows basic choreography; as she progresses she will acquire more sophisticated techniques and internalise principles concerning when and how to respond to cues from partners or musicians. And so on.

If someone with knowledge how to ϕ, when intentionally ϕ-ing, systematically observes the principles or procedures of the practice of ϕ-ing, it is natural to suppose that their doing so requires, and proceeds from, *knowledge* of those principles or procedures. Ryle appears to endorse this idea also. For example, he agrees with his interlocutor that someone who reasons intelligently "is knowing rules of inference whenever he reasons intelligently". But

[14] Some of Ryle's champions would no doubt add that there are none to be had. I am more optimistic.

he adds that "knowing a rule is knowing how. It is realised in performances that conform to the rule, not in theoretical citations of it."[15]

Ryle emphasises the difference between acknowledging these rules in the abstract—assenting to their formulations—and observing them in practice; and in particular, the independence of the two achievements. The man who reasons intelligently,

> Observes the rules of logic without thinking about them. He does not cite Aristotle's formulae to himself or to the court. He applies in his practice what Aristotle abstracted in his theory of such practices. …. The rules that he observes have become his way of thinking, when he is taking care; they are not external rubrics with which he has to square his thoughts. (*ibid*, p. 48)

Similarly, Ryle conjectures that a chess student might lose (or never have gained) the ability to recite the rules of the game, but

> so long as he can observe the rules we do not care if he can also formulate them, it is not what he does in his head or with his tongue that shows whether or not he knows the rules in the executive way of being able to apply them. (*ibid*, p. 41).

Applying in action and reciting verbally seem by Ryle's lights to require, and in that sense reflect, different ways of knowing or grasping: knowing in theory and knowing in practice, in the executive way. Thus: "the intelligence involved in putting the prescriptions into practice is different from that involved in intellectually grasping the prescriptions. There is no contradiction…in in describing someone as bad at practicing what he is good at preaching." (*Ibid*, p.49).

[15] It is utterly unclear to me what Ryle could mean by "theoretical citations" such that a theoretical citation of a truth suffices to demonstrate genuine knowledge, but a theoretical citation of a rule does not. My impression from passages like these is that Ryle took there to be an important difference between rules and truths, and believed that there is something especially practical about the knowledge of the former but not the latter. However, excepting the one remark just quoted, that claim is never made explicitly, or defended.

Indeed, Ryle appears to suggest that typically, rules, principles and procedures are formulated only *after* they have been applied in practice (*ibid*, p.30); in one of his more eccentric but nonetheless vivid metaphors: "rules, like birds, must live before they can be stuffed." (Ryle 1945-6, p.11). None of this shows, as Ryle once claimed it did[16], that knowing how is logically prior to knowing that. But it does suggest that, where rules and methods of practice are concerned, knowing is often prior to *saying*; knowledge-in-action is prior to knowledge stated.

Whatever Ryle's own position might ultimately have been, I find much in his discussions of knowing how that is plausible and suggestive. What I have extracted from reflection on these is the following, simple thesis:

Knowing how to ϕ is knowing the rules, principles, methods or procedures that govern ϕ-ing, *in the executive way of being able to apply them*.

This is not so distant from Stanley and Williamson's claim that knowing how to ϕ is knowing that *w* is a way for one to ϕ, *under a practical mode of presentation*. And the distance is further narrowed if we incorporate Pavese's (2015) conception of practical modes of presentation (discussed further in §6), on which grasping a way to ϕ under a practical mode of presentation entails the ability to follow a rule to ϕ.

[16] Ryle 1945-6, pp.5-6: "I want to turn the tables and to prove that knowledge-how cannot be defined in terms of knowledge-that and further, that knowledge-how is a concept logically prior to the concept of knowledge-that." The latter claim is not repeated in the 1949 book chapter.

Of course, from Ryle's point of view, most of what is interesting about knowing how concerns the 'executive way' of knowing that is involved—the conditions associated with possessing such knowledge—not its content. But this is a difference in interest and emphasis between Ryle and Stanley and Williamson, not a substantive point of disagreement. Both positions see knowledge how as an attitude to a content—propositional content, if one allows that rules are a proposition-type. And both grant that knowledge how is typically expressed in, and evaluated on the basis of, action or activity, reflecting the distinctive way that the propositional content is thought about or grasped. There are, however, some important differences, which I turn to below.

4. Knowledge as ability

On Ryle's account, it is not merely that knowing how has a dispositional aspect, is associated with certain dispositions, or affords certain abilities. Rather, knowing how *just is* a multi-track disposition (Ryle 1949, p.42). This would seem to prevent its assimilation to most accounts of propositional knowledge, but not all. The rough profile Ryle gives for knowing how in fact aligns best with John Hyman's (1999, 2015) view of propositional knowledge as the ability to be guided by, or do things for reasons that are, facts or truths—a view of knowing-*that* whose inspiration Hyman credits to Ryle, among others (Hyman 2015, p.164)[17]. Hyman argues that knowledge, like any ability, must be characterised by reference to what it enables us to do, rather than by reference to mechanisms through which it is acquired or factors which would make a knowledge claim justified. But propositional knowledge is extremely versatile, which appears to stand in the way of

[17] In *The Concept of Mind*, Ryle writes that '"know" is a capacity verb, and a capacity verb of that special sort that is used for signifying that the person can bring things off, or get things right.' (Ryle 1949, p.133).

delivering a satisfactory characterisation of knowledge by reference to its exercises or manifestations in action. My knowledge that the pub is due to close at 11 pm may be expressed when I order a last round of drinks at 10.30, when I advise a friend not to join us if they cannot arrive by 10.15 pm, and when I show frustration with the barman who attempts to usher us out at five to eleven. What feature could be common to all of this behaviour, which distinguishes it as an expression of my knowledge?

Propositional knowledge is indeed versatile, so its expressions in behaviour are heterogeneous, but, Hyman insists, that is not to say those expressions lack any definite pattern or distinguishing feature. When knowledge that *p* is expressed in what we do, think or feel, the fact that *p* is among the reasons why we do, think or feel what we do (Hyman 2015, pp. 167-168). Knowledge, then, is the ability to be guided by the facts, or more precisely, to do things for reasons that are facts or truths (*ibid*, p.169-70).

The account purports to explain both how knowledge differs from true belief, and why we value it more highly: only when we know, and not when we merely believe, are we guided by the facts or by the truth[18]. But it offers no grounds for a distinction between practical and propositional knowledge, or for thinking that 'knows how' and 'knows that' are constructions typically used to ascribe irreducibly different varieties of knowledge. On the contrary, the ability to do things for reasons fits Ryle's profile for knowing how well. He writes:

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

[18] See Hyman 2006 and Hyman 2015, chs 7-8, for a defence of the account's power to explain the distinctiveness and value of knowledge as compared with true belief.

> method in his cleverness. But his observance of rules, principles, etc., must, if it is there at all, be realized in his performance of his tasks. It need not (though it can) be also advertised in an extra performance of paying some internal or external lip-service to those rules or principles. (Ryle 1945-6, p.8-9, emphasis added).

My understanding of this passage is that exercises of know-how are guided by rules, principles etc., in the usual way that actions can be guided by premises or grounds: these are the reasons for the agent's actions. However, to count as knowing how the agent must *act* on these reasons, not necessarily be able to state them. This is, *prima facie,* a description of propositional knowledge as Hyman conceives it.

It is worth noting that the intuitive contrast between knowing how and knowing that can still be drawn even if we take Hyman's account to cover both. For his view is that propositional knowledge is the ability to be guided by the facts, and that leaves open the possibility that there are distinctive ways in which the facts may guide us. The fact that oil paint is more viscous than acrylic might be (one of) my reasons for telling someone else that this is the case; my reason for encouraging a novice to work with acrylic before using oil paint; my reason for handling the pallet knife differently when I work with these two paints. We can imagine a person with only 'knowledge that' who can access and respond to that reason in the first sense but not the last; and another with 'knowledge how' who is sensitive to that reason in the last sense but not the first. Ryle's distinction between knowing how and knowing that may be seen as a distinction between ways in which facts, truths, rules and so on are apt to guide action.

I can see only two sources of resistance to the inclusion of knowing how within Hyman's account. One is that knowledge-how, on the quasi-Rylean view I advance, *sometimes* involves knowledge of rules (sometimes methods or means—more on this below), and one

may think that knowledge of rules is non-propositional. This is an issue that has been discussed above, and I do not propose to get into the question whether there is a substantive rule/proposition distinction—as I have said, if there is one, it cross-cuts the distinction between knowing how and knowing that. But more pertinently, I do not think Hyman's account can be restricted so as to exclude knowledge of rules, for he maintains that reasons-explanations of action credit the agent with knowledge of the explanans, and rules can certainly be explanatory reasons. Thus Anne's stopping at a red light is partly explained by a traffic rule; Mary's raising her hand in a seminar to ask a question by a rule of etiquette, and Sarah's castling in chess by a chess rule (Cf. Hyman 2015, p.170)[19].

The second motivation for resisting the assimilation of Ryle's knowledge-how to Hyman's view of propositional knowledge would be suspicion of the notion of an *implicit* reason. Reasons are often given in justification or explanation of our own actions, and it may be thought that an agent cannot be said to act on a reason if he cannot give it. This would in reality be an objection to the conception of knowing how that I propose, for it is at root an objection to the notion of implicit propositional *knowledge*.

Typically, it is complained that ascriptions of implicit knowledge over-intellectualise the behaviour alleged to demonstrate it, but this charge seems to me quite misdirected. For its force depends on a conception of knowledge that is highly intellectualised, for example as

[19] There may be other salient explanatory reasons for these actions, but which reasons are found explanatory depends on the question we have raised about the action, or the aspect of the action we want illuminated. Thus we may take the defining rules of chess for granted, and ask why Sarah made some particular move looking for understanding of her strategy. But someone partly ignorant of the rules could ask, say, why she had not moved her rook diagonally instead (perhaps a wonderful advantage could be gained if only such a move were licit). The reason she must give is just that this "move" is disallowed—in this case her knowledge of the defining rules, on which she certainly drew, will be brought out as a reason for action. On the question- or interest-relativity of action explanation see Achinstein 1975 and Sandis 2012.

involving a conscious mental representation which is consulted whenever the knowledge is drawn upon in thought or action. (This is, for example, the conception of knowledge-*that* which Ryle assumes in the context of his attacks on the 'intellectualist' assimilation of knowing-*how* to knowing-*that*)[20]. With this demanding conception of knowledge in the background, the suggestion that agents are acting on implicit knowledge seems to require that they be thought of as possessing *unconscious* mental representations that are *unconsciously* consulted whenever that knowledge is drawn upon in thought or action. It is the critics' intellectualist presupposition concerning explicit knowledge that is responsible for the absurdity of the implicit knowledge thesis they target. If one eschews the intellectualisation of *explicit* propositional knowledge, then there is no sense in which an *implicit* knowledge claim involves over-intellectualisation. Neither acting on *explicit* knowledge nor acting on *implicit* knowledge need be thought of as involving the consultation of a mental representation.

Hyman's account of knowledge as ability is, in this sense, as un-intellectualist as they come, and there is no suggestion that only agents who can give their reasons act on them—i.e. that there is no such thing as implicit reasons-guidance and implicit knowledge. Indeed, such a restriction would prohibit the application of the concept of knowledge to animals and pre-verbal infants, a consequence that also follows from Alan White's[21] definition of knowledge that *p* as the ability to answer the question whether *p*, which is explicitly rejected

[20] As Stanley (2011, Ch. 1) explains, Ryle assimilates the ideas that knowledge-*how* is a species of knowledge-*that*, and that exercises of know-how are steered by the contemplation of a proposition—the latter, Ryle contends, is absurd. Arguably these two ideas are natural companions within what Ryle conceived as the prevalent 'Intellectualist Legend', but there is no necessary connection between them. Ryle's claim about the intellectualist aetiology of exercises of know-how requires that exercises of knowledge-*that* be thought of as steered by the anterior contemplation of a proposition. And this is, of course, a commitment disclaimed both by Ryle in his own positive discussions of knowledge, and by many of Ryle's intellectualist critics (e.g. Stanley and Williamson 2001, Ginet 1975).

[21] White (1982) Ch. 6.

by Hyman for that very reason (Hyman 1999, p. 437). Hyman does appear to think that animals and pre-verbal children can know things, and accordingly that they can act for reasons. By Hyman's own lights at least, then, I can see no grounds for excluding knowledge-how, as Ryle describes it, from the category of propositional knowledge.

I have felt it important to flag this connection with Hyman's views in order to drive home how *much* of Ryle's own philosophy of mind is congenial to an assimilation of knowing-how to knowing-that. But ultimately the main contentions I want to develop from Ryle's work have nothing to do with the identification of knowledge with a complex ability or disposition. The quasi-Rylean claim I endorse, that knowledge-how to ϕ is knowledge *of* the rules, methods, principles or procedures of ϕ-ing, is independent of the claim that this knowledge *consists in no more than* the ability to follow these rules, methods, principles or procedures (i.e. to take them as reasons for action). Like other intellectualist views, the one I prefer does not require any particular metaphysics of propositional knowledge in particular, or of the mind generally: one can just as well maintain that knowledge how *affords* rather than *consists in* the ability to apply the known rules, methods, principles etc. in action[22].

[22] I have slight reservations about embracing Hyman's identity thesis in favour of the thesis that knowledge merely *entails* the ability to do things for reasons that are facts, for two reasons. First, as Williamson (2000, p.64) and Setiya (2013 p. 511) have objected, there seems to be an explanatory asymmetry between knowledge that *p* and the ability to ϕ for the reason that *p*, insofar as the former intuitively explains the latter but not vice versa. Second, an important motivation for Hyman's view is that it purports to explain the pattern of uses or manifestations of knowledge (2015, pp. 167-9). But it seems that knowledge might in some cases *inform*, or be taken into account in, an action, without yet being among the agent's reasons for performing it. For example, as Erasmus Mayr has pointed it out in discussion, if one ϕs *despite* the fact that *p* one must know that *p* and take *p* into account, but *p* cannot here be counted among one's reasons for ϕ-ing. These are not decisive objections—Hyman gives a detailed response to the first in his 2015, pp.181-4—but they lend some support to the idea that knowledge is *more* than, and perhaps the basis of, the ability to do things for reasons that are facts.

5. Rules, principles, criteria… *vs* ways to ϕ

The most salient difference between the view I extract from Ryle's discussion of know how, and Stanley and Williamson's view, concerns *what* someone with know-how can be said to know.

According to Stanley and Williamson (2001), an ascription of knowledge how to ϕ contains an embedded question concerning how to ϕ, to which the knowledge ascribed supplies an answer. Just as knowing *where* the Post Office is, is knowledge of an answer to a question about the Post Office's location, and knowledge *when* the next Paddington train will leave is knowledge of an answer to the question when the next Paddington train will leave, knowledge *how* to play chess is knowledge that supplies an answer (or answers) to a question concerning how to play chess. This aspect of Stanley and Williamson's proposal strikes me as independently quite plausible, although there are now well known objections to the semantic argument for it, e.g. relating to the fact that use of the interrogative particle 'how' in this kind of knowledge ascription is an adventitious feature of English, not found in other languages (Rumfitt 2003; Wiggins 2012). What I think is less plausible is Stanley and Williamson's view of what the knowledge (how) *is* that addresses questions about how to ϕ.

Their official view is that an ascription to an agent A of the form 'A knows how to ϕ' is an ascription of knowledge that *w* is a way for A to ϕ. In the basic case, an agent knows (or is ascribed knowledge of) just one such way-proposition, as I shall call them. But she may

know several or all of them ("mention all" reading)—that is, she may know all the ways *w* for her to ϕ. Thus Stanley (2011, p.183) explains:

> When we say that a skilled outfielder knows how to field a fly ball, we do not mean that he knows, of at least one way to field a fly ball, that it gives him counterfactual success in fielding fly balls. That is, we do not intend the *mention-some* reading of the embedded question "how to field a fly ball". Rather, in such a case, we mean the mention-all reading of the embedded question. What we assert when we assert of a skilled outfielder that he knows how to field fly balls is that he knows *all* of a range of relevant ways.

The ways to ϕ that an agent with knowledge how to ϕ knows are properties of token-events of ϕ-ing, picked out demonstratively. Thus the ways to field a fly ball known by the professional baseball player are properties of events of his fielding, and they can be picked out demonstratively as he fields a ball, saying 'that → is a way for him to field a fly ball' (Stanley 2011, *ibid*).

The analysis applies to all ascriptions of the form 'knows how to ϕ', not only those in which the phrase '(some) way to ϕ' provides an intuitively acceptable paraphrase of 'how to ϕ.' Thus knowledge how to play chess, knowledge how to dance flamenco, knowledge how reason correctly, and knowledge how to compose a symphony are all taken to consist in knowledge that some way *w* is a way for one to do the thing in question. Thus Mozart's knowledge how to compose a symphony—that is, the content of *what* he knows—is also a property exhibited by a token event of his composition, and can be picked out demonstratively. Thus Stanley claims:

> The 8 year-old Mozart can assert the proposition that constitutes his knowledge how to compose a symphony; he can just say say, while composing it, the German translation of "this is how I can do it". (Stanley 2009, p. 6).

The lack of intuitive plausibility of this construal of sophisticated forms of know how, which do not seem reducible to knowledge of methods for or *means* to the achievement of

ϕ-ing, has I think two related sources. They have been best brought out in critical discussions by Wiggins (2012) and Hornsby (2011). First, an answer to a question is expected to be informative. If one has a query about how to compose a symphony, or even about how Mozart composes a symphony, to be told at some instant in his months of work that 'this is a way for me to compose a symphony', is unhelpful. Of course his knowledge how to compose may be being exercised and on display at that very moment, but it is not available to be appropriated to any degree by the audience. This is not a problem that modes of presentation can be appealed to in solving. We have not gone away with knowledge of a way to compose a symphony 'in theory', only failed to engage that knowledge under a practical mode of presentation. We are simply none the wiser.

Certainly it can be replied (as I do on Stanley and Williamson's behalf elsewhere[23]), that the impossibility of *communicating* knowledge–how by means of demonstratives does not show that the knowledge is not on display at the time demonstrative reference is made, or that the demonstrative reference fails to pick it out. An epistemological objection to the metaphysical proposal that knowledge how to ϕ *is* knowledge of a way to ϕ such as can be exhibited in a token event of ϕ-ing cannot be decisive. But given how the proposed conception of knowledge how is supposed to fit into a semantics of embedded questions, the fact that a canonical expression of the putative answer communicates next to nothing, does seem problematic. Stanley claims that Mozart expresses his knowledge-how when he sits at the piano and says 'this is a way I compose a symphony', but in what sense does he express it? Can a proposition be semantically expressed by a sentence, even if no hearer of the sentence uttered is capable of recovering the proposition in question?

[23] Waights Hickman unpublished b.

The second problem is deeper. When Stanley talks about different ways of fielding a fly ball, this suggests a fairly intuitive and fairly fine-grained way of individuating ways to ϕ. But if we apply that conception of ways to an analysis of (e.g.) knowledge how to run a business, or knowledge how to navigate at sea (two of Wiggins' central examples), it becomes clear that the ways of ϕ-ing that the agent will exhibit in the course of exercising his knowledge-how are indefinite, heterogeneous, and not predictable in advance—they may involve individuals and conditions that do not yet exist or that the agent has no acquaintance with beforehand. His knowledge how to ϕ prepares him for these eventualities when they arise, resulting in his ϕ-ing in whichever way or ways are appropriate under the circumstances. But no one of these specific ways of ϕ-ing (nor any finite number of them), exhibited in the face of some particular contingency, can be constitutive of the *general* knowledge how that ensures his counterfactual success in ϕ-ing, whatever the circumstances. It is precisely because the way that the agent ϕs on any particular occasion gives only the most partial glimpse of the general knowledge at his command, that an attempt to explain or convey his knowledge how by making demonstrative reference on the particular occasion is bound to be uninformative (Cf. Hornsby 2011, §2.2).

None of Stanley's responses, in his discussion of novelty and skill (2011, ch. 7), succeed in defusing the crux of the challenge. He points out that propositional knowledge-how, being either itself a multi-track disposition or associated with one, will be practically expressed in a range of different ways; that there are motor skills that one trains in the application of know-how, that are partly responsible for the fluid application of knowledge-how in

changing circumstances; and that experts know not only one way to ϕ, but a range of ways to ϕ (see the quote above concerning the baseball player). All these claims may be true, but they do not help to answer Hornsby and Wiggins. For one way of putting their point is that there is a mismatch between that aspect of knowledge-how *qua* state or power of mind that Stanley acknowledges—its indefinite versatility and aptness to be expressed in multifarious and fundamentally novel circumstances—and its putative *content* on Stanley and Williamson's view.

An analogy might help to drive the point home. The indefinite productivity of *linguistic* knowledge lead Chomsky and others to maintain that a speaker's linguistic competence must consist in knowledge of a generative grammar, a system of general rules from which the interpretation of any sentence of the speaker's language can be derived. Knowledge of that sort, i.e. with that sort of *content*, would indeed equip a speaker with the resources to understand and compose indefinitely many novel sentences of her language. For the knowledge is comprehensive and systematic, it is not knowledge of a set of facts concerning specific sentences, but knowledge of general underlying principles that allow the interpretation of *any* arbitrary sentence. Now, knowledge-how is in general indefinitely productive in a similar way to linguistic competence. It equips the agent to act correctly, rather than interpret correctly, in an indefinite range of unforeseen eventualities. What sort of knowledge could explain this productivity? Hornsby's and Wiggins' point is that it cannot be knowledge of any number of specific ways to ϕ exhibited in specific token-events of ϕ-ing. These ways are too finely individuated, too heterogeneous, too numerous, and unpredictable by the agent in advance of the moment of action, for any finite range of such ways to constitute general knowledge-how.

I think that Stanley and Williamson's view can be coherently defended from these objections, simply by maintaining (Stanley's comments about the outfielder no withstanding) that ways to ϕ are *not* finely individuated, indeed that they are so coarsely individuated as to be exhibited in all of an agent's indefinitely heterogeneous expressions of their knowledge how. That is, if the challenge is to come up with an account of Stanley and Williamson's ways to ϕ that will ensure that grasp of a way could constitute knowledge-how to ϕ, it can be done simply by stipulating that a way to ϕ must have whatever properties a plausible object of knowledge how to ϕ has, e.g. by denying that ways to ϕ are equivalent to methods, means or procedures for ϕ-ing (since not all knowledge how is plausibly construed as knowledge of methods, means or procedures) and by maintaining that ways are very coarsely individuated, coarsely enough that the same way can be exhibited in indefinitely heterogeneous range of ϕ-ing performances. But the resulting interpretation of 'way' is contrived and has no echo in ordinary usage, and the proposal remains fundamentally unilluminating and unexplanatory, in just the sense that an analogous proposal about language mastery would be unilluminating and unexplanatory, i.e. one on which mastery of the French language amounts to knowledge, for some ways w_1 to w_n, that they are ways for one to speak French. I do not think, then, that Stanley and Williamson's proposal can be made intuitively plausible, even if it cannot be decisively refuted or shown to be incoherent by the considerations so far adduced.

The view I extract from Ryle dovetails more closely with Chomskian cognitivism about linguistic competence. It would see a question about how to ϕ as addressed by reference to the methods, principles, procedures, rules or criteria whose practical application

constitutes ϕ-ing itself. The knowledge involved does not always take the same form. In some cases, one exercises know-how as a means to producing some independent good or product, as in baking, sculpting or violin-building, and here what one knows is a series of methods or procedures by means of which the final good may be produced. In other cases, the exercise of the know-how *is* the good or product, as in piano-playing, dancing, and speaking. Not every activity, craft or profession with respect to which one can have knowledge how fits the same mould, and accordingly, knowing how to perform in these does not everywhere entail knowledge of the same sorts of facts or rules. Thus there are procedures known for building a fence, sawing a dovetail joint or making an omelette, but there are few if any procedures for dancing flamenco or writing an essay. But flamenco dance and essay writing are nonetheless sophisticated activities constrained and sustained by principles, canons and techniques, and accordingly knowledge how to do them is not a brute capacity, it is a genuinely cognitive achievement. Justice is done to this when we see that to know how to do these things is to have internalised the relevant principles, canons and techniques—to know them practically, or in the executive way, such that one can apply them in practice.

The view I recommend allows for the heterogeneity in practices with respect to which one can possess knowledge-how. But concerning each practice, it also explains the unity of knowledge how to perform: it traces the indefinitely varied expressions of an agent's knowledge how to ϕ, in teaching others how to ϕ, writing about ϕ-ing, criticising and evaluating ϕ-ing, and ϕ-ing itself in myriad different circumstances, to a single stock of knowledge, namely knowledge of the general rules, principles, techniques, methods etc. whose application sustains the practice.

Finally, for an activity to be rule or procedure-governed is not for it to be so *determined*. No violence is done to the ingenuity and spontaneity manifested by agents exercising their knowledge how, by seeing that knowledge as propositional in content. It is the use to which knowledge is put that is creative and spontaneous; the knowledge does not consist in creativity or spontaneity.

6. Knowing Practically and Practical Modes of Presentation

One of the innovations in Stanley and Williamson's work is the idea of a practical mode of presentation: a way of thinking about propositional content that gives rise to, or is associated with, a distinctive disposition to action. They are postulated to explain from an intellectualist perspective, why there could be two agents who both know the relevant proposition(s) concerning how to ϕ, yet only one of whom is naturally credited with knowledge how to ϕ—typically, such an agent will also *be able* to ϕ. Thus according to Stanley and Williamson, if Hannah and John both know the same proposition(s) about bike-riding, but only John can ride, it will "seem false" to say that Hannah knows how to ride. This is, on their view, because unlike John, Hannah does not entertain the proposition concerning how to ride under a practical mode of presentation. (Stanley and Williamson 2001, pp. 426-8). They associate practical modes of presentation with the use of the construction 'knows how' (while remaining neutral as to whether these are part of the semantic content of the knowledge ascription), but as they acknowledge, modes of presentation are often thought to be introduced by *contexts* of ascription. In connection with ascriptions of know how, how much in addition to propositional knowledge is imputed to the agent by the ascription does seem plainly context-dependent. Thus one of the cases that Stanley and Williamson themselves cite as counterexample to the "capacity

thesis" discussed in §§2-3 above, is that of a ski instructor who can teach others how to perform a complex jump, though she is not able to perform it herself. According to Stanley and Williamson, the instructor knows *how* to perform the jump, although she isn't able to perform it (Stanley and Williamson 2001, p 416). I think this is the right verdict, but it seems clear that whether an ascription of knowledge how to her will "seem false" depends on context. If the context is one in which we are looking for a candidate to perform the jump in a competition, an ascription of knowledge how to the instructor will seem false or at least misleading[24]. On the other hand, if we are looking for an instructor to lead a group of students preparing for such a competition, the ascription will seem true and pertinent.

As for what practical modes of presentation or ways of thinking themselves entail, Stanley and Williamson never provide a substantial account; the only serious attempt to construct one is due to Carlotta Pavese (2015). Pavese's discussion is largely devoted to a defence of construing computer *programs* which represent *algorithms* for the computation of *functions* as an example of Fregean practical senses. She goes on to suggest that programs are but one species of the genus of practical senses, but of course that certain features of programs generalise to the genus, and so can teach us something about practical modes of presentation as such (Pavese 2015, pp. 12-13).

Pavese begins from the idea that all knowledge how to ϕ may be construed as knowledge how to perform the *task* of ϕ-ing (*ibid,* p.2). A task may be modelled as a function. In the cases we are interested in, i.e. tasks of bike-riding, wood-chopping and all the other

[24] Stanley and Willimason (2001, pp. 427-8) remain neutral as to whether practical modes of presentation have semantic or merely pragmatic import. Thus they are neutral as to whether an ascription of know-how where the knowledge is *not* engaged under a practical mode of presentation is, in certain contexts, false or merely misleading.

activities one can know how to engage in, we may suppose that the function is one that takes relevant scenarios as inputs, and yields (a range of possible) successful ϕ-ings as outputs (*ibid*, p. 3). But now, a function may be computed in more or less circuitous ways, by means of different *algorithms*. There can be a range of different algorithms for the computation of a function, just as there can be a range of different ways to field a fly ball (*ibid*, pp.2-3). Further, algorithms can be represented in different ways, by means of different *programs*. Programs are not, Pavese insists, bits of language but rather what that language expresses (p.5); thus, she maintains, they qualify as *modes of presentation* under which an algorithm can be represented. In particular, the semantic values assigned to programs under operational semantics are good candidates for practical senses. Operational semantic values (OSVs) are inferential rules of a certain sort, and executing a program is then a matter of performing an inference according to these rules (pp.5-6). Pavese maintains that OSVs are distinctively *practical* senses because their grasp presupposes certain distinctively practical capacities, *viz*. the capacity to (come to be able to) follow the inferential rules, and yields certain distinctively practical abilities, *viz*. the abilities to follow the inferential rules. Indeed, one cannot grasp the OSVs without having the ability to follow these rules, just as, Pavese says (citing Rumfitt 2011), one cannot grasp the inferential rule of modus ponens without being able to follow it (*ibid*, pp. 10-11). In fact, assuming there are no internal or external impediments, "an interpreter cannot understand an OSV to ϕ without acquiring the ability to carry out the task of ϕ-ing."(p.12).

So, what does all this tell us about the practical mode of presentation under which someone with, say, knowledge how to swim purportedly entertains a proposition of the form *w is a way for me to swim*? Not all practical senses are programs, according to Pavese,

and not all referents of practical senses are algorithms. Rather they are *methods*, of which algorithms are one species. But practical senses *are* all constituted by inferential rules of a certain kind, which present parts "of the relevant task *in the form of a command*, according to a certain method's way of breaking that task into parts." (p.13). Grasping a practical sense affords one the ability to follow a rule to ϕ, and presupposes the capacity to come to acquire that ability. And since the ability to follow a rule to ϕ entails, in a qualified sense—i.e. barring internal impediments such as injury or external impediments like the absence of necessary tools—the ability to ϕ, knowing how to ϕ entails, in a qualified sense, the ability to ϕ, as opponents of intellectualism have long claimed (*ibid*).

To sum up, combining Stanley and Williamson's account with Pavese's yields the following position:

- Knowledge how to ϕ = knowledge, for some way *w*, that *w* is a way for one to ϕ.
- If *w* is thought about under a practical mode of presentation, it is known 'through' grasp of a rule to ϕ, where this entails the *ability to follow* a rule to ϕ.
- Knowledge how to ϕ entails the ability to ϕ in a qualified sense, i.e. barring internal or external impediments. (*ibid* pp.12-13).

I do not think the last contention is crucial to Pavese's overall account, but I shall register my scepticism on this score nonetheless. The claim seems to me untrue unless those capacities lack of which would constitute an 'internal impediment' are assumed to include *specific ϕ-ing capacities* that have been honed by ϕ-ing. And these are constitutive of the ability to ϕ, not extraneous prerequisites for its exercise. Consider, for example, the fine

motor skills and dexterity of a pianist—the skills that enable her to translate her knowledge into a controlled, fluent and articulate movement of her fingers. These are not abilities that she has *in advance* of the ability to play the piano, rather they are partly constitutive of that ability. So they cannot be counted among the capacities absence of which would constitute a mere 'internal impediment' to the exercise of the ability to play, since absence of these *just is* absence of *the ability to play*. If the pianist is injured and those abilities are lost over time, recovering from the injury will not be enough to ensure that she can immediately play as well as she once could. Thus knowledge how plus normal physical function plus occasion or opportunity does not add up to ability; knowledge how does not entail ability in any interesting sense.

I turn now to Pavese's main proposal concerning practical senses. Pavese's starting point, the idea that all knowledge how may be thought of as knowledge how to *perform a task*, I think is also somewhat dubious and needs argument: playing chess, playing jazz, drawing, kissing, driving, singing, comforting a friend, making people laugh, running a business and commanding authority are not naturally construed as tasks, but they are things one can know how to do. Moreover, although in all of the mentioned activities there are certainly either techniques employed, principles applied or criteria met, I do not think it is true that grasp of a *method* is everywhere key to knowing how—e.g. chess, driving, and jazz are not activities or arts which seem to me characterised by methods. As I stressed earlier, the activities with respect to which one can have knowledge how are highly diverse, and that diversity is reflected in the content of the knowledge-how as well as, I shall argue below, in the capacities that are relevant to *practically* knowing the relevant techniques or principles. But I am not entirely unsympathetic to Pavese's proposal. For I agree that the exercise of some forms of knowledge how entails rule-following. It follows

from the account I recommend that someone exercising knowledge how to play chess follows the rules of chess, for example. (Although I am unsure whether chess rules are of the right sort to fit Pavese's proposal—they are supposed to be rules that represent parts of the 'task' of chess-playing in the form of a *command* (Pavese 2015, p. 13)). Considering that Pavese's account purports to be Fregean, so that the practical sense is a *constituent of a practical proposition* grasped by someone who knows how, this brings her version of intellectualism very close indeed to the conception of knowing how I find in Ryle. I also agree with Pavese that one cannot know a rule without being able to apply, use or follow it.

However, and this underlies my core motivation for rejecting Pavese's account, like any knowledge, knowledge of a general rule can be used, applied or followed in multifarious specific ways, and with different levels of practical or theoretical ingenuity. Just as someone who knows that *p* cannot be expected to be able to draw *all* the inferences that *p* licenses or perform *all* the actions *p* provides grounds for, someone who 'grasps' a rule *R* cannot always be expected to be capable of making *all* its conceivable applications. Someone who grasps a rule of trigonometry may apply it competently in response to familiar theoretical questions, but fail to see its bearing and put it to use when faced with a practical problem. Such a person, it must be allowed, has some ability to follow or apply the rule, but not in practice or in practical contexts. Some such difference in the way an agent's knowledge (or 'grasp') of a truth or principle can be exploited and manifest is precisely what practical modes of presentation were invoked to explain.

The point I am making is essentially Rylean[25]. Mere reference to the rules that an agent grasps is not illuminating of what equips her to apply those rules with special proficiency or practical efficacy. To take Pavese's own example, consider the rule of Modus Ponens. Pavese maintains, rightly in my view, that this rule like all rules is one that an agent can be said to grasp only if he can follow it. But to what degree must one be able to follow it? In what range of cases? We can easily imagine two agents, one more trained and proficient in complex reasoning than the other, who differ in the range arguments in the formulation, consideration or assessment of which they can follow modus ponens. It would be wrong to say that the the person less trained in argumentation does not, or does not "really" grasp modus ponens, for we may assume that he follows the rule without difficulty in a wide range of cases. What we want is some account of what distinguishes the two reasoners in the use they are apt to make of the rules they grasp, that is, differences in what they are able to do. This is the sort of difference in competence manifested despite common information that Stanley and Williamson's practical modes of presentation were supposed to explain, but they cannot explain it if what a practical mode of presentation affords the knower is just grasp of the inference rules themselves. The question remains: what differs in the *kind of grasp* the two agents have of these rules? I reject Pavese's account ultimately because I do not think that practical modes of presentation on Pavese's construal can do the work that intellectualism needs them for.

[25] Indeed, Pavese's is the kind of intellectualist move that Ryle's regress arguments purport to target. The crux of those regress arguments is that crediting an agent with possession of infinite information can never explain his ability to deploy it (in whatever special ways he does). Thus we can imagine an attempt to explain the difference between the more and the less capable reasoner, following something like Pavese's approach, in terms of further meta-rules for the application of, inter alia, Modus Ponens, which the more capable reasoner grasps but the other doesn't. And the Rylean reply will be that a less capable reasoner might grasp all those rules, but differ in the extent to which he is able to exploit *them*. Following the same strategy of explanation, we will have to attribute the difference between them to the expert's grasping some further tier of higher order regulative rules, launching us on a regress of higher and higher order regulative rules, each set postulated to buttress the practical efficacy of the set before.

To take another familiar example, consider the ski coach, let's caller her Annie, who can train agile students to perform a certain jump, though she is unable to perform it herself. What distinguishes her from the students who, with her training, are now able to perform the jump? Is it a difference in what rules they grasp or what they can do on the basis of that grasp? Stanley and Williamson's approach recommends the difference be accounted for in terms of modes of presentation[26], which also gives a plausible explanation of the apparent context sensitivity of an ascription of know-how to Annie, mentioned above. And yet it is implausible that the students grasp rules for performing the ski-jump that their teacher ignores. Whatever rules are relevant to performing the jump are precisely the ones she has been teaching the students to follow, something she could not do if she lacked a grasp of them herself. Pavese's account cannot succeed in explaining what is *practical about the grasp* of how to ϕ that one has under a practical way of thinking.

My own view, spelt out in previous sections, is that rules, principles, methods, techniques of ϕ-ing and so forth, are what someone with knowledge how *knows*. But importantly, on my view this is what knowing how to ϕ consists in irrespective of whether one is, in addition, *able to* ϕ—grasp of rules is not what gets one from mere knowledge to ability. I admit that knowing (or 'grasping') these rules entails *some* ability to apply them[27], just as knowing any fact or truth entails *some* ability to exploit it. But as this ability is gradable

[26] Compare their illustration involving Hannah and bike riding (Stanley and Williamson 2001, pp. 428-9).

[27] Thus I agree with Stanley and Williamson that Annie the ski coach knows how to perform the jump, for she can apply her knowledge of principles and techniques in teaching and assessment, but I deny that Hannah has in any sense acquired *knowledge how to* ride a bicycle merely by being pointed in the direction of John on his bike and told 'that is a way for you to ride a bike'. There is nothing Hannah can do with knowledge that John's is a way she herself could ride a bike that people with genuine knowledge how can do ϕ, even when they have lost their ability to ϕ.

along different dimensions—one may be more competent in applying the rules, principles or procedures in some contexts rather than others—to credit someone with its possession will not yet suffice to explain their special proficiency *in ϕ-ing*. I maintain that someone can have knowledge of the relevant rules (/methods etc.) and yet lack the ability *to ϕ* or to ϕ well; and thus that the question of what it takes to know practically, such that one is able to deploy one's knowledge effectively in ϕ-ing, remains so far unaddressed.

To begin with, it is worth considering whether what is at issue is anything like a mode of presentation. As I said earlier, ascriptions of knowledge-how appear to be context-sensitive, sometimes but not always carrying an imputation of the corresponding ability. The context-sensitivity of attitude reports is one way to motivate an appeal to modes of presentation, but it does not make that appeal compulsory. For if an attitude report is context-sensitive, it may be either the complement-clause or the verb that is responsible. Thus 'Annie knows how to perform the ski jump' may be context-sensitive because 'how to perform the ski jump' is context-sensitive, or because the verb 'know' is[28]. If it is the latter, then it seems we should not be theorising in terms of modes of presentation. Instead, we should be seeking to understand the special properties of the knowledge relation denoted, e.g., in the ascription 'Annie knows how to perform the ski jump' in a context in which we are looking for a candidate to perform the jump in the competition. The issues here are subtle[29] and this is not the place to settle them, but it seems to me that postulation of practical *modes of presentation* is not the only option we have, assuming intellectualism is

[28] This is not to say that 'knows' or 'knows how' is ambiguous, as Wiggins (2012) suggests. Wiggins' contention is that 'knows how' sometimes ascribes propositional knowledge, but sometimes a non-propositional, irreducibly practical competence. The question I am raising, on the other hand, is whether 'knows how' may in some cases be taken to ascribe a distinctively practical knowledge relation to a propositional content, rather than a practical way of representing the content known.

[29] See e.g. Dorr 2014.

true, to make sense of the distinctive conditions associated with some imputations of knowledge how to ϕ. My aim in what follows is to make progress in understanding those conditions, and I shall say nothing further about the choice between a practical knowledge relation and a practical way of thinking about propositional content.

Let us first revisit the most familiar example for which a practical mode of presentation was invoked: the case of Hannah, who in some sense knows how to ride a bicycle—she knows the relevant propositions—but is not able to ride (Stanley and Williamson 2001, p.428). It is natural to assume that she will not succeed in putting her knowledge how into action because she lacks the requisite balance, muscle memory, and specific fine motor skills, all of which might be developed through practice (assuming she has no injury or disability). What does all this have to do with failure to grasp a practical mode of presentation? I think the natural way to conceive the mode of presentation that would equip Hannah to do what she currently cannot with the knowledge she has, is not as consisting in rules, but as consisting in some form of *haptic* information or feedback, which experience of physically riding *herself* would afford. If Hannah practices riding a bike, she will not only know what she has to do to ride a bike, she will also have experienced how it *feels* to ride, for her body to be appropriately positioned and subject to air resistance, friction and gravity as the bike moves. The more she practices, and the broader the range of positions, speeds, terrains and moves she tries out, the more information of this sort she will acquire. This plausibly affords Hannah a new way of thinking, a *haptic* way of thinking, through which she grasps the propositions she knows concerning how to ride; and because she grasps them in that way she comes to be able to translate her knowledge into the right movements when on the bike, and to ride.

It seems plausible that haptic information is at least partly constitutive of the practical grasp agents have of their knowledge how to ϕ in a wide range of cases, e.g. when someone knows how to swim, how to plane, how to tie their shoelaces, or how to climb. But not all the actions or activities one can know how to perform are physical in this way, depending, or depending so crucially, on fine motor skills, balance, and so on. People also know how to add, how to navigate, how to give speeches and how to build cabinets. Haptic information may often be partly relevant to the agent's ability to translate knowledge into action, but it cannot always be sufficient.

If we want to better understand the conditions that ensure someone's effective use of their knowledge how ϕ *in ϕ-ing*, the natural question to ask is what, besides *some* kind of grasp of propositions concerning how to ϕ, goes into someone's standing *ability* to ϕ, or general 'competence' in ϕ-ing. The richest discussion of this question that I am aware of is in David Wiggin's (2012) paper on practical knowledge. Wiggins presents his account in explicit opposition to Stanley and Williamson's intellectualism, for he identifies knowing how to ϕ (in some cases) with general competence or expertise in ϕ-ing—what he calls 'practical knowledge' or 'practical-cum-agential knowledge' . But as long as we distinguish knowing how to ϕ and expertise in ϕ-ing (see §1 above), there is no obstacle to endorsing intellectualism about know how *and* some of Wiggins' positive contentions concerning the conditions of expertise.

Wiggins (2012) argues that mastering an art, craft or profession involves habituation into its characteristic *ethos* or outlook. This means not only acquiring knowledge but a whole complex of goals, concerns, skills, emotional sensibilities and dispositions, attitudes and

character traits, under whose integrated control judgments will be reached and actions undertaken in the full spectrum of possible circumstances in which the agent's expertise may be called on (Wiggins 2012, p. 100). Internalisation of its *ethos* is, according to Wiggins, what is *essential* to success in a profession or activity, for it is this that gives agents a grasp of "the sprit in which they are to act" in indefinitely heterogeneous and unforeseen circumstances (*ibid*). Through habituation into an ethos, it becomes second nature to the expert to attend to particular kinds of consideration and see their bearing on her own ends; to judge correctly where she should focus and when she should act, as well as how (*ibid*, p. 112).

Wiggins's appeal to habits, sensibilities, and values united under an ethos is supposed to provide a competing account of practical knowledge to that offered by intellectualist. He maintains that no amount of propositional knowledge could budget in advance for the indefinite range of contingencies that the expert must be able to navigate, and that this capacity is irreducibly practical, acquired through habituation. As I said above, however, I think the considerations to which Wiggins draws attention do not tell against intellectualism concerning *knowledge how.* First, by my lights knowledge how to ϕ consists in knowledge of *general* principles, methods or procedures that, being general, have application in an indefinite range of different circumstances; in that sense, *pace* Wiggins, they *do* budget for them. Second, it should be uncontroversial that an expert's counterfactual success in ϕ-ing in heterogeneous circumstances is not sustained by propositional knowledge alone; the exploitation and use of knowledge, however rudimentary, invariably depends on other psychological and physical conditions of the agent. This is no threat to intellectualism. The fact that other attitudes and qualities of mind are necessarily exercised in concert with knowledge how when it is put to use, does

not show that knowledge-how is not itself propositional. Nevertheless, the specific kinds of conditions on competence or expertise that Wiggins identifies may yet be of interest for intellectualism. What I propose is that some of these conditions enter into what it takes to engage one's knowledge under a practical way of thinking, and accordingly that they can help to explain, in cases where appeal to a haptic mode of presentation proves unsatisfactory, the gap that can emerge between knowing the relevant propositions concerning how to ϕ, and being able to ϕ—that is, being able to apply one's knowledge effectively in ϕ-ing.

In this context, the notion of an *ethos* of ϕ-ing is best dispensed with. Not only is it vague, but there is plainly no hope of it shedding light on the practical ways of thinking involved in, e.g. knowing how to tie one's shoelaces or how to bake a carrot cake—these would not normally be deemed arts and are not the sorts of activities with which Wiggins' account is concerned. I think instead we need to look at specific attitudes, dispositions or values of the sort alleged by Wiggins to sustain expertise, and ask whether and how they relate to the grasp the agent has of the knowledge at her command.

It seems incontestable that mastery of a craft or profession very often does, as Wiggins maintains, involve acquiring habits, sensibilities and concerns as well as knowledge. A good woodworker learns the methods and principles of his craft and develops a steady hand, but also learns to exercise patience, diligence, attention to detail, innovativeness and ingenuity in applying his knowledge. He must be concerned with and appreciative of particular features of the wood and tools he has to work with. Similarly, while a professional flamenco dancer must certainly perfect techniques which require agility and dexterity, she must also show the right spirit or temper—pride, impetuousness, and

sensuality. A good teacher needs, among other things, empathy to notice and respond appropriately to aspects of her students' behaviour; a successful businessman must be fair, prize innovation and look to profit; a surgeon must remain calm under stress; a tennis player must be alert.

Moreover, if such attitudes and qualities of mind as have just been mentioned can sometimes be integral to expertise, we can reasonably assume that they are so precisely because they improve the practical efficacy with which the agent is able to to exploit her knowledge how to ϕ. But of course, not any factor that plays this role of supporting the practical efficacy of knowledge need determine or constitute a (practical) *way of thinking*. The absence of depression may be a condition on a flamenco dancer's being able to deploy her knowledge in dancing, but, intuitively, this is because depression would impede the exercise of her standing ability to dance, not because mental wellbeing is partly constitutive of a practical grasp of how to dance. Moreover, some dispositions and concerns may be acquired as expertise is cultivated because they *issue from* a practical way of thinking, not because they *issue in* that way of thinking. Nevertheless, there are cases in which values, goals, concerns, sensibilities and habits plausibly govern the agent's grasp and use of knowledge how, in such a way that an attribution of knowledge how also imputes these specific psychological conditions to the agent. Thus consider knowledge how to comfort a friend in distress. By my lights, the exercise of this knowledge involves following principles and methods. But the general principles will not be applied with success in a specific case by an agent who lacks sympathy and compassion; in their absence, the agent may well be blind to the fact that this is a situation in which comfort is needed. Thus among other things, sympathy and compassion dispose the agent to attend to, or notice, features of a situation on which her standing knowledge bears, and to attend

to features that determine the specific ways in which that knowledge should be applied. The critical role played by these emotional dispositions in the exercise of knowledge how to comfort a friend in distress, is mirrored in what I take to be the fact that a person who has this knowledge is necessarily someone endowed with sympathy and compassion. These are constitutive conditions on their knowing how in the usual sense, i.e. in the sense that carries an implication that they have also the corresponding *ability*. Similarly, while it is not immediately clear *which* qualities of character are presupposed of someone who knows how to babysit a toddler, it seems clear that there are *some*—qualities that are critical to the conversion of knowledge into ability, and that we are not entitled to assume any arbitrary human being will possess.

Turning to crafts and professions, it seems to me that the diligence and care exhibited by a good woodworker, virtues that he must possess or cultivate in order to master his craft, are ones that support the effective use of his woodworking knowledge in practice, rather than being constituents or consequences of that knowledge. Do they flow from a practical way of thinking about how to work wood, or issue in that way of thinking? One reason for thinking it is the latter is that if a person shows care and diligence in woodwork, he is likely to show care and diligence elsewhere, suggesting these qualities are not products of his special ability but are among its conditions.

The pride, impetuousness and sensuality of a good flamenco dancer seem to me also to govern, and so be reflected in, the specific manner in which she interprets and manifests the principles and techniques of flamenco dance. A dancer lacking these qualities is, by my lights at least, not apt to pull off compelling performances, and in that sense to to say

someone knows how to dance is also to say, or at least imply, that they do possess these qualities to some degree.

This still leaves us some way from a precise and comprehensive account of practical ways of thinking. What I am putting forward is an approach, or a direction of enquiry. First, I maintain that conditions on general ability or competence in ϕ-ing hold the key to the practical grasp an agent must have of *how to ϕ* if he is to have counterfactual success *in ϕ-ing*. Second, I have argued that among these conditions, in many cases, is the haptic feedback or information afforded by physical experience in ϕ-ing, information in light of which the knowledge how may be differently grasped. Third, I have argued that values, emotional dispositions and qualities or virtues of character can be conditions on knowing how to ϕ *in the executive way*.

7. Conclusion

I have argued that Ryle's *positive* discussions of knowing how to ϕ suggest a conception on which this knowledge consists in the practical or 'executive' grasp of the rules, principles, methods or procedures that characterise ϕ-ing itself. There is a surprising proximity between this Rylean view and the intellectualism of Stanley and Williamson, although there are important points of difference, canvassed in the preceding sections. Combining Pavese's view of practical senses with Stanley and Williamson's proposal brings their account especially close to the the one endorsed here, for on the resulting view knowing how (under a practical mode of presentation) involves grasping a rule to ϕ, and in turn being able to follow it. However, I believe that Pavese's conception of practical modes of

presentation prevents them from doing the explanatory work intellectualism needs them for: that of accounting for the difference between two agents with common knowledge (how) who nonetheless diverge with respect to the practical efficacy with which they deploy it. Instead, I suggest that a practical way of thinking can involve both haptic information and the possession of values, sensibilities, and qualities or virtues of character, which influence the access to and use of knowledge, among other things, by the guidance of attention.

Novelty and Knowledge (How)

Introduction: A Problem for Intellectualism

The purpose of this paper is to examine and defuse what I believe is the most powerful and philosophically interesting source of resistance to contemporary intellectualism[1], the thesis that knowledge how to do something, e.g. how to play the violin or build a cabinet, is a species of propositional knowledge—that it is a species of knowledge *that*.[2]

This challenge is one seldom-acknowledged thread in Gilbert Ryle's first purported *reductio* of intellectualism, but it has received its most explicit and forceful elaborations in the critiques of Jenifer Hornsby (2011) and David Wiggins (2009; 2012). Sometimes referred to as a problem of novelty[3], the challenge begins from the versatility of knowledge how. Someone who possesses knowledge how is able to act appropriately in the face of novel and heterogeneous conditions in which that knowledge is called upon. For it is *in general*, across an indefinite range of circumstances, that one knows how to ϕ if one does know how. From this it can seem to follow inexorably that knowledge how cannot reduce to propositional knowledge: no finite number of propositions can embody information sufficient to budget in advance for all the eventualities in which one must know how to act if one knows how to ϕ. As David Wiggins has put it, there can be "no proposition, no conceivable form of

[1] This contrasts with "Intellectualism" and "The Intellectualist Legend" as Gilbert Ryle invoked these terms, which stand for a cluster of pseudo-Cartesian doctrines concerning the relation between cognition and action. I return to the differences between Ryle's target and the positive thesis advanced by his critics today in §3 below; see also Stanley's 2011, ch. 1.

[2] I assume here that factual knowledge is propositional, but adherents to a substantial fact/proposition distinction may understand intellectualism as the thesis that knowledge how to ϕ is knowledge of facts or truths.

[3] Toribio 2008; Stanley 2011, pp.181-185.

words, that answers to and condenses…the whole substance of [the agent's] control of an indefinite spread of possible contingencies." (Wiggins 2012, p. 120).

The first section below elaborates this challenge in greater detail. Thereafter, I shall be concerned with dismantling the challenge *as directed against intellectualism*, or the view that knowledge how is propositional. Its true purport and interest, I believe, lies elsewhere, in its connection with the versatility of knowledge in general, and in particular knowledge with inherently general content, e.g. rules, types, properties or universals. My positive aim in this paper will be to redirect attention towards these issues. My contention is that the proper target of the challenge under discussion is not an intellectualist view of know-how, but what Searle described as the idea that "all intentionalistic mental life and all cognitive capacities could be entirely reduced to representations" (1991, p.152), and an associated conception of the aetiology of action that exhibits that mental life. This was subject to persistent critique by Wittgenstein in connection with a variety of intentional attitudes including imagining, understanding and, most notoriously, knowing a rule. I shall refer to this conception of intentionality as representationalism, or representationalist, and draw out how the *reductio* arguments against it in Searle and Wittgenstein parallel those in Ryle's, and neo-Rylean, critiques of intellectualism.

If my diagnosis is correct, it explains why the putative category distinction between "knowing how" and "knowing that" has had such extraordinary influence across the philosophy of mind, language, and action in the last half-century, why its interrogation has generated such heated controversy, and why it has appeared to divide two deeply opposed

metaphysics of mind[4]. For on my account, the contention that *knowing how* is irreducible to *knowing that* was and continues to be a misdirected attempt to register a quite different and genuinely foundational, paradigm-shaping commitment: that representationalism is false, for it is only through human intelligence that representation (e.g. of a fact or state of affairs) is possible—it is only to human intelligence that a representation *represents*—and that intelligence cannot itself, without generating vicious regress, be reduced to the possession of a mental representation. The reduction of knowing how to knowing that is *not* a reduction of this sort, because knowing *how* is not a species of intelligence, and knowing *that* is not the possession of a mental representation.

1. The Argument from Versatility

The challenge is how to do justice, on an intellectualist view, to the fact that knowledge how to ϕ is knowledge that equips one to negotiate a potentially infinite range of contingencies in the practice or activity of ϕ-ing. If Carmen knows how to dance flamenco, she knows how to dance to different pieces, with different partners, accompanying different musicians; and indeed under all sorts of circumstances, the details of which she cannot anticipate in advance. When the next occasion for dancing arises, she will know how to act: her general knowledge how to dance will 'work', so to speak, precisely by enabling her to respond appropriately to unforeseen contextual differences.

Why is this fact about knowledge how supposed to cause trouble for intellectualism? Well, don't propositions embody finite information? So how could knowledge of propositions carry one through potentially infinite contingencies? Moreover, how can any such

[4] See the introduction to Bengson and Moffett's (2011) *Knowing How*.

knowledge equip one for contingencies that are *unknown in advance?* Knowledge how equips one to negotiate particular *future* conditions one might never have imagined; for example, in a possible future Carmen will bring to bear her knowledge how to dance in her interaction with another individual about whom she presently knows nothing. The information about that individual to which she will, bringing her know-how to bear, appropriately respond, cannot be available beforehand—not until the circumstance of their meeting in fact materialises. So how could being armed now with any amount of propositional knowledge prepare her, as her present know-how surely does prepare her, for what must be fundamentally *novel*?[5]

On the one hand, propositional knowledge that is narrowly tailored to the particularities of some one occasion of ϕ-ing, say knowledge that *this* is a way for one to ϕ, seems to be too specific to amount to knowledge how to ϕ in general (cf. Hornsby 2011, §2.2). Armed with such context-specific knowledge, one will not know how to act in *novel* circumstances, i.e. as soon as the context of ϕ-ing changes in some respect. And indeed, it seems that knowledge of *n* ways to ϕ (*n* way-propositions) will not add up to knowledge of the requisite generality as long as *n* is finite. There will always be some possible contingency which has not been captured under the *n* ways, in which one *would* respond appropriately if one knew how to

[5] As in debates over the nature of language mastery, the central question is perhaps not how an *infinite* capacity could be encompassed in *finite* resources, to echo Davidson (2001[1984]), but rather how finite knowledge could equip one for future *novelty* (Evans 1981). This analogy already suggests that the "problem of novelty" cannot speak unambiguously against intellectualism. For in context of the linguistic debate, that very issue has been appealed to, since the Chomskian revolution in linguistics, to *motivate* the ascription to speakers of complex propositional knowledge, in this case, of general rules governing the language as a whole. If this linguistic knowledge is not postulated it remains unexplained how individuals can, on the basis of a very limited linguistic evidence base, come to be in a position to formulate and understand an indefinite number of *novel* sentences.

ϕ; as has been said, it is of the essence of knowledge how that it prepares one for the unforeseen.

If it is absurd to suppose that in knowing how to ϕ one knows an infinite number of context-specific propositions pertaining to possible eventualities one cannot anticipate in advance (Wiggins 2012, p.p100; 105), it must be better for the intellectualist to assume that one knows certain *general* propositions pertaining to ϕ-ing. But now, on the other hand, knowledge of general propositions pertaining to ϕ-ing won't be specific enough to carry one through all the contingencies in which one's knowledge how might be realised, for they cannot, as Ryle noted, "embody specifications to fit every detail of the particular state of affairs."(Ryle 1949, p.20; cf. Wiggins 2012 p.104-5). It might be supposed that one simply sees how the general propositions are to be applied as each new contingency arises. But then some knowledge how has (apparently) been assumed: the knowledge how to apply the general propositions in the particular circumstances that confront one. If this knowledge is construed again as propositional we shall be back where we started, with indefinite context-specific propositions, except now a vicious regress also threatens: if one's knowledge how to apply the general propositions is embodied in additional, higher order propositions, won't one now also need knowledge how to apply those?

To sum up: it seems as if intellectualism cannot do justice to the fact that one's knowledge how to ϕ carries one through indefinite and diverse contexts of ϕ-ing: knowledge how to handle an indefinite range of contingencies could not be condensed in a finite number of propositions. The impossibility of squaring this circle may be brought out as follows: a general proposition doesn't give one enough detail concerning how to ϕ on *this* occasion; a

specific proposition gives one *only* that. Whatever type of proposition we go for, when we subtract the postulated knowledge of such propositions from knowledge how to ϕ, it seems there is a remainder that is stubbornly irreducible. And the irreducibility of this remainder is brought out most clearly by threat of vicious regress: "knowing how to apply truths cannot, without setting up an infinite process, be reduced to knowledge of some extra bridge-truths" (Ryle 1945-6, p.7).

The argument form versatility is, I believe, of great interest, but not because of its connection with intellectualism. Wittgenstein asked in connection with knowledge of meaning, and most notoriously knowledge of a rule, how it can equip one to apply it appropriately in an indefinite range of as yet uncontemplated eventualities; or how knowledge that *does* so equip one is possible. Closely analogous concerns can in fact be raised in connection with knowledge that is not conspicuously general in the way that knowledge of a rule, principle or methodology is thought to be. It can be asked how in coming to know a simple empirical fact one is thereby enabled to draw upon that fact in addressing all manner of as yet uncontemplated questions and tasks on which it bears; for if one were not so equipped, to some degree at least, one would not count as really knowing the fact in question (more on this in §3 below).

2. Intelligence, Ability, Knowledge and Knowing

Some preliminary points are in order, to clarify what the intellectualist thesis really is.

One issue to flag at the outset is this. Intellectualism says that knowledge how to ϕ is knowledge of a proposition or propositions. But what qualifies as a proposition? One could

object to intellectualism on grounds that knowledge how is a form of general or generic knowledge, and that while such knowledge *can* be put into words (perhaps even in a declarative sentence), what is then put into words is *not a proposition*. It might be a set of rules or principles, and rules and principles are sometimes thought not to be proposition types, because propositions are truth apt, while in some domain a rule may hold or not hold, but it cannot be true or false. Or knowledge how to ϕ might be knowledge of a property or properties of events of ϕ-ing: knowledge of a *way* to ϕ, as Stanley and Williamson (2001) maintain. But knowledge of properties may be thought to be non-propositional, even if it is possible to devise a sentence of declarative form corresponding to it[6]. According to many interpreters, Aristotle held this view of knowledge of universals, for he maintained that it is not possible to be in error about universals, only to grasp them or fail to grasp them[7]. The anti-intellectualist critique of §1 and the defence I will offer prescind from these issues. The argument from versatility may be thought of as targeting not only the idea that knowledge how is propositional in some narrow sense, but the idea that there is any form of words at all apt to express what someone knows if they know how to ϕ—there is no way or method of ϕ-ing, no set of rules, principles or procedures that condense an agent's knowledge how to ϕ; instead this knowledge is supposed to be *irreducibly* practical.

---

[6] Adrian Moore (2003, p.177; cf. 1997, pp.134-5) makes an objection along these lines against Stanley and Williamson's account, considering a treatment of knowledge of greenness on the analogy of their treatment of knowledge how to ϕ: "Consider my knowledge of what it is for an object to be green. On their view this is knowledge, concerning something, that that thing is what it is for an object to be green. But concerning what? … If I try to express my knowledge by indicating a green object and saying, 'This is what it is for an object to be green' what can I be referring to by 'this'? There does not seem to me to be any good answer. Nothing short of an unacceptable Platonism, it seems to me, can subserve the extension of their account to this case. I do not think that my knowledge of what it is for an object to be green is knowledge that anything is the case."

[7] Metaphysics Θ10 and Metaphysics E 4 (Aristotle 2006).

Against this contention, the view I implicitly defend through sections 3-4 is the view that knowledge how to ϕ is knowledge of the rules and/or principles, methodologies and procedures that characterise ϕ-ing itself; those that are sometimes put into words in the the teaching, evaluation and correction of ϕ-ing, but are most often explicitly articulated by critics and theorists of ϕ-ing rather than by practitioners[8]. Unlike most intellectualists today, what I take to be the most important motivation for intellectualism—and for this view in particular—is not any semantic analysis of knowledge ascriptions using an interrogative particle. Rather it is the fact that knowledge how is exhibited in activities and practices that are rule, principle or procedure-guided[9]. Many have argued that an agent's action cannot be guided by a fact or proposition *p* unless the agent knows that *p*.[10] If this holds for guidance by empirical facts, one would expect it to hold for rules, principles and procedures also. Thus an agent who is guided by principles in the course of an activity knows them, irrespective of whether he can articulate them in words—as I will explain below, I believe this was also Ryle's view. It will be more immediately clear how my preferred view relates to Wittgenstein's discussion of novel rule applications, and might withstand the argument from versatility. The points I make in its defence, however, hold for intellectualist views in general; and in §5 I will defend Stanley and Williamson's (2001) proposal from a version of the versatility argument advanced by Jenifer Hornsby against their account in particular.

[8] Cf. Ryle (1949, p.60): "The rules which the agent [exercising knowledge how] observes and the criteria which he applies are one with those which govern the spectator's applause and jeers."
[9] Ryle stresses this on virtually every page of both his 1945-6 and 1949 essays; see e.g. 1945-6, p.8.
[10] Stanley and Hawthorne (2008); Hornsby (2007); Hyman (1999, 2015).

Now more must be said to distinguish contemporary intellectualism in general, the view that *what* someone with knowledge how knows is a proposition or propositions, from other theses with which it is often conflated.

First, when Gilbert Ryle argued that "knowing that" and "knowing how" correspond to distinct and mutually irreducible categories, his first interest was in what he called intelligence concepts, e.g. shrewdness, acumen and prudence (Ryle 1949, p.16), which, he argued, belong to the second of these categories, while philosophers under the combined pressures of Platonism and Cartesian dualism have routinely attempted to force them into the first. Contemporary intellectualists are not concerned with intelligence concepts or with the question whether intelligence is a species of propositional knowledge[11]. My own view is that it is not—Ryle's remark that "stupidity is not the same thing, or the same sort of thing, as ignorance" (1949, p.25) strikes me as needing no argument. But Ryle's equation of intelligence with knowing how is mistaken. Intelligence is not equivalent to knowing how to do anything, or with knowing how to do things in a particular way, other than *intelligently*[12]. It is true that some intelligence is a necessary condition of knowing how to do anything, but it is just as much a condition of knowing that something is the case. And it is true that one may show special intelligence in the rapid acquisition and ingenuous exploitation of knowledge *that*, and that this intelligence is not itself reducible to further propositional

[11] The same goes for "practical wisdom", good judgment, professional virtues and craftsmanship, aspects of mind that have been the focus of an incisive discussion by David Wiggins (2012), which nevertheless fails to support the anti-intellectualist stance he endorses there, principally for the reason just indicated: intellectualism is silent on the main topics of Wiggins' discussion.

[12] One might think that this concedes Ryle's point, for his claim was that intelligence amounts to knowing how *to do things in an intelligent manner*. Thus intelligent action is distinguished not by proceeding from contemplation of propositions but by exhibiting a certain method or *modus operandi* :"what distinguishes sensible from silly operations is not their parentage but their procedure."(1949, p. 38). But my point is that Ryle's contention that intelligence is dispositional and manifest directly in action does not entail that intelligence is a kind of practical *knowledge*. Intelligence is no more *knowing how* to do things intelligently than elegance is knowing how to do things elegantly, dexterity is knowing how to do things dexterously or punctuality is knowing how to arrive on time.

knowledge (Ryle 1945-6, p.7, quoted above). But it is not reducible to knowledge how either. One can just as well show intelligence in the rapid acquisition and ingenuous exploitation of knowledge how (e.g. how to play chess or tennis), and this intelligence surely is not a further form of know-how. Varieties of intelligence or intellectual skill are not reducible to knowledge how to do something, as opposed to knowledge that something is the case. The differences between (varieties of) intelligence and propositional knowledge therefore do not bear on assessment of the intellectualist thesis.

Second, as is now widely recognised, knowing how to ϕ is not having the ability to ϕ or being skilled at ϕ-ing[13]. Knowledge how to ϕ, unlike the ability or skill of ϕ-ing, may survive serious physical injuries[14], and knowing how to ϕ, unlike the ability or skill, may be demonstrated in thinking about ϕ-ing, teaching others how to ϕ, appreciating and criticising the ϕ–ing of others, etc. The failure of the equation of knowing how and ability/skill is important, because philosophers often reason as though ϕ-ing were *just* an expression of an agent's knowledge how to ϕ, and therefore imagine that every detail of an agent's ϕ-ing performance must be creditable to their knowledge how to ϕ. And that is really to equate knowing how to ϕ with being able to. In fact, when an agent who knows how to ϕ does ϕ, they draw upon countless other physical and mental capacities, such as daring, stamina, dexterity, grace, sensitivity, prudence, imagination, strength or flexibility. A professional flamenco dancer will demonstrate much more than knowledge on stage: agility, aggression, pride, humour, sensuality, etc. (cf. Wiggins' discussion of craftsmanship

[13] For the widest range of counterexamples to the equation of knowledge how and ability see Snowdon (2003). Others can be found in Stanley and Williamson (2001) Ginet (1975), White (1982), and Moore (1997).

[14] There may be exceptions to this, for example knowledge how to speak English may not survive if the ability to speak English is lost, but this scarcely shows that in general knowledge how can be equated with ability, for in general the two *can* come apart.

and expertise in his 2012). The integrated exercise of these under appropriate circumstances is needed to bring off her performance: they *enable* her to dance in the distinctive way she does. But her knowledge how is not her ability. Not one of these factors nor yet the whole package is her knowledge how to dance. If she twists her ankle, goes on a diet or falls in love her performances may suffer, but she will still know how to dance. The intellectualist thesis, then, is *not* that the package of sensibilities, dispositions etc. that together guarantee someone's general ability to ϕ is reducible to propositional knowledge. It is merely the thesis that *knowledge how to ϕ*, one ingredient in the mix, is propositional knowledge.

The final point is elementary but important. The term 'knowledge', like other cognitive attitude terms, has both an "object" and an "act" or attitude sense[15]. It may be used to refer to the objects or contents of knowledge, i.e., to what is known, such as that the Battle of Hastings took place in 1066, that water boils at 100 degrees centigrade, or that the bishop in chess moves diagonally. In its "act" sense, on the other hand, 'knowledge' refers to the state (/power) of the person who knows these things: to his ("act" of) *knowing*. Accordingly, the suggestion that the construction 'knows how' is typically used to ascribe a distinctive kind of practical knowledge may mean either one of two things. According to the first, the construction is used to ascribe knowledge of a distinctively practical, or non-propositional "object"; according to the second, it ascribes a distinctively practical grasp of what is known, or a distinctively practical way of *knowing*.

[15] Philosophers since Aristotle have distinguished act, power and object uses of cognitive attitude terms (see e.g. Aquinas (1953[1475], Q. 17, 1st reply), where the 'act' use is reserved for occurrent states. Following Alvarez (2010) I transfer the 'act' label to what Aquinas and his contemporaries regarded as the power use: to denote an enduring state or disposition. The same distinction is drawn and emphasised in philosophical logic without appeal to the language of acts or powers: what is believed, judged or stated, i.e. a proposition, is contrasted with a believing, judging or saying of that proposition—what might now be called propositional *attitudes*. See Kneale & Kneale ([1962] 1984) and Prior (1971).

The view rejected by intellectualists is the first.[16] Intellectualists argue that knowledge how does not have a non-propositional object: *what one knows* when one knows how to ϕ is a proposition or propositions. But this leaves open the possibility that distinctive conditions are associated with the possession of knowledge how, that is, with *knowing* how. And indeed a critical, albeit underdeveloped, component of the best known intellectualist account (Stanley and Williamson 2001; Stanley 2011) is its appeal to a "practical mode of presentation" associated with knowledge how—a distinctively practical way of engaging propositional content, postulated to explain why knowledge how is typically expressed in action and not in linguistic statements. Stanley and Williamson say little about what contemplating propositional contents under a practical mode of presentation amounts to, except that it "undoubtedly entails the possession of certain complex dispositions" (Stanley and Williamson 2001, p.429). The core thesis of intellectualism is that the features traditionally appealed to in the individuation of different (contents of) knowledge—contrasts between knowledge-*how* and knowledge-*that*—in fact reflect differences in the way propositional contents are thought about or engaged (differences in mode of presentation). Thus intellectualists do not deny that knowledge-how is distinctive in being typically acquired through training and practice, demonstrated in action, and independent of any ability (on the part of the knower) to put the knowledge into words. What they deny is that a knowledge state with this profile is non-propositional.

The view originally advanced by the arch-critic of intellectualism, Gilbert Ryle, was the second canvassed above, relating to know*ing* how in the "act" sense: he insisted on the

[16] The opposing view that knowledge-how is not a relation to a proposition, but to an activity type, is nowhere in Ryle's work, but it has been defended elsewhere in the literature. See, e.g. Carr 1981, 1979; Hornsby 2011.

distinction between, e.g., knowing *that* the bishop in chess moves diagonally, and knowing *how to* apply that rule in practice (see, e.g., Ryle 1945-6, p.6); "between acknowledging principles in thought and intelligently applying them in action (*ibid*, p.8).[17] But he did not pretend that only the former cases, but not the latter, entail genuine knowledge of the rules or principles in question, i.e. that thinking about or stating a rule requires knowing it, while applying the rule in practice proceeds from a brute capacity, consistent with ignorance of the very rule applied. On the contrary, Ryle's point was just that rules are often learned and applied in practice "without ever hearing or seeing the rule written down at all". For example, he wrote, "we all learned the rules of hunt-the-thimble and hide-and-seek, and the elementary rules of logic and grammar in this way." (1949, p.41). And of the chess student Ryle remarked, "it is not what he does in his head or with his tongue, but what he does on the board that *shows whether he knows the rules* in the executive way of being able to apply them." (*ibid*, my emphasis). So Ryle was distinguishing not between kinds of object or contents of knowledge, but between two ways or modes of *knowing*, one abstract and theoretical ("knowing that"); one "executive" and practical ("knowing how"). If we keep the act/object distinction in view, it is unclear where the substantive tension between Ryle's positive views and contemporary intellectualism lies.[18]

[17] Rumfitt 2003, p.159 and 2011, fn. 2, has also attempted to register, in slightly different terms, how Ryle's concerns ran orthogonal to those of his neo-intellectualist critics.

[18] Stanley (2011, ch 1) makes clear that whether one thinks of knowledge as a disposition or ability instead of as a state is independent of whether one subscribes to intellectualism. I have argued elsewhere (omitted for blind review) that Ryle's profile for knowledge how (see esp. Ryle 1945-6, pp. 8-9) exactly matches John Hyman's (1999, 2015) account of propositional knowledge as the ability to be "guided by" the facts, more precisely to do things for reasons that are facts or truths.

3. Regress and Representationalism

The crux of what is wrong with construing the challenge from versatility set out in §1 as a problem with, or for, intellectualism can be stated quite briefly: knowledge how is not *distinctively* versatile; someone with knowledge how is not *distinctively* equipped for novelty. An agent possessed of *any* kind of knowledge can be expected to see its bearing and draw on it appropriately in an indefinite range of as yet uncontemplated contingencies—that is what we have knowledge *for*. But this is, of course, much too brief and superficial a reply. In this section I will discuss how the challenge from versatility is integrated with Ryle's first purported *reductio* of intellectualism, and show that the proper target is a representationalist conception of intentionality in general. To do this I shall first show how Ryle's regress argument is insensitive to a distinction between knowing how and knowing that, and refer to a parallel, explicitly anti-representationalist variant of the regress argument in Searle (1991).

What Ryle called 'intellectualism' is not merely the thesis that knowing how is a species of propositional knowledge; it is this thesis interpreted against a background view of intentional states (e.g. imagining, believing, thinking) as defined by the permanent presence or momentary occurrence of a representation in or before the mind, so that the realisations of these states in action are construed as practical operations piloted by a cognitive operation on the representation. The doctrine Ryle was attacking combined the ideas that intelligent behaviour (by Ryle's lights, exercises of knowledge-how) is behaviour that is piloted by cognition, and that cognition is a matter of contemplating propositions (Ryle 1945-6, p.1-2). Hence the idea that intelligent behaviour demonstrates cognition of a proposition was elided with the idea that it is piloted by a "shadow act" of contemplating a

proposition (*ibid*, pp. 2-4), and what is in reality an argument against the second of these ideas was mistaken for an argument against the first.

Critics have made much of Ryle's apparently gratuitous attribution to the intellectualist of the idea that "shadow acts" of propositional contemplation precede uses of know-how and, by disclaiming this commitment, dismantled his argument that regress of these shadow acts follows inexorably from the intellectualist view (Stanley & Williamson 2001, pp.413-416; Stanley 2011, pp.12-14). But the crux of Ryle's regress arguments is ignored rather than neutralised by this response. It is that propositions (conceived of as objects of thought) are not self-applying: it takes intelligence to apply them, and that intelligence is not reducible, on pain of vicious regress, to the mental containment of further propositions. As I shall explain in more detail shortly, Ryle's problematic regress is of propositions, which are impotent to generate any action except in the hands of human intelligence, not of mental acts of contemplation.

According to the conception of mind that Ryle attacked, propositional knowledge consists in having facts and truths registered somewhere accessible to the mind—stored in some mental archive or database. If we assume this conception, and that knowing how to do things is at bottom knowledge of facts and truths, it seems that putting that knowledge how into practice must mean bringing before the mind the facts and truths that are relevant and considering them before or as we act. If we wish to bake a cake, say, we must access the propositions about cake-baking and consult them (perhaps lock onto them in unconscious thought)—the mental analogue of pulling a cookery book off the kitchen shelf, finding the right recipe and reading the instructions. A sophisticated action or activity is now cast as the execution of two tandem operations: on the one hand, engaging propositional

knowledge (finding the right file, and considering the pertinent cake-baking facts registered there) and on the other hand, subsequently or concomitantly acting.

This version of intellectualism, one can agree with Ryle, is absurd. But what is absurd is the way it models the connection between knowledge and action—knowledge in general, not merely knowing *how*.[19]

Consider knowing how first. Ryle observed that any given fact, maxim, or rule which a practical performance may be said to draw upon will: first, not *itself* indicate that it is to be drawn on in the present action or activity, in preference to indefinitely many others one may have at one's disposal—which is to say intelligence is required *to* draw upon the fact, rule or maxim in question at all (Ryle 1949, p.19); second, it will not indicate how an action which draws on this knowledge is to be performed, or in other words, how the fact, maxim or rule is practically applied even where it pertains to one-pattern procedures (e.g. tying clove-hitches) (Ryle 1945-6, p.8); third, though this is no more than an expansion of the second point, being an item of knowledge which one must be able to draw upon repeatedly in indefinitely diverse situations, it must be general in nature, so it will not contain specifications as to how it must be applied in *this* case (Ryle 1949, p.20).

Pointing to the presence of some propositions underlying an intelligent performance is, therefore, insufficient to explain its author's knowledge how to perform. Those propositions do not have a list of clauses explaining how and when they are to be applied in every possible circumstance tucked into their contents (this is one way to put the challenge from versatility). So the reduction of knowing how to propositional knowledge, if

[19] For a discussion of the double standard at work in Ryle's arguments see Stanley 2011, ch. 1.

that is construed as having propositions registered somewhere in the mind, is plausible only by taking for granted that same variety of knowledge or intelligence—knowledge *how to* use or apply information—under whose control alone productive use can be made of the propositions in question. But this cannot be taken for granted, for people sometimes do produce facts when they are irrelevant, assent to rules but fail to follow them in practice, and understand technical principles in general but fail to see how they may be applied to the practical problem at hand (Ryle, e.g. p.1945-6, pp.2-6). Having a stated fact or truth before one is not yet to understand it, know when it is pertinent or how it is to be applied in a given practical context. The representationalist intellectualist of the sort I have been discussing doesn't have the resources to explain the transition from mental containment of a proposition to practical competence exploiting it, except through the postulation of higher order regulative propositions of which one must be apprised, and this leads in short order to vicious regress. For the knowledge how to apply propositions must be explained as possession of further propositions in one's mental archive. But if the bare presence of the first tier of propositions in the mind can't ensure their own productive use, the same must be true of the second tier, so that we must go on postulating tier behind tier of regulative propositions *ad infinitum*, each called in as a buttress for the set before (Ryle 1945-6, pp.6-7; 1949, pp.19-20).

Ryle's argument shows the quagmire we get into if we imagine that the conscious or unconscious introspection of a representation is what makes sophisticated action possible. But it does not show that knowing how is not propositional knowledge, because propositional knowledge cannot be understood as the bare containment of a representation in the mind *either*. For one does not *know* a fact or truth unless one can draw on that fact or truth in thought, speech or action (Ryle 1945-6, p.16). And so the regress argument will

apply equally to the representationalist conception of knowing *that*. For example, if I know that *p* I will normally draw on that knowledge when I answer the question whether *p*, or the question whether *q*, given that if *p* then *q*, and every time that I ϕ *because* (or for the reason that) *p* . But when I am asked whether *p*, what ensures that I will consult the file in which '*p*' is recorded, instead of any of the other files? And answering a question is also a practical application of knowledge that *p* which, it now seems, I must also know how to make. There is no way of limiting in advance the range of possible contexts in which my knowledge that *p* might serve me, e.g., of circumscribing the range of possible questions it could be called upon to address—they might be infinite! On the representationalist picture, mere knowledge that *p* will not enable me to cope with any of these; without the buttressing of a further intermediary between knowledge and action, the mere knowledge that *p* must be *entirely* idle. And if we conceive of what is needed to put knowledge that *p* into action as possession of some further propositions stored in the mind, we shall need an infinite number of them before we can answer a single question (cf. Ryle1945-6, p.2).

Ryle's regress arguments parallel Wittgenstein's less systematic invocations of vicious regress, in discussions of other intentional states such as remembering, intending and understanding in the *Philosophical Investigations* and the *Blue Book*[20]. I will briefly revisit his notorious "rule following considerations" in the next section. The Rylean regress is also echoed by Searle's (1991) arguments for the dispositional "background" of intentional states. The proper target of these arguments is not, and is not tied to, the thesis that knowledge how to do things is a species of propositional knowledge. Here is Searle's argument for comparison (p.152):

[20] See the remarks on understanding an order in the *Blue Book*, p.3 and *Philosophical Investigations* §451, §239 and surrounding discussions; §186-201; §213.

> Suppose that all Intentionalistic mental life and all cognitive capacities could be entirely reduced to representations: beliefs, desires, internalised rules, knowledge that such and such is the case, etc. Each of these representations would be expressible as an explicit semantic content ... But the semantic contents with which the conception provides us are not self-applying. Even given the semantic contents we still have to know what to *do* with them, how to apply them, and that knowledge cannot consist in further semantic contents without infinite regress.

The target is representationalism. Searle's idea is that there must be a dispositional aspect of mind, for representations are idle by themselves, and a mind exploits information—a mind is, for example able to *understand*, *appreciate the bearing of and apply* information—it does not merely contain it. Searle's regress argument, like Ryle's, really shows that no mental attitude to *p* is reducible to storing '*p*' in a mental notebook (cf. Kenny 1989, p.108; Stroud 1991). Thus the cognitive *attitude* that knowing is cannot be reduced to the presence to the mind of its *object* or contents.

Knowing is, as Ryle might say, itself an intelligent capacity, and so it is not the case that propositional knowledge is idle unless buttressed by some further practical intelligence: rather, practical intelligence makes the difference between merely containing information (as en encyclopaedia might) and *knowing it*.[21] As Jason Stanley has urged, the way to answer Ryle's anti-intellectualism is to agree with him that knowing how has a dispositional or practical aspect, but deny that this distinguishes it from knowledge *that*—deny that propositional knowledge is "behaviourally inert"(Stanley 2011, p.27). This is a point that Ryle surely would have been obliged to accept, for he argued in *The Concept of Mind* that knowledge-*that* is a multi-track ability, i.e. an ability with indefinitely heterogeneous manifestations (Ryle 1949 p.117).

[21] See again Kenny 1989, pp.108-9; Hyman 2015, 183-184.

Of course, in a population of agents who know that *p* there will be variations along multiple dimensions in their responsiveness to *p*, i.e. in what they can do (understand, theorise about, achieve in practice) on the basis of *p*[22]. But the Rylean argument, properly understood, rules out the possibility that any one of these agents can do *nothing* on the grounds that *p*. If a fact is in principle inert in an agent's mental and practical life, he cannot qualify as knowing it. So if knowing by itself equips one for action, how much action and of what kind? I don't believe there is any single, non-arbitrary answer; what an agent can do with what he knows depends in complex ways on, *inter alia*, the rest of his cognitive resources and on his circumstances. True, when an agent knows a rule or principle, it may be that the central or standard case of knowledge entails a more or less circumscribed range of dispositions, while when he knows a fact of colonial history, the central or standard case of knowledge will entail dispositions of a different kind. But differences in the kind of grasp that agents can have will be present whether what they know is a fact of chemistry, a rule of chess, a method of farming or an algebraic principle; whether what they know is *who* Aristotle was, *how* to navigate, *when* to add the flour to the cake mixture or *that* water boils at 100 degrees centigrade. Differences in grasp appear in connection with knowledge *how* to do something as well as knowledge *that* something is the case—it is not know-how itself that makes the difference.

In sum, if more careful attention to the Rylean regress arguments leads us to acknowledge that knowing a fact *by itself* entails being in a position to draw on that fact in practice—in

[22] In fact, a single agent may not be constant in these respects across contexts. For example, A may know that *p* and be capable of addressing the question whether *p* in casual conversation, but go blank in an exam, or when the question is posed in an unfamiliar guise. Or A may know that *p*, *q*, *r*…but not know that she knows these propositions—she discovers when examined that she has the knowledge after all; or A may disclaim the belief that *p*, but behave as if she believed that *p* (e.g. the implicit sexist). Cases of this sort have moved some philosophers to conjecture that our cognitive systems are divided or fragmented, preventing us from "putting two and two together". See Egan (2008), Lewis (1982), Stalnaker (1984; 1991).

connection with at least *some* range of queries and tasks—and we notice that the ways a fact may be practically drawn upon are multifarious and unpredictable in advance, then the argument from the indefinitely heterogeneous applications of knowledge how to the absurdity of construing it as propositional cannot convince. All knowledge is indefinitely versatile[23].

I do not pretend that there is nothing at all puzzling about this. But the puzzle is independent of intellectualism. I want now to go briefly over how this feature of knowledge was brought to light and problematized by Wittgenstein in his best-known discussion of rule-following in the *Philosophical Investigations*, to bring out more clearly how the cogency of the challenge in §1 tacitly relies on representationalist notions, and support the idea that knowledge how to ϕ is knowledge of the general rules and principles of ϕ-ing.

4. The Application of a Rule: From the General to the Particular

In Wittgenstein's discussion of rule-following the two concerns—representationalism as a target and indefiniteness as a problem—interact. Ostensibly the central puzzle is how knowledge of a general rule or formula could equip one in advance to act appropriately (in conformity with the rule) in an open and indefinite future. Wittgenstein then tries and fails responses recommended by a naïve representationalism, finding they lead inexorably to vicious regress. These regress arguments parallel Ryle's and Searle's; indeed, it seems probable that they provide the blueprint for them. But first, intellectualism is not at issue in these arguments; and second, their apparent upshot is that knowing a rule and being in a

[23] Indeed, Gareth Evans (1981) famously argued that we can *only* credit an an agent with belief if this versatility is present; if a belief can combine with the agent's desires in the service of indefinitely varied projects.

position to apply it in a potential infinity of instances cannot (on pain of vicious regress) be regarded as distinct achievements. This result is congenial to intellectualism, suggesting that where discrepancies arise between what two agents can do, understand or conclude on the basis of a known rule or principle, this is not to be understood as a discrepancy in their knowledge but as a discrepancy, as it were, in in the way they know it—that for which Stanley and Williamson have invoked the notion of a practical mode of presentation.

It is difficult to ignore Wittgenstein's special preoccupation with the connection between knowledge of a rule and the rule's *indefinite* applications. Given a rule for a series of numbers (Wittgenstein's central example is '+2'), it is fixed in advance how the series should continue up to infinity. If one knows the rule, one has knowledge that fixes every legitimate application of it before the application is made. How is this possible? Must all the applications be somehow present in one's grasp of the rule now? Must understanding the rule involve running through these applications to infinity in thought? (§§184;186-8). The rule itself makes no mention of its *nth* application, so seemingly knowledge of the rule is a deficient resource when one comes to apply it. One must in addition know what to do at the *nth* stage of the series. Must one then not only know the rule, but also, for each of its applications, a proposition that tells one what to do at *that* point (§187)—that is, infinite propositions that cover all the future applications in advance? In the absence of these propositions, one would surely be without grounds to apply the rule one way or another at each turn; nothing could guarantee that the next application will be correct.

This is much the same difficulty critics of intellectualism seize upon in an attempts to show knowledge how cannot be propositional. How could knowledge of general rules of warfare or flamenco or perspective drawing carry one through the specific circumstances of the

task or performance one must address oneself to *now*? If all the knowledge how to perform is propositional, won't one need to know some further propositions, which tell one how to connect the general rules with the specific situation at hand? And this lead to the absurd conclusion that it would be necessary to know infinite context-specific propositions concerning eventualities one does not yet anticipate. [24]

Wittgenstein did not conclude that in addition to knowing the general rule or formula, one must also know how to apply the rule at *this* point, and also the next, *ad infinitum*. Like Ryle, he pointed out that postulation of the infinite propositions supposed to guide one at each and every stage of application does not help, for even if the applications of the rule were laid out by an infinite series of propositions, one might not know how to follow *them* (§§218-219). Wittgenstein considered a variety of candidate go-betweens which might take one from the *general* rule or formula to the *particular* application: a derivation rule; a specific proposition ('after 1008, 1010'); an interpretation; an intuition. But none of the intermediaries introduced to bridge the gap between knowing the rule and knowing how to apply it at the *n*th stage of a series help, because if one could doubt how to follow the general rule in that case one could also doubt how to follow the go-between instruction. And so the appeal to intermediaries is rejected by appeal to vicious regress in §201 (the intermediary here is an *interpretation*):

> It can be seen that there is a misunderstanding here from the mere fact that in the course of our argument we give one interpretation after another; as if each one contented us at least for a moment, until we thought of yet another standing behind it.

[24] The analogy suggests itself strongly in a critical discussion of Crispin Wright's (2007 pp.490-497) of what he calls the 'modus ponens model' of rule following. On the modus ponens model an agent always combines the rule known with additional contextual information when they come to apply it, so rule following is not "pure". Wright concludes that this model cannot be coherent in general, for in general the putative "extra" information cannot be characterised except by reference to the rule together with which it is applied.

According to one standard and, I think, correct interpretation of Wittgenstein's proposed solution to the rule following problem, the way out of this quandary is to reject the idea that knowing how to apply the rule in a potential infinity of instances is a separate competence from knowing the rule[25]. Thus §201 continues[26]:

> What this shews is that there is a way of grasping a rule which is *not* an *interpretation*, but which is exhibited in what we call "obeying the rule" and "going against it" in actual cases.

To insist, on the contrary, that in addition to knowing the general rule we must know how to apply it in each particular case, casts knowledge how in the role of intermediary: it is assigned the brute task of carrying one from apprehension of the general rule or formula to producing the particular application that agrees with it. But this is to sanction the representationalist picture Wittgenstein apparently set out to upset. It means accepting the disassociation of cognitive apprehension of the rule, and making some particular application ("in practice") that agrees with the rule: and now what bridges the gap is knowing how. The tacit dependence on representationalism here may be seen from the fact that, once we have subtracted the knowledge how to apply the rule at every point, there is really nothing left for 'the purely cognitive apprehension of the rule' to refer to, but the presence of a rule formulation in a mental pocket.

The "solution" to this puzzle about knowledge of rules appears in slogan form in Ryle's own work. "Knowing a rule," he wrote, "is knowing how. It is realised in performances which conform to the rule, not in theoretical citations of it." (Ryle 1945-6, p. 7).

---

[25] This is the interpretative line advanced by Gordon Baker and Peter Hacker in their *Analytical Commentary* (Baker and Hacker 2009), and endorsed by numerous commentators including Severin Shroeder (2006), Hans-Johan Glock, and, with qualifications, Crispin Wright (2001, 2007), although Wright is emphatic that Wittgenstein's response is not enough dissipate the perplexities raised by his rule-following discussion.
[26] Cf. *Zettel* §§300-306.

This averts the threat of a vicious regress, though as a response to Wittgenstein's puzzle, of course, it cannot yet be satisfying. On the epistemological side, being told that to genuinely know a rule *just is* to be able to apply it correctly in a potential infinity of instances does nothing dissolve puzzlement at how an ability of this sort could be acquired. How can knowledge (of a rule) one acquires now entail constraints on one's actions in some unforeseen future eventuality? On the metaphysical side, the puzzle remains how an object of knowledge of this sort could be constituted: how should the metaphors of containment (of infinite applications within the rule) and determination (of infinite applications by the rule) be unpacked? And however one looks at it, the salient puzzle about what could rule out a "deviant" interpretation, cannot, in light of the under-determination of a rule by finite instances, be resolved by appeal to an "internal relation" between the rule for an infinite series and its applications in practice *up to this point* (the point emphasised by Kripke 1982) [27]. But puzzling or not, this is what knowledge of a rule is like. It has content we can put into a finite form of words, and its applications extend in an open and indefinite future. The puzzles are independent of intellectualism.

5. Ways to ϕ and the Generic Nature of Knowledge How

In previous sections, I have taken as the target intellectualist thesis the view that knowledge how to ϕ is knowledge of the general rules, principles, methods or procedures that characterise ϕ-ing itself, and the anti-intellectualist contention I have focused on debunking is that (practical) knowledge of *general* rules and principles is a deficient guide for action in any *specific* circumstances. But the challenge of §1 can be brought from the other direction:

[27] The nature of the difficulties and the inadequacy of existing responses is brought out most clearly in Crispin Wright's discussions; see Wright (2001, Parts I-II) and (2007).

it can be argued that propositional knowledge that *is* sufficient to guide performance in any one specific circumstance will not help one know how to act when the circumstances change, thus knowledge of sufficiently circumstance-specific propositions could not constitute an agent's general knowledge how to ϕ. A *prima facie* more appropriate target for this kind of argument is Stanley and Williamson's (2001) intellectualist proposal; it is their view, and its defence from an argument along these lines advanced by Jenifer Hornsby, that I turn to now.

Although they remain a minority, philosophers of diverse persuasions have rejected Ryle's distinction between "knowing how" and "knowing that", and argued that the propositional content of knowledge how is the proposition or set of propositions that answer a question about how to ϕ. Thus Alan White (1982) and Adrian Moore (1997), both of whom have been prepared to go along with Ryle in thinking of knowledge in general as a complex ability or disposition, argue that 'knows how' sometimes simply ascribes propositional knowledge by characterising that knowledge in terms of a question about what to do.[28] But the most sophisticated arguments for this view are due to Stanley and Williamson (2001), elaborated somewhat in later work by Stanley (2009, 2011). Their proposal involves developing Kartunnen's (1977) semantics for embedded questions to provide an analysis of knowledge ascriptions using the interrogative 'how', which yields the following schema, where ways are properties of token events of ϕ–ing, to be picked out demonstratively:

[28] Moore (1997, ch.8) argues for a distinction between practical and propositional knowledge, but this cuts across the knowing how/knowing that distinction. Some states of knowing how but also some states of knowing that (e.g. knowledge of necessary

A knows how to ϕ iff A knows a proposition of the form '*w* is a way for A to ϕ'. [29]

Stanley and Williamson's linguistic *argument* for their view has been challenged on empirical as well as philosophical grounds[30], but I shall be concerned only with assessing the view itself, and only in relation to the objection from versatility. I hope that the defence of Stanley and Williamson's view will make clear why *any* intellectualist proposal can withstand that objection; as far as that challenge goes, it makes little difference what is postulated as the proposition(s) someone who knows how to ϕ knows, as long as the propositional content is understood as being sufficiently coarse-grained as to cover the indefinite and varied manifestations of knowledge how to ϕ.

On first consideration, it may seem impossible to interpret the propositional content Stanley and Williamson postulate as inherently general in this way. For in most cases it is a single proposition involving a way to ϕ *such as can be exhibited by a single token event of ϕ-ing*. The knowledge they postulate might seem therefore too tailored to the specific details of just one of a potential infinity of ϕ-ing occasions; accordingly their proposal may seem hopelessly unable to do justice to the essentially general nature of knowledge how to ϕ. Hornsby (2011) argues precisely this: "what is known by a person who knows how to ϕ needs to be somehow *generic,* and that is why it cannot be captured by citing particular

[29] As mentioned in §2, the idea that some ascriptions of knowledge how state or imply that the *way to ϕ* is engaged under a "practical mode of presentation" or "practical way of thinking", entailing dispositions to act, is a powerful component of the account, but not relevant to the present discussion. See Stanley and Williamson (2001), pp. 428-30 and Stanley 2011, ch. 4.

[30] Some have challenged the methodology of drawing conclusions about the metaphysics of know-how from linguistic analyses (See Noë 2005, Devitt 2011); others question the empirical support for the linguistic analysis itself, given that the use of an interrogative for the ascription of know-how is a notoriously adventitious feature of English (see Wiggins 2009 and especially Rumfitt 2003). Stanley (2011, ch.6) attempts to respond to these objections; for a reply to that attempt see Brown (2013).

instances of the person's ϕ-ing." To illustrate, she gives the following example (*ibid*, pp.91-2):

> Consider Clare, an excellent gardener, who is pruning the roses. As she cuts with the secateurs, you say, 'That's a way for Clare to prune roses.' Then, as she is examining a bit of a plant in order to determine where to cut next, you say again, 'That's a way for Clare to prune roses.' The ways you denote with your successive 'that's' evidently have very little in common. When you first say 'that,' Clare's hand is in motion; when you say 'that' a second time, her brow is furrowed as she contemplates the next step. But each of these samples of rose pruning must be supposed to serve equally well to show a way to prune roses that is known by Clare.
>
> ….If the demonstratives…behaved as they would need to in order to fit an account of what it is to know how to do something, then they would go hand in hand with ways of ϕ-ing such that their being known by an agent to be ways for her to ϕ amount to her knowing how to ϕ. But in [this] example …only one piece of knowing how is exercised (Clare's knowing how to prune roses), but it is exercised in two evidently different ways (in moving her hand, in determining what to do next).

Hornsby's contention is that Clare's general knowledge how is what *affords* her knowledge "of such propositions as enable her to see her way through the contingencies that she needs to negotiate to complete the task", but that these propositions, pertaining to the particularities of the one occasion, are not available to her in advance. (*ibid*, pp. 94-5).

Considering Hornsby's example of the rose pruner, one might think it uncharitable to assume that such slender time slices of Clare's rose-pruning are the token events that Stanley and Williamson have in mind as apt to exhibit a way to prune, especially as in some of these Clare is lost in thought and apparently not engaged in pruning at all. Must Stanley and Williamson's proposal require that a single chess move exhibit the way to play constituting a chess-master's know-how, or that a pianist's hitting a single key should demonstrate a way of playing the piano, such that her knowing that *that way* is a way for her to play, exhausts her knowledge how to play the piano? Any proposal making use of the device of deferred ostension must take for granted a primitive notion of contextual

salience, and the contextually salient event of rose-pruning when Clare is gestured at and we are told, 'that's a way to prune', is not, I think, the event of her cutting a single stem, or the event of her crouching beside the bush with secateurs in hand. But nothing Stanley and Williamson say rules out this possibility, and attempts to block it through general principles would only betray the need for independent criteria for distinguishing manifestations of knowledge how to ϕ, suggesting that, in an account of knowledge how to ϕ, *ways* to ϕ do no explanatory work.

Moreover, if the thrust of Hornsby's argument is right, appeals to contextual salience will not help. The difficulty Hornsby is pointing to is supposed to be one of principle, which will bite *whatever* stretch of an agent's ϕ-ing one selects as the event that exhibits a way to ϕ, for the reasons set out in §1: the ways of ϕ-ing that someone with knowledge *how to ϕ* is apt to demonstrate are indefinite and indefinitely heterogeneous; they must differ substantially from occasion to occasion (consider how different two flamenco performances by the same dancer might be). Any given event of ϕ-ing only reveals a more or less *particular* expression of the agent's knowledge how; it cannot fully reveal the *general* source of those particular manifestations (*ibid*, p.93).

One way of being drawn back into the frame of mind in which Hornsby's critique will seem persuasive is to focus on what kind of practical performance would constitute decisive *evidence* that an agent knows how to ϕ. Ryle observed that it is often impossible to distinguish, from a single observation of an agent's ϕ-ing, a demonstration of genuine knowledge how from a fluke performance (Ryle 1949, p.42). Hornsby turns a similar point against Stanley and Williamson when she notes that a token-event of typing the word

'Afghanistan' does not reveal a difference between an individual exercising their general knowledge how to type, and an individual exercising their more limited or specific knowledge *how to type the word 'Afghanistan*'. (Hornsby 2011, p.91). Like the example of the rose-pruner, this example is supposed illustrate why "ways in which [an individual] ϕ-s could not be constitutive of [their] knowledge how to ϕ." (*ibid*). But it does not really illustrate this, any more than Ryle's acknowledgement that *we may not be able to distinguish* a fluke performance from an exercise of knowledge how can be taken to show that those performances do not *in fact* differ in their procedure (Ryle certainly never drew any such conclusion). We may not be able to tell, for example, whether a single French utterance at a bar is a use of knowledge how to speak French or merely of knowledge how to ask for a beer in French (acquired minutes earlier by consulting a phrase book), but there is a fact of the matter as to which it is, and accordingly, following Stanley and Williamson's approach, a fact of the matter as to which *way* or *ways* are being exhibited.

To assume that the *ways to ϕ* exhibited in each pair of cases *do not differ*, while stipulating that the *knowledge* drawn upon in them *does differ*, would beg the question against Stanley and Williamson. If we do not beg it, then all that such examples show is that it may not be transparent to the onlooker, from a single observation of A's ϕ-ing, whether A is exhibiting a way to ϕ (e.g. a way to type in general; a way to speak French) or a way to φ (e.g. a way to type the word 'Afghanistan'; a way to ask for a beer in French).

There is one seemingly awkward consequence of this line of defence, relating to the so-called paradox of the progressive[31]. If ways to ϕ can be exhibited at any instant of an agent's ϕ-ing (though an onlooker might not cotton on to the fact that such a way was

[31] Thanks to Jennifer Hornsby for the point and example discussed here.

being exhibited), then it seems that Mary can exhibit a way to make an omelette when she is breaking eggs at T1, even if at T2 she is interrupted, or abandons her intention to make an omelette and goes on to make scrambled eggs instead. At T1, that is, there is supposedly a token-event of omelette making exhibiting a way to make an omelette, even though no omelette was ever made. Likewise, if Stanley and Williamson's official view is that Mozart can, over a period of two minutes at the piano, say, exhibit a way of composing a symphony (cf. Stanley 2009, pp. X), then surely Mozart exhibits that way even if no symphony is ever composed, because he looses interest and decides to study a fugue instead. This might seem counter-intuitive—it is one consequence of allowing that the exhibition of a way to ϕ need not be communicative, i.e. it need not make immediately evident to an onlooker that such a way *is being exhibited*. But however unintuitive, I think this is entirely correct. For at T1 Mary is, as far as she is concerned (and surely she knows best), making an omelette, and drawing on her knowledge how to make an omelette; likewise, before Mozart abandons his work he is exercising and displaying his knowledge how to compose a symphony, even if in the end nothing comes of it.[32]

Consideration of *what an onlooker can discern* from observing a single token-event of ϕ-ing cannot, then, decide the assessment of Stanley and Williamson's contentions as to *what is exhibited* by a token-event of ϕ-ing. But Hornsby's objection is not just epistemological; her contention is not just that we can't *see* Clare the rose-pruner, at a given instant, as exhibiting a way to prune which is plausibly constitutive of her knowledge how to prune. It

[32] Other responses are possible, if one wants to deny that there is any token event of ϕ-ing suitable to exhibit the corresponding way in cases where an agent's ϕ-ing is incomplete (e.g. interrupted omelette making). In the omelette-making case, one could appeal to contextual salience and argue that the contextually salient event of omelette-making must be one that incorporates the completed procedure, so an utterance at T1 of 'that's a way to make an omelette' would just be false. But this response would fail in the symphony case, for it might take months even for Mozart to compose a symphony, and it is implausible that the contextually salient event of symphony-composition comprises months of scattered activity. So alternatively, one might argue that whether an agent is at T1 exhibiting a way to ϕ depends in part on what happens in the future.

is that no such way to prune is being exhibited. Hornsby's critical assumption is just that individual ways of ϕ-ing are much too finely individuated to constitute someone's general knowledge how to ϕ—e.g. the ways of pruning the roses that Clare exhibits from moment to moment and occasion to occasion "have very little in common", as Hornsby puts it. But Stanley and Williamson are at liberty to deny this; that is, to deny that ways to ϕ are any more (or less) finely individuated than the knowledge states whose exercise exhibits them.

The absence of comment on the nature and individuation of ways to ϕ is no small lacuna in Stanley and Williamson's account, but they do rule out the idea that ways are *tropes*: "Since ways are properties rather than particulars…demonstrative reference [to a way] involves deferred ostension to a property. That is, what [one] actually ostends is an instance of a way of [ϕ-ing], and [the] demonstrative expression refers to a property of which that is an instance." (Stanley and Williamson 2001, p. 428, fn. 29). Their comparison, instructively, is with demonstrative reference to colours. (*ibid*). It is true that a single instant of an agent's ϕ-ing may not suffice for the onlooker to cotton on to the fact that it is a way to ϕ being exhibited, not something more particular or determinate, but it is equally true that a single sample of blue fabric may not suffice for a child to cotton onto its being *blue* that is being exhibited, and not only, say, a determinate shade of indigo. But this does not tell against an explication of someone's knowledge of what it is for an object to be blue by demonstrative reference to a single blue sample.

Even without reflecting on their analogy with colour, it is natural to suppose that Stanley and Williamson's ways to ϕ will be much less finely individuated than Hornsby assumes. For the most natural interpretation of the expression 'way to ϕ' in this context, I suggest, is as denoting a method, procedure, technique or strategy for ϕ-ing—i.e. a *means* of one sort

or other for the achievement of ϕ-ing. These are inherently general in application; accordingly, they are coarsely individuated. A single procedure for tying a clove-hitch or for stringing a bow may be applied with slight variations from case to case depending on the materials used, the skills and style of the agent, and perhaps other factors. The nature of the procedure tolerates these variations in application: we say the *same* procedure is applied across a range of cases. Still any one of a multitude of, in certain respects, different events of bow-stringing (say) would do as an illustration of the one procedure for stringing a bow—i.e. would be suitable to exhibit a single *way* to string a bow.

Interpreting Stanley and Williamson's *ways* as methods or procedures thus renders Hornsby's contention that Clare's ways of pruning at different points in time "have very little in common" incredible. On the contrary, it is easy to imagine that Clare exhibits the very same way of pruning—applies the same procedure or technique—*whenever* she prunes[33]. This interpretation of 'way' can, of course, only be regarded as a heuristic, for it is not plausible than in every case knowledge how to ϕ (by Stanley and Williamson's lights, knowledge of a way to ϕ) is knowledge of a means or method for ϕ-ing. But it allows us to see that instances of ϕ-ing are capable of demonstrating a sufficiently coarsely individuated property as to give content to someone's general knowledge how to ϕ, and thus to show that Hornsby's challenge can be answered.

[33] At least, whenever she prunes that type of bush. It might be that Clare would have applied a different procedure, for example, had she been pruning climbing roses instead of small bushes. But this level of heterogeneity in procedure is easily accommodated by Stanley and Williamson's proposal; Stanley suggests that experts in ϕ-ing will know several ways to ϕ, not only one. The idea that Clare knows and exhibits three or four different ways to prune is consistent with ways to prune being sufficiently coarse-grained as to be exhibited by an indefinite multiplicity of token-events of rose-pruning.

A second line of response is in fact available to Stanley and Williamson[34], one proceeding from the premise that ways to ϕ are just as finely individuated as Hornsby assumes, but this would require a substantial revision of their semantic analysis of ascriptions of knowledge how.

It is not implausible that like many other verbs, 'know' has a habitual aspect or reading, so that in some cases, ascriptions of propositional knowledge should be understood to register the agent's *knowing* of a certain proposition-type as habitually occurring—or more precisely, reliably occurring in relevant circumstances[35]. For example, one might think that the statement 'Watson knows when to take his revolver' ascribes propositional knowledge to Watson, but not standing knowledge of any one proposition, the proposition that he should take his revolver at such-and-such a time (or under such-and-such circumstances). Instead it may be that we should understand this sentence as stating that Watson "habitually" knows (can be relied upon to come to know), in circumstances that warrant his taking a revolver, that this is an occasion on which he should take his revolver. Similarly, 'Holmes knows whether a client is lying,' perhaps means that Holmes reliably knows, when he is faced with a client who is lying, that the client is lying; and when presented with a client who is telling the truth, that they are telling the truth. He doesn't, of course, know in advance of meeting them whether specific clients are lying, but when faced with a lie he detects it immediately. It may well be that in the right circumstances Watson

[34] This move was suggested to me by Tim Williamson, without endorsing it, in discussion in 2013. My impression is that the idea implicitly informs Stanley and Williamson's (2016) paper on skill, but the revision it implies of their original semantic analysis is not discussed there.

[35] Drawing an analogy with particularism, Jason Stanley (2011, pp.181-2) considers an idea that dovetails with this linguistic analysis: the idea that skilled action might be action that proceeds from reasons, even though the reasons are *not available to the agent in advance of the context of action*. Stanley does not opt for this position in discussing the problem of novelty, however. His ultimate stance is not altogether perspicuous, but at least one prong of his response coincides with the argument of this paper: the idea that all knowledge states either *are* or are associated with multi-track dispositions, and so will be manifest in indefinitely heterogeneous conditions (*ibid,* p.183-185).

and Holmes come to possess these items of propositional knowledge in virtue of exercising an indefinite range of cultivated skills and sensibilities, but the expressions 'Watson knows when to…', and 'Holmes knows whether to…' simply ascribe *propositional* knowledge that Watson and Holmes habitually possess.

There are certainly ascriptions of knowledge *how* that also invite this type of reading, for example, 'Holmes knows how to mollify Mrs. Hudson'. This might be understood as stating that Holmes knows, whenever Mrs Hudson is at the end of her tether, of some way *w* that *w* is a way to mollify her. If such a reading *is* appropriate for some, or perhaps all ascriptions of the form 'A knows how to ϕ', then Stanley and Williamson can agree with Hornsby that ways to prune a rose bush are finely individuated, highly context-specific properties, and that indeed the rose-pruner does *not* come armed to the task of pruning with knowledge of those ways. Instead, they could now argue, the rose pruner *comes to know* propositions concerning the specific ways to prune on each occasion of rose-pruning[36]. Still the knowledge will be propositional. 'Clare knows how to prune roses' still ascribes propositional knowledge; only it is not standing propositional knowledge of a *general* way to prune, but rather "habitual" knowledge of *specific* ways.

Appeal to this analysis might shift the point of contention between Stanley and Williamson and their critics. For both sides could then allow (a) that someone with knowledge how to ϕ comes to know, in each ϕ-ing context, propositions concerning what do and how in order to ϕ *only in that context* and (b) that such a person could be relied upon to attain this

[36] Hornsby might reject even this. According to Hornsby, the propositions that Clare comes to know when faced with a specific bush do not relate to ways of pruning it, but to such matters as the variety and dimensions of the rose bush to be pruned, her general knowledge how then equips her to respond to. In other words, the agent's knowledge *what to do* in response to new information is non-propositional.

propositional knowledge when needed in virtue of some background complex of dispositions, skills, interests and attitudes peculiar to the practice or profession of ϕ-ing. The disagreement would concern whether this latter dispositional complex itself deserves the honorific 'knowledge.' –Hornsby, Wiggins and others on the Rylean side would, I imagine, want to say yes; Stanley and Williamson would say no. No, for that which sustains an agent's coming to possess some propositional knowledge is no part of the knowledge itself. On a habitual interpretation of 'knows how', 'Holmes knows how to mollify Mrs Hudson' will be true just if in each (or enough) of the relevant circumstances, Holmes possesses knowledge that some way *w* is a way for him to mollify her; Holmes' background skills and dispositions will be irrelevant to the truth-conditional content of the ascription which, by Stanley and Williamson's lights, is what we should be interested in.

Conclusion

Against the critique of §1, I have wanted principally to show three things. First, it is doubtful whether knowledge of a finitely statable set of procedures, rules, principles or methods, is by itself a deficient guide for their practical application in a potential infinity of unforeseen contingencies. But second, if such a practical deficiency in the face of indefinite novelty *does* attach to knowledge of this sort, we should expect an agent to be similarly unequipped for action in unforeseen contingencies when what he has to guide him is, say, scientific or historical knowledge, and the kind of actions he is called upon to undertake in light of what he knows are principally linguistic. There can be no principled circumscription in advance of the range of possible tasks and questions that propositional knowledge *does* equip an agent for, and if it equips him for *none* then it seems that in learning a fact or truth an agent has so far gained *nothing*. Third, the idea that knowledge of a rule or principle is extricable from a capacity to apply it in a potential infinity of instances

is the target of a regress argument that directly parallels Ryle's; and it is indeed an assumption that is directly controverted by Ryle in his positive discussion of knowledge, when he claims that knowledge of a rule is proven in action not in speech, and that (propositional) knowledge is a "multi-track disposition", i.e. a practical power manifest in indefinitely heterogeneous ways[37].

Where does this leave us? I have tried to emphasise a way of thinking about the intellectualist position, and the possibilities for further investigation about the embedding of knowledge amid a network of other mental capacities and attitudes that this opens up. One may see intellectualism as simply shifting the locus of the distinction between knowing-*how* and knowing-*that*: there is no difference in type of content (no difference in type of *knowledge*), but instead (in some cases at least) a difference in the type of grasp an agent has of what he knows, and accordingly, of the types of questions and tasks we can expect him to be capable of calling on his knowledge to address. If this is how contrasts between knowing-*how* and knowing-*that* should be understood, we can still ask what is distinctive about knowing how, only the answer is to be sought in the distinctive repertoire of capacities, attitudes and sensibilities that are typically brought to bear in the acquisition, retention and deployment of the knowledge in question, i.e. those that afford the subject a distinctively practical "grasp" of what is known. Here anti-intellectualists, who have tended to be more interested in the special conditions associated with knowing how to do things than in its putative content, have done most to illuminate the terrain[38], but their insights can be brought to bear within a developed intellectualism, as part of the theory of practical guises or modes of presentation.

---

[37] See the positive discussion of knowing how in Ryle (1949, ch.II) and of multi-track abilities in (1949, ch. V).

[38] e.g. Hawley (2003), Rumfitt (2011) and Wiggins (2015).

Devitt on Linguistic Competence

Introduction

According to linguistic cognitivism, speakers have propositional knowledge of language and such knowledge is constitutive of their linguistic competence. In particular, cognitivism credits speakers with knowledge of a grammar, a finite set of axioms and rules rules assigning sentences of the language phonological form, and structural and semantic interpretations (Johnson and Lepore 2004). *Semantic* cognitivism in particular credits speakers with knowledge of axioms and rules competent to deliver an assignment of *meaning* to any sentence of the language—a compositional "semantic theory" which permits the interpretation of any sentence to be derived from the meanings of its constituent expressions and their mode of combination.

Cognitivist views have always been controversial. Two familiar sources of resistance to them relate to putative conditions on knowledge ascription. First, cognitivism imputes complex propositional knowledge to speakers which is unavailable to consciousness and which they lack the linguistic (perhaps conceptual) resources to articulate themselves; to the extent that the invoked notion of implicit or tacit knowledge is held to be dubious, so are cognitivist views of linguistic understanding[1]. Second, indefinitely many grammars for a language L may be extensionally equivalent, i.e. assign interpretations to all sentences of L. The upshot seems to be that it is indeterminate which of the competing grammars a speaker (putatively)

[1] See e.g. Baker and Hacker 1989; Hornsby 2005; Wiggins 1997; and discussion in Wright 1986. Gareth Evans (1981) also declares the idea that speakers have full propositional knowledge of language a 'mysterious and confused position' and thus dubs it "tacit knowledge".

knows, and it has been objected that under such circumstances an ascription of knowledge (of some grammar or other) is senseless[2].

Those sympathetic to cognitivism have attempted to allay these concerns, defining and defending notions of implicit or tacit knowledge sufficiently robust for some form of cognitivism but weaker than the ordinary notion of (explicit) knowledge, and indicating ways in which the indeterminacy problem might be mitigated[3]. Moreover, these concerns over the conditions under which knowledge is legitimately ascribed, while significant, fail to engage with the powerful *motivations* for cognitivism, in particular its ability to explain an otherwise mysterious dimension of linguistic understanding: its indefinite productivity (more on this below).

More recently, Michael Devitt (2003, 2006a, 2011) has argued against cognitivism on other grounds[4]. Devitt maintains linguistic and semantic theory is not concerned

---

[2] Quine 1972 and Lewis 1975. If the indeterminacy turns out to be a matter of principle, the cognitivist might be forced to maintain that implicit linguistic knowledge is so coarse grained that from an epistemological standpoint, extensionally equivalent grammars are identical—there is indeed no difference between knowing one and knowing another. But it is unclear that the indeterminacy is a matter of principle: there may be behavioural or psychological data that helps to narrow down, in principle even to uniquely identify, which grammar is in fact known by speakers: See Evans 1981; Davies 1987; Antony 2003. Chomsky (1986) also suggests that neurophysiological information may bear on the choice of grammar.

[3] On indeterminacy see note 2. In response to concerns about ascribing linguistic knowledge that speakers cannot put into words Chomsky at one time adopted the term 'cognize' in place of 'knows' to denote the propositional attitude speakers stand in towards grammar, but he is clear that (i) cognizing shares with knowing all features except articulability and availability to consciousness (thus whatever one explicitly knows one also cognizes; some things one cognizes but does not explicitly know—Chomsky1986, pp. 265-9) (ii) nothing substantial turns on the *choice* to accept or refuse application of the 'knowledge' to what he calls linguistic competence.—the issues are terminological (*ibid* and 2000, pp.93-97). Chomsky is then one proponent of the view that knowledge of language is *implicit*: unconscious and unarticulable but nonetheless integrated (at least partly) with conscious beliefs in reasoning—see also Chomsky 1980, Heck 2006a; Knowles 2000; Lepore 1997; Weiss 2004. Others have argued that linguistic knowledge is *tacit*: it is sub-doxastic and not integrated with reasoning (Stich 1978; Evans 1981; Davies 1981, 1987, 1989; Peacocke 1986, 1989). This paper prescinds from the tacit/implicit distinction; only the argument of §4 can requires the assumption that linguistic knowledge is *implicit* and not merely tacit. The issues raised there and the tacit/implicit contrast are discussed at greater length in Waights Hickman unpublished d.

[4] See also Devitt and Sterelny 1987, 1989.

with the competence of speakers at all, but with the linguistic products of that competence, token utterances and inscriptions (under certain idealisations[5]).Linguistics and semantics are not subfields of psychology, as has long been maintained in the Chomskian tradition of generative grammar—according to Devitt, they are disciplines concerned with *language*. This point, Devitt avers, seems rather obvious, and yet cognitivism ultimately rests on failure to heed it. It rests on failure to distinguish language and linguistic competence as domains of enquiry, and accordingly to distinguish the kinds of theories and explanations appropriate to each: it thus misconstrues the 'structure rules' governing linguistic outputs as 'processing rules' governing the exercise of linguistic competence. In fact, Devitt maintains, linguistic competence requires only 'something we-know-not-what…that respects the rich and complicated structure rules of a certain natural language' (Devitt 2006a, p.25). And the only *knowledge* of language that speakers possess is 'mere know-how', which is non propositional (Devitt 2011a).

The aim of this paper is to examine and defuse Devitt's line of argument and variants upon it in Soames (1984, 1989) and Levin (1977). This line of objection goes hand-in-hand with a critique of the self-understanding of generative linguistics as a branch of psychology; while I will be concerned throughout with the defence of cognitivism as a substantive thesis about the nature of linguistic understanding, and not with vindicating the methodology and self-conception of linguistics, it is worth spelling out the connection between these in order to get Devitt's target in view.

[5] Abstracting from performance errors: (Chomsky 1965, p 3-4, 1986, p.36; Devitt 2006a, ch. 1).

1. Cognitivism and the methodology of linguistics

The question whether speakers possess implicit knowledge of the syntax and semantics of their language is an old and foundational one. At the most abstract level, this is a question about how speakers and languages are related, and accordingly how theoretical work concerning language relates to an account of linguistic understanding. It has therefore occupied theorists who are principally concerned with how linguistic theory and the facts to which it is responsible should be conceived, as well as those independently interested in characterising the nature of our linguistic abilities (Smith 2006). Indeed, cognitivist conceptions of linguistic competence have developed alongside branches of linguistics and semantics in which the task of those disciplines is conceived in explicitly cognitivist terms, as one of articulating the linguistic rules implicitly or tacitly known by speakers, and in that sense revealing the cognitive foundation of speakers' linguistic abilities. Chomsky, whose revolutionary program of generative linguistics has given cognitivism its main impetus, thus maintains that the central concern of linguistic theory is a speaker's *knowledge of language* (Chomsky 1986, p. 3). Every speaker, he has long maintained, has internalised a grammar expressing his knowledge of language, and this grammar is what the linguist seeks to characterise:

> Clearly, a child who has learned a language has an internal representation of a system of rules …using the term "grammar" with systematic ambiguity (to refer, first, to the native speaker's internally represented "theory of language" and second, to the linguist's account of this), we can say that the child has developed and internally represented a generative grammar. (Chomsky 1965, p.25).

In this sense, grammar is "psychologically real" in speakers.

The 'core problem of human language' for which linguistic theory must account, Chomsky argues, is the indefinite productivity, or creativity, of language mastery[6]:

> Having mastered a language, one is able to understand an indefinite number of expressions that are new to one's linguistic experience, that bear no simple physical resemblance and are in no simple way analogous to the expressions that constitute one's linguistic experience; and one is able, with greater or less facility, to produce such expressions on an appropriate occasion, despite their novelty and independently of detectable stimulus configurations, and to be understood by others who hare this mysterious ability. The normal use of language is, in this sense, a creative activity. (Chomsky 1972, p.100)

The cognitivist view that speakers implicitly know a finite system of rules which delivers interpretations of any and all sentences of their language, purports to explain linguistic behaviour in a way that accounts precisely for its creative aspect. It is because speakers implicitly know the underlying principles in virtue of which any sentence possesses the syntactic and semantic properties that it does, that they are able to interpret and compose well-formed and meaningful sentences which are novel to their linguistic experience. Thus in semantics, a secondary impetus for what Lepore (1985) calls the 'epistemological turn' came from the work of Donald Davidson, in which the indefinite productivity of semantic understanding is also a central concern[7]. Davidson argued that our semantic theory must reveal how a language, despite its infinite productivity, is learnable by finite beings. An adequate semantic theory will be one that articulates finite knowledge possession of which would suffice to enable a speaker of a language to understand indefinitely many of its sentences (Davidson 2001[1984], pp.1-65)[8].

---

[6] Cf. e.g. Chomsky 1965, pp. 6, 15, 1966 pp. 59-7; 1986; Davidson 2001[1984], *passim* Essays 1-4).

[7] Despite Chomsky's increasing scepticism about the fruitfulness of traditional (referentialist) semantics, his work had a considerable impact on the field—see Partee 2015.

[8] Davidson's own view of the relation between natural language semantics and speaker competence is unclear. At several points he described his semantic proposal as complementary to or continuous with Chomsky's program (e.g. Essay 4, pp. 62-3)—i.e. as corresponding to the deep structure of generative grammar originally postulated within the framework of Chomsky's *Syntactic Structures* (Chomsky 1957)—and seemed to identify semantics with a theory of speakers' semantic competence (p. 25). But later he apparently accepted John Foster's suggestion that his theory be seen as articulating knowledge that *would* suffice for a speaker to understand any sentence of her language, rather than articulating knowledge that *does in fact* underlie actual speakers' linguistic abilities. See Foster 1976 and Davidson's 'Reply to Foster' in Davidson 2001[1984]. Meaning theorists following Davidson have typically, however, been happy to identify

Chomsky's conception of linguistics entails further commitments. One is the important distinction between what he calls 'competence' and 'performance'. A speaker's linguistic competence, on Chomsky's understanding, is just her *propositional knowledge of language*, i.e. the system of internalised rules or grammar which is taken to be the linguist's object of study. 'Competence' therefore does not designate the speaker's linguistic abilities or practical capacity for language use, but rather the system of propositional knowledge underlying those abilities[9]. 'Performance', and its contrast with competence, can seem more elusive (Knowles 2000, §2). It is described as 'the actual use of language in concrete situations,' and Chomsky adds that only under certain idealisations 'is performance a direct reflection of competence'. (Chomsky 1965, p.2). It is might be tempting to conclude that competence is equivalent to performance idealised, or that a theory of competence is theory of performance under idealised conditions. In fact 'performance'; refers to the way a speaker exploits and makes *use* of her knowledge, or competence, in actual linguistic transactions. Thus Chomsky goes on to explain:

> To avoid what has been a continuing misunderstanding, it is perhaps worthwhile to reiterate that a generative grammar is not a model of a speaker or a hearer. It attempts to characterise in the most neutral terms the knowledge of language that provides the basis for actual use of language by a speaker-hearer… When we say that a sentence has a certain derivation with respect to a particular grammar we say nothing about how the speaker or hearer might proceed, in some practical or efficient way, to construct such a derivation. *These questions belong to the theory of language use—the theory of performance*. (Chomsky 1965, p. 9, emphasis added).

The 'continuing misunderstanding' to which Chomsky refers is the tendency to misunderstand 'generative grammar' as denoting a theory of performance, a theory of the

semantic theory with the theory of actual semantic understanding—see e.g. Lepore 1985, Higginbotham 1995, Larson and Segal 1995, Heck 2006a, 2007; Antony 1997, 2003; Weiss 2004, 2010.

[9] The contrast appealed to here between propositional knowledge and practical ability may be somewhat misleading, for it has recently been argued that propositional knowledge is a species of ability (Hyman 1999, 2015), and that practical 'know-how' is a species of propositional knowledge (Stanley and Williamson 2001; Stanley 2011). The point is that linguistics is, in Chomsky's view, concerned with *what a speaker knows* and not with *how he is able to make use* of what he knows.

processes by which the speaker herself arrives at or "generates" the interpretations assigned to sentences by the grammar—a misunderstanding that, despite the efforts of Chomsky and others to dispel it, I shall argue continues to hamper debate over cognitivism. The confusion is understandable to some degree, given Chomsky's insistence that generative grammar is structure of the mind/brain, the language faculty[10]. This brings me to other commitments of Chomsky's, which are independent of cognitivism.

What in *Aspects of the Theory of Syntax* is referred to as the 'grammar' or 'generative grammar'–the system of rules known by the speaker and underlying her language use—is later called the speaker's internal language, or I-language[11]. This is to contrast with external language (E-language), a mapping of words or sounds to meanings (Lewis 1975), which is not the concern of generative linguistics (Chomsky 1986, pp. 26-7). The notion of I-language, Chomsky maintains, is in fact closer in certain respects to the folk, or common-sense, notion of language. For 'when we speak of a person as knowing a language we do not mean that he or she knows an infinite set of sentences, or sound-meaning pairs taken in extension, or a set of acts or behaviours; rather we mean that the person knows what makes sound and meaning relate to one another in a specific way'. (*ibid* p.27). The rules of language known by speakers are thus rules that 'form or constitute the language, like…rules of chess (not a set of moves, but a game, a particular rule system).' (*ibid,* pp.27-8).

[10] E.g. Chomsky 1986 pp. 3-5, 23, 26-7, 29, 44-5 and *passim*. See Fodor 1983, ch. 1 on the potential of talk of the 'faculty', 'structure' and 'module' in connection with linguistic competence to obscure the fact Chomsky sees a speaker's innate competence, as well as her mature attained competence, as a body of *information* or *knowledge*.

[11] See Chomsky 1986, pp. 26, 29. Correspondingly what was in *Aspects* referred to simply as 'language' is later denoted 'E-language' or 'external language'.

So far the move to talking of an 'I-language' instead of grammar is only a shift in terminology, and as Heck (2006) suggests, the prefixes 'I' and 'E' might be more illuminatingly understood as standing for Intensional and Extensional, rather than Internal and External. Thus

> An E-language is a set of grammatical sentences, or a function (in extension!) from sounds to meanings; an I-language is a set of syntactic or semantic principles that generates an E-language, much as a set of axioms and rules generates a set of theorems.

However, Chomsky's understanding of an I-language, what I shall from hereon revert to calling a grammar, is radically (and increasingly) internalist. The rule system that is taken to constitute the individual speaker's knowledge of language is held to be an internally represented part of her mind/brain, which is somehow physically realised (1986, p.24). Thus one of the advantages Chomsky claims for the study of I-language over E-language is that, while the latter is a theorist's idealised construct, the former is an element of 'particular mind/brains, aspects of the physical world, where we understand mental states and representations to be physically encoded in some manner' (*ibid*, p. 26). Moreover, the speaker's knowledge of language is not thought to be accountable to anything in the external environment. It is not knowledge of syntactic and semantic principles and properties possessed by external symbols *independently* of the speaker's language faculty (Chomsky 2000, p.170-1).

It it is partly Chomsky's idea that grammar must be *eo ipso* a psychological phenomenon, pertaining to nothing without the mind, that Devitt seeks to call into question. But his core target is cognitivism, the view that speakers implicitly know the grammar of their language, and the central motivations for *cognitivism* do not require a mentalist ontology of grammar or internalist conception of linguistic

properties. If Chomsky is right that knowledge of a system of generative rules comprises the cognitive basis of a speaker's ability to use and understand indefinitely many novel sentences of her language, it does not follow that the rules, as objects of speakers' knowledge, are internal elements of the mind/brain. As Alexander George observes,

> The grammar is not itself a characterisation of a system of mental representation; it is the object of a speaker's knowledge, not a description of how that object is represented by the speaker (if it is). ... The grammar is what is represented, not what does the representing. (George 1989, p. 91)

Moreover,

> The foregoing is agnostic with respect to which particular stripe of materialism is true; it is even agnostic about whether any brand of materialism is true. (*ibid*).

Neither is cognitivism an implicitly internalist view[12]. Consider the familiar language-chess analogy that Chomsky himself exploits. We may suppose that knowledge of the rules of chess, a human artefact, is (at least partly) constitutive of a player's ability to make and assess indefinitely many novel moves; and that imputing this knowledge in fact best explains the creativity exhibited by players in making and assessing such moves. It follows neither that a theory of chess is *eo ipso* a theory of an aspect of the mind/brain of chess players, nor that the properties that the rules account for, e.g. the property of being a licit move, have no existence beyond the mind of the player and as such are not ones about which the individual player can be mistaken. In the case of chess, the more plausible view is surely that these

[12] Indeed, some Chomsky commentators have questioned the cogency of an epistemological conception of grammar given radical internalism, for here the speakers 'knowledge' is not *of* anything mind independent: it has no externally imposed correctness conditions – See Collins 2004, 2007; Rey 2003; and Smith 2006 for discussion.

properties are fixed institutionally, and for this reason individual players *can* be mistaken about them[13].

2. Language and its mastery; theory of language and theory of understanding

Devitt takes Chomsky's cognitivism and associated conception of linguistics to rely on a 'very fast argument for the psychological reality of linguistic rules—we have good evidence for the grammar so we have good evidence for the psychological reality is describes' (Devitt 2006a, p.9). His contention is that grammars are rule systems governing languages, not the competences of speakers, and that the assumption that speakers must know the grammar of a language in order to be competent is unmotivated and highly implausible (independently of any skepticism about implicit or tacit knowledge states as such). On the contrary, as the title of his recent book suggests, Devitt argues that speakers can be entirely *ignorant* of language (Devitt 2006a, p.5).

Devitt's critique rests on a series of analogies and a set of distinctions. The core distinctions are as follows:

(1) Competence *vs* the product(s) of competence

(2) Processing rule *vs* structure rule.

[13] I agree with Chomsky's contention that, in contrast to the case of arithmetical truths, 'there is no initial plausibility to the idea that…there is an additional domain of fact about [language], independent of any psychological states of individuals' (Chomsky 1986, p. 33). But this intuition is consistent with some form social externalism, according to which the semantic properties of expressions are partly fixed by social convention. On such a position it can be granted that it is unintelligible that a whole speech community should be mistaken about what an expression of their language means, even if is perfectly intelligible that an individual should be in error about its meaning (Fricker 1983). For attempts to reconcile some form of externalism within a broadly Chomskian framework see Ludlow 2003, 2011 and Longworth 2007.

(3) Identity/isomorphism with structure rules *vs* respecting structure rules.

(4) Rule represented and applied *vs* rule embodied.

A series of analogies then serves in part to elucidate these contrasts. The competence of a blacksmith, compared with horseshoes, the product of that competence, provides Devitt's first illustration of distinction (1). This example best fits the interpretation of 'competence' that Devitt offers. A competence is a productive capacity that is partly individuated by what it is a capacity to produce: 'the very nature of the competence is to produce its outputs: producing them is what makes it the competence it is' (Devitt 2006a, p.22). 'Linguistic competence', it follows, will not in Devitt's parlance automatically acquire the technical meaning that Chomsky adopts, as the speaker's underlying (propositional) knowledge of language—that is, a grammar.

Of course, Devitt's understanding of 'competence' is not straightforwardly incompatible with Chomsky's. The cognitivist line concerning Devitt's notion of linguistic competence will be that the productive competence of speakers of a language L to produce (and presumably comprehend) sentences of L *just is* a state of propositional knowledge whose content is a grammar. Indeed, according to Devitt 'the theory of a competence explains what it is about an object that makes it competent'(*ibid*, p.18), and on one reasonable interpretation this is precisely what a theory of the speaker's internalised grammar achieves, by Chomsky's lights, as a theory of *linguistic* competence. This is important if Devitt's starting point is not to be entirely question-begging—he cannot simply *define* linguistic competence in a way that excludes a cognitivist interpretation. Nevertheless, it is worth noting that Devitt's conception of competences as productive capacities individuated by their products is tendentious to the extent that (1) it most naturally coheres

with Devitt's own view of linguistic competence as a skill, or non-propositional 'mere know-how' (Devitt 2006a, 2011) for skills and abilities are individuated by what they are skills and abilities to do; (2) conversely, it is uncongenial to a cognitivist view on which propositional knowledge *of* language is distinct from, and the basis of, knowledge *how to use* and understand it (e.g. Matthews 2006a). Subject to these caveats, we can accept Devitt's distinction (1) as a more-or-less neutral starting point from which to build a framework for assessment of linguistic cognitivism.

The array of other examples of competences and products to which Devitt appeals is heterogeneous, and includes the competence of the honey bee to "produce" its waggle-dance as output; the competence of a chess player to produce licit chess moves, and the competence of a calculator to produce correct calculations (*ibid* pp. 15-21). In every case Devitt stresses that while the characterisation of the competence and of its product are related in the way indicated above—the competences are defined in part by what they are competences to produce—to give an account of one and of the other are in general quite distinct enterprises. An account of the calculator's competence to produce correct calculations will, for example, appeal to whatever internal mechanisms mediate between the inputting of, say, '7+5' and the calculator's outputting '12' on the screen, but these are not mechanisms appealed to in an account of arithmetical calculations as such. Conversely, an account of the blacksmith's horseshoes (the output or product) will make reference to its physical properties, properties that are not part of the blacksmith's competence to produce them (*ibid*, p. 15).

This observation is then applied to the case of *linguistic* competence and language as the external object that is its product: token written or uttered strings (*ibid*, p.24-5). What

comparison with examples such as the blacksmith and his horseshoes is supposed to impress upon us is that language and linguistic competence are *obviously* distinct phenomena: 'competence in the language is not the language any more than the blacksmith's competence is a horseshoe, the chess player's competence is chess, or the bee's competence is its dance. The linguistic competence is in the mind/brain, the language is not.' (*ibid*, p.23). The enterprises of characterising each phenomenon are, accordingly, obviously different; and there is no reason to anticipate that the two characterisations will converge, i.e. that in the linguistic case a generative grammar will provide *both* a theory of language *and* an account of linguistic competence (*ibid*, p.24).

Distinction (2) is the most elusive in Devitt's discussion, but also the most critical. It is supposed to align with the first contrast: processing rules are precisely those that govern the exercise of the competence, while structure rules are rules that in some sense define or characterise its product (*ibid*, p. 18).

The first and clearest illustration Devitt offers of the contrast is a 'logic machine' which takes well formed formulas (*wff*s) as inputs, and produces *wff*s as outputs, by applying transformation (or derivation) rules. The *wff*s themselves are defined by formation rules (rules that determine what is to count as a *wff*), thus all outputs of the logic machine respect them: these are the "structure rules". However, these are not the rules that govern the process by which the output is produced—the "processing rules" in this case are the transformation rules. Devitt's notions of processing rule and of structure rule are, however, supposed to have purchase in connection with *all* competence-product pairs. Thus the following is cited by Devitt (p.20) as a characterisation of a structure rule of the honey bee's waggle dance:

> To convey the direction of a food source, the bee varies the angle the waggling run makes with an imaginary line running straight up and down…if you draw a line connecting the beehive and he food source, and another line connecting the hive and the spot on the horizon just beneath the sun, the angle formed by the two lines is the same as the angle of the waggling run to the imaginary vertical line (Frank 1997, p.82).

We can then ask what processes the bee goes through to produce the dance from its "memory" of the direction of the food source: whatever rules 'govern this unknown process'; are the processing rules governing the bee's competence (Devitt 2006a, p.20). And in this connection we should not, Devitt urges, 'rush to the judgement that the structure rule itself must govern this unknown process. It might be the *wrong sort* of rule to play this role.' (*ibid*). In the case of a pocket calculator, the electronics responsible its production of correct additions are characterised by processing rules, as compared with the addition rule defining the produced calculation (*ibid*, p.48); and in the case of chess, the defining rules of the game are the "structure rules", but even these need not be processing rules, Devitt maintains. A player need not go through, e.g., 'a process of inferring "*x* moves diagonally" from "*x* is a bishop"' when considering a move with the bishop (*ibid*, p.19). It is hard to say exactly what Devitt imagines the processing rules *are* which 'govern the process of producing the output' (ibid, p. 18) in the case of the bee's dance, the calculator's calculations, and the chess player's moves. But it seems clear that they must govern a process, be it physical or computational, that goes on *in* the competent agent, and which eventuates in its production of an output.

The point of principle on which Devitt insists is that being a processing rule and being a structure rule are different roles or properties (*ibid,* p.19) and that for there to be a competence to produce an output with certain structure rules, what is required is that the processing rules governing the exercise of the competence *respect* the structure rules, not

that they include them (distinction (3)). That is, the processing rules must be such as to guarantee the production of an output conforming to the structure rules (*ibid,* pp.21-2). With the conceptual point made, the illustrations referring to insects and machines are presumably to sow doubt as to whether structure rules governing a product are *ever* among rules applied in the exercise of the competence that produces it—certainly to normalise the phenomenon of mismatch. For, as Devitt repeatedly stresses, these seem to be different *kinds* of rules (*ibid*, pp. 37, 53), making bizarre the idea that the structure and processing rules must be identical or isomorphic.

Turning back to the linguistic case, a grammar is construed by Devitt as the system of structure rules governing the products of linguistic competence: 'the outputs of linguistic competence, physical sentence tokens, are governed by a system of rules, just like the outputs of the chess player, the bee, and the logic machine.' (p.24). The processing rules are whatever rules govern the linguistic competence of the speaker, such as to ensure a product of the right kind—one conforming to the system of structure rules—is produced. And if the assimilation of structure and processing rules seems bizarre in the case of the bee and the calculator, it should also seem bizarre in the linguistic case: we should be just as suspicious of the idea that the rules of grammar are psychologically real (*ibid,* pp.64-5).

Finally, Devitt's distinction (4) recognises two ways in which a processing rule may figure in the exercise of competence; in particular, two ways in which a *structure rule* could conceivably be included as one of the *processing rules*. On the one hand, the processing rule may be one that is represented and applied by the competent agent or system—this is what we would expect if the rule is known. On the other hand, the rule may be "merely embodied" in the competent agent/system, without being represented (*ibid,* p. 45). Devitt

illustrates this contrast with the example of a pocket calculator. A particular arithmetical rule is not represented by the calculator, Devitt says; the calculator doesn't *know* the rule—and we can be 'quite confident about this', he affirms, 'because the calculator is not the sort of thing that can know anything' (*ibid,*p. 48). On the other hand, 'in virtue of the calculator being governed by those rules we can say, if we like, that the arithmetical "information" that those rules reflect is embodied in the calculator even though the calculator does not represent them.' (*ibid*). Similarly, Devitt says of a kingfisher diving to fish, that is is 'not plausible to suppose that the kingfisher represents any of [the] facts about refraction and angles' that govern its fishing; rather the 'mechanisms, reflecting water to air refraction, are simply built into the kingfisher "by nature" just as the rules for the calculator are built into it by an engineer.'(*ibid,* p.49). It is not entirely clear to me what Devitt means by the rules or information in these cases being 'embodied' by the system or organism, that goes beyond the fact that, in Devitt's terms, their *output* is governed by these rules as *structure rules*. But whatever he does intend, it seems clear that neither the notion of a (structure) rule that is embodied, nor of a (structure) rule that is represented and applied by the competent agent, fits the fix we are given on the notion of *processing rule* by Devitt's other illustrations. I shall return to this point below.

The considerations that Devitt adduces to make on linguistic cognitivism appear implausible and unmotivated if not downright confused, have some echoes in the last few decades of anti-cognitivist literature, particularly in the work of Soames (1984, 1989). Like Devitt's, Soames' critique of cognitivism proceeds from an insistence upon certain distinctions. First, Soames maintains, there is a 'conceptual' distinction between the subject matter of linguistics and semantics which is 'the structure of natural languages, considered in abstraction from the cognitive mechanisms causally responsible for

language acquisition and mastery' (Soames 1984, p.157), and the subject matter of psychology, which is the 'the states and processes mediating sensory inputs and behavioural outputs' (*ibid* p. 163); accordingly there is a distinction between the 'structures' postulated by linguistic theory and those to be uncovered by a computational theory of language processing (*ibid*, p. 167-8). –These are reminiscent of Devitt's distinctions (1) and (2). Like Devitt, Soames argued that there is no reason to expect isomorphism between linguistic and mental 'structures'—indeed, that there is positive reason to expect divergence, since processing efficiency is likely to be important at the psychological level, while simplicity and elegance are criteria that may bear on our choice of linguistic theory (*ibid* pp.169-70).

Levin (1977) accuses cognitivists following Chomsky and Davidson of conflating two notions of *explanation*: the 'analytical' explanation that generative grammars and semantic theories provide of the linguistic strings they structurally/semantically interpret; and the causal explanation provided by psychological structures or processes involved in language use and comprehension. A generative grammar, he contends, 'explains why' a particular string is grammatical only in the logical sense that this string is entailed by the principles of the grammar. In this sense a function also explains a set of coordinates on an *x,y* plane, or argument-value pairs. Levin's point is that this kind of explanatory power is not the kind claimed for grammars by Chomsky and his followers—it is not causal or psychological. One of his illustrations could easily be incorporated in Devitt's collection of analogies and it does much the same work: the example of a machine that plots the graph of a certain function (*ibid*, p.129). While the function may, in the logical or analytical sense, explain the co-ordinates plotted by the machine, it does not explain the mechanisms by which the machine is able to plot the graph. Similarly, Levin thinks, the "explanations"

provided by linguistic theory are not (causal or psychological) explanations of speakers' linguistic interpretations or judgements at all:

> To see that the generation of a string is not an [causal] explanation of why the string is grammatical, consider again the analytic geometer. Given an infinite array of points on an empirically given *x,y* plane his problem is to find a function $\lambda x f(x)$ such that precisely the ordered pairs $[x,f(x)]$are the co-ordinates of the points. … causal questions about how a pair (m,n) got into the array have not arisen yet. Someone who wants to know why $(m,n) \sim \lambda x f(x)$—the non-eventival analogue of why S is grammatical—is someone who doesn't understand the computation of n from $f(m)$or does not believe that $f(m)$ is n…"Explaining why" $[(m,n) \varepsilon \lambda x f(x)]$,or why $f(m) = n$, amounts to tracing the *f*-computation of *n* from *m*. Tracing the derivation of a grammatical string from the rules is this kind of explanation. It shows "how" the string is generated in the same way a computation shows "how" *n* manages to be $f(m)$ or a proof of a theorem X shows "how" X follows from premises. Thus, if the question "Why is S grammatical?' is interpreted as "How did S get on the grammatical list?" it is not an explanandum of the empirical grammarian's theory.

Like Devitt (and Soames), Levin urges that an explanation of speakers' intuitions 'would be one which cited the causal processes leading up to these intuitions, something [the grammar] is not competent to do.' (p. 130). And like Devitt (and Soames), Levin complains that the idea that the derivational procedure of the grammar matches the psychological processes underlying language use and comprehension is arbitrary and unwarranted: 'It requires further argument, and a separate premise that does the whole explanatory job, to show that the speaker executes processes isomorphic to the generation process' (*ibid*, p.131).

To sum up, the argument brought against cognitivism by Devitt and others is this.

First, grammars and semantic theories are, crudely put, *linguistic* theories, not theories of *mind*. They articulate principles for a language from which it is possible to derive structural or semantic interpretations of any sentence of that language, in this way making clear the underlying structural principles of the *language* as a whole, not the psychological principles underlying speech and comprehension in speakers. Second, as Devitt's range of analogies purports to show, it is arbitrary to assume that the principles of linguistic theory ("structure

rules") *correspond* to or match in any way the psychological "processing rules" underlying speech and comprehension. It would be strange, after all, if the rules that govern the bee's waggle-dance also governed the internal processes of the bee which ensure its behaviour "respects" the rules of the waggle dance. Third, the difference in the role of structure and processing rules can be seen from the difference in kind of *explanation* each kind of rule affords. The "structure rules" of a grammar or semantic theory are explanatory in the analytical sense: just as a function explains the derivation of a set of values from a set of arguments, a grammar explains how an interpretation can be derived from an uttered string. By contrast, a processing rule of linguistic comprehension contributes to a causal explanation of how a speaker in fact arrives at a given interpretation.

3. Dialectics

The form of attack articulated above is not only popular but ostensibly the most serious for cognitivism. For the other prevalent complaint against cognitivism, namely that it entails an implicit knowledge hypothesis that it impossible to swallow, makes absolutely no contact with the primary motivation for cognitivism: its claimed explanatory power with respect to the productivity of linguistic understanding. By contrast, the line of objection pressed by Devitt et al challenges the cogency of cognitivism's explanatory aspiration, seeing it as premised on conflations of language and linguistic competence, and accordingly of the theories that can illuminate each phenomenon.

It is worth treating this form of objection with some patience, then, but I believe it is entirely unsuccessful. This is not because it raises challenges to cognitivism that have familiar *answers*—for example, because the 'assumption' of isomorphism between

processing and structure rules is not at all as wild and arbitrary as Devitt would have us believe. Rather, my contention is that there *is no challenge* to cognitivism that the considerations introduced in §2 succeed in bringing to the cognitivist position.

The idea that cognitivism might rest on a failure to notice that external language and speakers (or competences) are categorically different is surely naïve. From the standpoint of Chomsky's internalism, Devitt's contrast between language as output and the competence form which it issues, is most naturally construed as one between performance episodes (uses of language) and the chosen focus of generative linguistics: I-language, or grammar, an aspect of the mind/brain (Collins 2008, p. 22)[14]. On one construal, the generative grammarian's focus on an aspect of mind, the speaker's I-language, is simply a methodological *choice* or stipulation, neither a premise nor an upshot of some independent theory. As John Collins puts it, 'what Devitt regards as a fast and dirty move to the "psychological reality of grammar" is in fact simply a demarcation of a domain of enquiry—linguistic cognition—and a proposal for its study—theories of the function computed.' (Collins 2007, p. 418)[15]. The legitimacy of such a methodological choice, moreover, is not something to be settled *a priori*, but to be determined by the fruitfulness of research in the chosen domain of investigation (Collins 2008, p.21).

There is surely some force to these considerations. Devitt's narrative can seem to supress the fact that theories of generative grammar are developed with the explicit goal of

[14] Perhaps alternatively, between I-Language and E-language, but then the centrality of the distinction in Chomsky's writings and the irrelevance of Devitt's topic to Chomsky's concerns is even clearer.

[15] Cf. Chomsky 1989, p.4: 'Generative grammar is…not a theory, any more than chemistry is a theory. Generative grammar is a topic, which one may or may not choose to study. Of course, one can adopt a point of view from which chemistry disappears as a discipline…Similarly, one may argue that the topic of generative grammar does not exist, although it is hard to see how to make this position minimally plausible.'

explaining an aspect of the mind. Instead of objecting that linguistics, *qua* branch of psychology concerned with the language faculty, is providing unconvincing answers to one of its central questions ('what constitutes knowledge of language?' (Chomsky 1986, p.3)), Devitt simply denies that linguistics *is* a branch of psychology and that its theories speak to psychological research questions—and this while admitting that linguistics as currently practiced is enormously successful (e.g. Devitt 2006a, p. 3).

Nevertheless, the reply that generative linguistics is just a topic of enquiry is not enough to impugn the pertinence of Devitt's challenge. Domain delineations are not incontestable, and hard as it may sometimes be to credit, Devitt's challenge to the self-conception of linguistics is explicit and unabashed. (It is also not unique—cf. Soames' critique discussed above, Katz 1984 and Stich 1971). Moreover, the received conception of generative grammar Devitt contests is motivated *in part* by precisely the idea that Devitt takes issue with, namely the idea that that if our linguistic behaviour and judgements conform to a certain rule system, we should assume that the system is psychologically real (Chomsky 1980, p.130).

To be fair to Devitt, then, since his aim is to challenge the view that grammar is known or psychologically real, he cannot sanction the internalist interpretation of distinction (1), as opposing I-language or linguistic knowledge, with the external outputs of that competence. His contention is precisely that the view of grammar as an internal aspect of mind is optional: one can see a grammar as a generative procedure from which all linguistic strings *can be* derived and interpreted, without seeing it as a psychologically real procedure that competent speakers *in fact* exploit when they produce and interpret those strings. Still, assuming an externalist interpretation of the contrast makes the starting point far from

neutral, it cannot be admitted by (the internalists among) Devitt's opponents. On the contrary, half the argument against Chomsky's conception of linguistics depends on sanctioning an externalist construal of distinction (1).

Cognitivism, the view of interest to me here, however, does not rely upon internalism. The imputation of linguistic knowledge that cognitivism recommends, just like any other imputation of knowledge, is neither sustained nor undermined by acknowledging that the agent's cognitive systems ('competence') are one thing, while the domain of their putative knowledge (in this case language) is another. It is not an ontological thesis concerning the realm to which linguistic properties belong or the domain to which linguistic theories are immediately answerable. It is an epistemological thesis, a thesis concerning the knowledge that speaker's have at their command, not the ontological status of the *contents* of linguistic knowledge. The idea that speakers have knowledge of linguistic rules is thus no more premised on a failure to properly distinguish language (and its theory) from linguistic understanding (and its theory) than the claim that chess players *know* the rules of chess is premised on failure to properly distinguish the game of chess (and its theory) from chess players' competence (and its theory).

4. Competence, knowledge and language processing

As I mentioned at the beginning of §2, Devitt's understanding of 'competence' is in itself not entirely innocent, although it can be made consistent with cognitivism. In assessing the claim that a theory of grammar is a theory of linguistic competence, one danger of uncritically adopting Devitt's notion of competences as productive capacities partly individuated by their outputs, is that it suppresses the important distinction discussed

above, between the theory of competence in Chomsky's sense, i.e. the theory of the speaker's underlying knowledge of language, and the theory of performance, the theory of how that knowledge is put to use in concrete speech and comprehension situations. This is because one could reasonably expect a theory of competence in *Devitt's* sense of competence to explain *how* the competent agent goes about producing a product or processing an input, and not merely to characterise any information that enables the agent to do so. This turns out to be critical.

The crux of the anti-cognitivist critique outlined above relies on distinction (2), the distinction between structure rules and processing rules. The central complaint against cognitivism is that it rests on an arbitrary assumption of isomorphism between language processing rules and the rules of a grammar ("structure rules"). But in assessing such an isomorphism 'assumption' it matters a great deal what a processing rule *is*. I believe that the force of Devitt's critique depends almost entirely on assuming a conception of processing rules that has nothing to do with the cognitivist thesis; ultimately it rests on projecting a misconstrual of grammar as a theory of performance, and not of competence or knowledge.

Cognitivism claims that linguistic rules are *known* by competent speakers and that this knowledge is exploited in language use and comprehension. So the rules exploited in speech and comprehension are certainly taken to be isomorphic with rules of grammar, on a cognitivist view, insofar as they are taken *to be* those very rules.

But the label 'processing rule' is unhappy for a rule that speakers know and apply or exploit. Consider the claim that a pre-verbal infant comes to know the rule that red blocks

are placed in one container while blue blocks are put in another, and that she is guided by this knowledge when she puts each kind of brick in its correct container; or the claim that a chimpanzee knows and observes, in an experimental situation, that a stone from the chimp trades for a piece of cucumber from the researcher. Whatever the strength of such claims, they are only claims about putative mental processes in the minds of the child and chimp in so far as they postulate particular *items of knowledge* as inputs. It is no more apt to dub these items of (postulated) *implicit* knowledge 'processing rules' than it is in cases where the contents are known explicitly—than it is, for example, to describe the cosine rule explicitly referred to and used by a student in a trigonometry test as a 'processing rule'.

The issue is not merely one of infelicitous terminology. While Devitt's distinction (4) ostensibly makes room for one appropriate conception of a processing rule—a rule "represented and applied" by the competent agent or system—such a conception is at odds with his illustrations and glosses of distinction (2) between structure rules and processing rules, and his stress on the idea that it contrasts two different *kinds* of rules with different *roles*. Processing rules are, Devitt says, rules that 'govern a process that the system actually goes through' (*ibid*, p.53), while structure rules govern the outputs of that system. Thus the processing rules responsible for ensuring the bee performs its waggle dance are rules that govern the processes internal to the bee which eventuate in its performing certain movements, while the structure rules govern its dance; and the processing rules ensuring the calculator produces correct calculations are rules that govern electronic or computational processes eventuating in those calculations appearing on the calculator's screen. Devitt is, I think, correct that structure rules 'seem to be the wrong kinds of rules' to be processing rules (*ibid*), but is disingenuous in implying that Chomsky and other

proponents of linguistic cognitivism are guilty of identifying language processing rules with grammar (or the rules of a semantic theory).

As Devitt himself eventually acknowledges (*ibid* pp.64-5), Chomsky has repeatedly warned against interpreting generative grammar as an account of a mental process gone through by the speaker in reaching a linguistic interpretation—an example is the passage from the first chapter of *Aspects of the Theory of Syntax*, quoted in §1. Chomsky dismisses this idea as absurd in another passage which Devitt quotes, pointing out that the grammar-as-processing-model view would 'take generative grammar to be a model of performance rather than a model of competence, thus totally misconceiving its nature…[grammar] can be regarded only as a characterization of the tacit knowledge or competence that underlies actual performance.' (1965, p.140-1; Cf. 1986, p.67 and 1965, p. 9, quoted in §1). But Devitt seems to miss the point and take Chomsky to fail to recognise his own commitments. Devitt makes much of Chomsky's use of terms such as 'generation' and 'computation'[16] and insists on interpreting their use in connection with grammar in a way that Chomsky explicitly disavows (Devitt 2006a, p. 62). In Chomsky's remark that the conformity of our behaviour and judgements with certain rule systems is best explained 'by the assumption that computation involving such rules and the representations they provide takes place in the mind' (Chomsky 1980, p. 130), Devitt contrives to find confirmation that Chomsky believes 'those rule systems *govern processes* that really go on in the mind.' (Devitt 2006a, p. 63, emphasis added). Thus Devitt goes on attacking an unoccupied position according to

[16] But see Chomsky 1965, p.9, following on from his emphatic rejection of the idea that grammar provides a 'model of a speaker hearer':

> Confusion over this matter has been sufficiently persistent to suggest that a terminological change might be in order. Nevertheless, I think that term "generative grammar" is completely appropriate and have therefore continued to use its. The term "generate" is familiar in the sense intended here in logic…since this use of the term "generate" is well established both in logic and in the tradition of linguistic theory, I can see no reason for a revision of terminology.

which at some level of abstraction or idealisation, grammars or I-languages 'are psychological processes…going on in real time in the mind.' (*ibid*, p. 66)[17].

This is not at all consonant with Chomsky's remarks. Consider analogous statements concerning mental computation and the rules of traffic, chess, archery, cooking, or etiquette. If we say, for example, that drivers perform (largely unconscious) computations involving traffic rules and that this best explains their conformity with those rules in driving, we cannot be accused of implying that traffic rules govern the mental processes of drivers: they govern the traffic! This is the sort of point Devitt is at pains to make about structure rules: they govern the products or outputs of a competence; and they are the *wrong kind of rules* to govern the internal processes that eventuate in the production of those outputs. For them to be the right kinds of rules, the mental processes that underlie an agent's driving competently would have to exhibit the same patterns as the movements of cars on the roads, an idea that Devitt rightly finds absurd in connection with any competence-product pair, including linguistic competence and its outputs.

Soames' and Levin's complaints seem similarly premised on a misinterpretation of the respect in which grammars and sematic theories purport to identify the basis of speakers' linguistic abilities and accordingly to explain their speech and interpretation performances. Soames, for example, compares the received view of grammar as a theory of competence with the idea that 'mathematical formalizations of elementary number theory' provide 'psychological theories of the *cognitive processes* underlying mathematical reasoning' (Soames 1984, p. 155-6 emphasis added), and in articulating the ways in which

[17] According to Geurts and Rubio-Fernández (2015) a similar confusion has hampered discussion of Grice's theory in pragmatics—i.e. it is misconstrued as a theory of the psychological processes undergone by participants in a conversation, and then argued to be psychologically implausible.

he thinks linguistic theory does and does not constrain psychology he writes: 'Universal Grammar, which attempts to characterize the class of humanly possible natural languages, constrains theories of language acquisition. Any such characterization implicitly specifies the scope of the human cognitive capacity to learn languages, without speculating about the *causal mechanisms* that implement this capacity.' (*ibid*, p. 157 emphasis added). Again, the conception of linguistic theory he recommends is one on which it abstracts from 'the cognitive mechanisms casually involved in language acquisition and mastery'(*ibid*). This recommendation, considered in light of the mathematical analogy, is compelling. But it also seems to be one that the cognitivist subscribing to Chomsky's vision of linguistics can, for a part, embrace: on her view the theory of grammar *qua* theory of speaker competence is not the theory of a cognitive *mechanism* or cognitive *process*, it is the theory of the propositional content of a cognitive state, of the *information* available to the speaker.

Levin's complaint concerns the notions of explanation appropriate to linguistics and psychology as he conceives them: a grammar can only provide an analysis, or analytically "explain" the interpretation assigned to particular strings, or why they are accepted as well formed. A genuine psychological theory of linguistic judgements must be *causally* explanatory, by contrast. And his claim is that grammar cannot provide this kind of explanation unless it is, just as Devitt contends, a model of language processing, which there is no reason to believe it is—cf. Levin's example of the graph-plotting machine, whose internal workings are not described by the function that defines the plotted graph. But the cognitivist, once more, sees grammar as the content or object of linguistic knowledge, not as a psychological process that eventuates in linguistic behaviour or judgements. If linguistic knowledge or competence enters into causal explanations of behaviour, it does so in the way that any other knowledge state does: if *knowing* that *p* is a

causally efficacious attitude involved in bringing about certain behaviour, it will (partly) causally explain the behaviour *ex-post. What* the agent knows, i.e. that *p*, need not be anything with independent causal efficacy. To return to the well-worn chess analogy, if an agent knows that the bishop in chess moves diagonally, and this knowledge state is implicated in causal processes eventuating in her making a licit move with a bishop, it must figure in a causal explanation of that move. But it is not *what* she knows, *viz*. the fact that the bishop moves diagonally, that is causally explanatory of her behaviour, but her possession of the knowledge; and by analogy, it is not the grammar itself that is potentially *causally* explanatory of language use and understanding, even if *knowledge of* the grammar is. Most importantly, the sum total of the defining rules of chess is not, insofar as it is computationally deployed by players making licit chess moves, a system of rules that governs chess players' mental processes and in this sense causally explains the player's moves. Neither is a grammar a system of rules governing the mental processes eventuating in linguistic judgements.

5. 'Creatures to whom can be ascribed intention and purpose'

Devitt's focus on analogies involving bird, insect and machine behaviour exploits our pre-theoretical resistance to attributing any kind of intentionality in these cases. Thus he affirms that we can be 'quite confident' that the calculator does not *know* arithmetical rules, 'because the calculator is not the sort of thing that can know anything.'(*ibid*, p.48). It is unclear whether current psychology in fact supports Devitt's intimations that birds and bees are ignorant of the principles that govern their complex behaviour[18], but there

[18] See Antony (2007) on this point. Indeed, rather oddly, Devitt himself asks (apparently rhetorically) whether the dancing bee can 'even manage the necessary representations of the food source, of the spot on the horizon, and of the angles' for performing a computation involving structure rules of the waggle-

is certainly a pre-theoretical intuition that however these complex behaviours might be explained, it can't be in terms of intentional states. The question is: why should we take instruction from these kinds of examples when theorising about human language mastery? What Devitt presumably seeks to illustrate is that a system or organism can produce outputs and behaviours that are enormously complex and sophisticated, without that sophistication and complexity being mirrored in the cognitive systems of the competent agent. Ignorance is no obstacle to sophisticated performance. So we should not feel under pressure to credit human beings with knowledge of the complex "structure rules" that their linguistic performances (under certain idealisations) conform to: they may be entirely ignorant.

In this connection it should be noted that the modest observation made by Devitt and many other critics of cognitivism, that conformity with rules does not by itself entail knowledge of them[19], is independent of any misapprehensions about the kind of model of competence that a grammar is purported to provide. That the psychological reality of grammar must be understood as a hypothesis about speaker's *knowledge* and not about language *processing*, does not, of course, place it beyond dispute. But equally, cognitivists have never maintained that their thesis is deductively *entailed* by the fact that (idealised) language use conforms with syntactic and semantic rules. It is motivated by adbuctive arguments, as the hypothesis about the nature of linguistic knowledge that provides the best going explanation for linguistic creativity[20].

dance; but then, in the footnote, Devitt refers to a plausible theory on which the bee *is* capable of these representations (Devitt 2006a, p.21and fn. 5).

[19] E.g. Quine 1972, Foster 1976, Wright 1986.

[20] See e.g. Partee 1975, 1979; Higginbotham 1988, 1995; Larson and Segal 1995.

I want to argue, however, that *pace* Devitt, there is an independent source of pressure to credit human beings with knowledge of the semantic and syntactic rules of the languages they speak, which connects with the distinctive kind of explanation that reference to linguistic knowledge provides. Whatever should turn out to be true of certain kinds of birds and insects, we know that much of human agency is purposeful and intentional; and indeed, our linguistic agency is of this kind.

If we alter the anti-cognitivist's analogies to focus on various forms of sophisticated, intentional *human* behaviour, the hypothesis that this behaviour is blind, ignorant, uninformed, ceases at once to appear the obvious or default position. If we dwell on a *human being's* performance of a dance, plotting of a graph, or calculation of a sum, this does not yield the same unequivocal intuition that knowledge of the relevant "structure rules" is not what guides the behaviour: certainly these are not things that human beings "just do", so to speak mechanically, they must be *somehow* motivated or rationalised. There is nothing absurd about the idea that a dancer draws on knowledge of the very principles that his dancing conforms to—that this conformity is indeed secured by, is the product of, his knowledge. And while a student *might* be plotting co-ordinates on an *xy* plane without knowing the function that describes them—e.g. by consulting a table in which the coordinates are listed, if she produces, for any arbitrary value of *x*, precisely the value of *y* yielded by applying a particular function *f(x)*, there is nothing wild about the supposition that she knows and is *applying* that function. The imputation of knowledge seems still more natural in the case of arithmetical calculations. If we cannot be sure that a human being who comes to reliably add correctly has come to *know*, and in such calculations deploy knowledge of, the addition rule, what *would* suffice to warrant an imputation of knowledge of this kind? It is at least a substantial, controversial thesis which

declares human actions of this sort, to the extent that they are intentional, to be compatible with *ignorance*. And so in Devitt's collection of analogies for the assessment of linguistic cognitivism, the most pertinent is is the least rhetorically effective: the chess player's making of licit chess moves. Of course the rules of chess are not processing rules, and much else besides knowledge of chess rules will be recruited by a player in the course of a game—as Devitt notes, a chess player will, e.g., rely on memory and imagination (Devitt 2006a, p. 19). But a chess player's making of a licit chess move is a paradigm of knowledge in action: knowledge of the rule(s) that sanction the moves is applied in the making of these moves. For as Ryle once observed, it not what a player 'does in his head or with his tongue, but what he does on the board that shows whether or not he knows the rules in the executive way of being able to apply them.' (Ryle 1949, p. 41). It may be that it is in connection with rules in particular, whose function it is to impose normative constraints on action and thought, that the central criterion for knowledge of them is heeding the rules in practice, not being able to cite them. And so as Crispin Wright suggests, even students who lacked the vocabulary to explicitly articulate chess rules would naturally be taken to know them, and exploit that knowledge in play, as long as they play according to the rules[21]:

> It does not happen, but might, that small children could learn to play chess long before they could learn speech… And then, remarkably, some of them acquire the ability to play not merely legally but well, responding with subtlety and inventiveness to board configurations which they have never encountered before….It is difficult…to see how such insight is supposed to function if not essentially as a faculty of *inference*: inference which goes to work on premises including, *inter alia*, the rules determining the powers of movement and capture of the various types of piece.(Wright, 1986a, p.23).

[21] Wright considers this thought experiment in context of a discussion of semantic cognitivism, the view that speakers implicitly know *semantic* rules for the languages they know. Although he is ultimately sceptical of cognitivism, the chess illustration is part of Wright's attempt to show that resistance to the notion of implicit knowledge is not a good ground for it. Wright's own grounds for rejecting cognitivism relate to the idea that, because linguistic competence is always deployed at the level of *sentence* use and interpretation, there can no behavioural evidence that motivates imputing knowledge of a semantic theory, that is not just as well accounted for by postulating only knowledge of *sentence* meaning. I discuss this objection in Waights Hickman unpublished d; see also Weiss 2004.

In appealing to the notion of inference, Wright here identifies what I think is a crucial motivation for seeing the players' moves as knowledge-guided: the behaviour strikes us as *rational*, guided by reasons and, moreover, reasons that are facts (or rules that obtain)—reasons that are seen not only to motivate the moves, but are fit to justify them. This is an important dimension along which the heeding of chess rules in chess playing is analogous to the heeding of syntactic and semantic rules in language use, from a cognitivist perspective. From this perspective speakers' conformity with linguistic rules is a reflection of their linguistic knowledge, which is implicitly exploited in 'computations similar to straight deduction' to arrive at conscious knowledge of facts about meaning, ambiguity and form. (Chomsky, 1986, p.270). As Fodor explains:

> According to generative linguistic theory, your ability to detect syntactic ambiguities, distinguish well-formedness from ungrammaticality, respond selectively to the noun-phrase that has been topicalized, and so forth are to be explained by reference to *what is entailed* by the grammar that you learned when you learned your language ... You can spot the ambiguity of 'they are flying planes' because, so the story goes, (i) you have learned the grammar of English, and (ii) it follows—deductively—from what you have learned that 'they are flying planes' has two well-formed parsings. (Fodor 1983, p.7)

The imputation of linguistic knowledge makes *rational* sense of speakers' linguistic judgements, interpretations, and uses of language; the primary sense in which syntactic and semantic rules are purported to *explain* linguistic judgements etc. is not, *pace* Levin, causal, but *rational*. From Devitt's perspective, it must seem either that this rationalisation is gratuitous—that our sophisticated linguistic abilities do not call for a rationalising explanation any more than the bee's dance apparently does—or that crediting speakers with a non-propositional linguistic skill, or knowledge *how* to use the language, suffices to present their language use, as it were, in a rational light. I address these ideas in turn.

Whence the pressure to seek a rationalising explanation of language use and linguistic judgements? As I suggested above, the use of language is intentional and purposive, and a widespread view going back to Davidson and Anscombe in the philosophy of action is that intentional actions are just those that are guided by reasons, with some uninteresting exceptions: voluntary actions that are non-purposive, but merely expressive of emotion, such as kicking the door in a fit of temper or throwing up one's arms in frustration[22]. As a rule, then, intentional actions should be explicable in terms of the agent's reasons, and language use is paradigmatically intentional and purposive behaviour[23]—it is essential to pursuit of perhaps the majority of our goals. And so, as Dummett maintained, an account of that behaviour should show it to be conscious, reason guided activity: 'rational activity on the part of creatures to whom can be ascribed *intention* and *purpose*' (Dummett 1993, p.104). This felt requirement contrasts starkly with our folk intuitions about the behaviour of the dancing bee or the diving kingfisher, and this is significant for the question whether linguistic behaviour need be thought to exploit knowledge of language. This is precisely because we are in a position to see it as rational if the knowledge is imputed, and seemingly not otherwise (Lepore 1997, Heck 2007, Matthews 2006a).

Devitt might well contest the view that appealing to known grammatical or semantic principles in the way Fodor indicates is the *only* way to provide a rational explanation of linguistic behaviour and judgements. He might take reference to the agent's *knowledge-how* to be sufficient to rationalise linguistic behaviour[24]. But if knowledge-how is assumed to be

[22] E.g. Raz 1999, ch. 2; Mele 1992, ch.6. The view goes back to Davidson 1963 and Anscombe 1957.

[23] Importantly, it is intentional *as such*. It is not only that our making promises, communicating knowledge, raising questions and so forth *by means* of language is intentional behaviour, our choice of linguistic means for those ends is also intentional. It is along this dimension that the intentional character of language use is rationally explicable by reference to speakers' linguistic knowledge. This is an issue I discuss at greater length elsewhere (Waights Hickman unpublished d); see also Heck 2007.

[24] Hornsby (2005) advances precisely this view, but I am unable to make out her argument for it (see pp. 125-8). She stresses the *complexity* of our (allegedly non-propositional) linguistic abilities, observes that

non-propositional knowledge, I do not think this position is defensible, simply because a reason is the kind of object apt to figure as a premise in practical and theoretical *reasoning*, i.e. it must have propositional form[25]. A reason or ground for action (/thought, feeling) is a *proposition or fact* that, from some perspective (perhaps the agent's, perhaps objectively) speaks in favour of that action, or explains why the action is reasonable or right. When it comes to action *explanation* it seems clear that the explanatory reason must be a fact or truth (as falsehoods are not explanatory), and recently there has been growing support for the idea that it must be a fact or truth that the agent not merely believes but *knows*[26]. And so the kinds of attitudes that are apt to rationalise intentional action are necessarily *propositional* attitudes—in the linguistic case, either linguistic beliefs or (propositional) linguistic knowledge.

Why will it not do to refer to the agent's knowledge-*how*? Isn't a speaker justified in, e.g., using a string of words composing a sentence S in order to say that *p*, simply in virtue of her ability to use the words of the language correctly? The question here is what *reasons* explain the speaker's intentional use of language. If her linguistic abilities consist in or are underpinned by propositional linguistic knowledge, then we can cite as among her reasons facts about the meaning of the constituent expressions of S and the semantic significance

semantic (practical) knowledge is manifested in concert with myriad other attitudes and goals, and that speaking is not something that 'just happens' but is subject to rational control; and she praises semantic theories for explaining 'how the users of the language are able to do all of the things that their semantic knowledge equips them to do' (*ibid*, p.126)—but she insist that semantic knowledge is *not* knowledge of such a theory. Although the features of linguistic competence to which Hornsby draws attention in her discussion are certainly ones we would expect of a distinctively rational competence, I cannot see any motivation for the claim that her non-cognitivist view (on which semantic competence is practical knowledge) is competent to *account* for the intrinsically rational character of language use.

[25] See Alvarez 2010, ch.2; Hyman 2015, ch. 6; and Raz 1999, ch. 1. Alvarez and Raz claim that a reason must be a fact or truth, it cannot be a believed falsehood; Dancy (2000, 2011) argues that a reason is a ground, whether true or false; Hyman (2015) argues that what an agent adopts as reason or premise can be a falsehood but a reason competent to *explain* action must be a fact or truth.

[26] Hyman 2015; Hornsby 2007; Hawthorne and Stanley 2008. At the very least, the reason(s) that guide an agent must be *believed* propositions (Dancy 2000). An agent cannot be motivated by a reason that is available but of which she is in no sense aware.

of combining them to form S, which make S an appropriate vehicle for conveying that *p* to her interlocutor. But if her abilities are, as it were, brute, without any propositional content or underpinning, there are no *reasons* to cite in explanation here. Note that the putative fact that the agent has (non-propositional) knowledge how to use the language cannot be *her reason* for using S to communicate that *p*: what motivates the adoption of particular linguistic expressions for her purposes must facts about those expressions, not facts about *her*.

Devitt's use of illustrations involving bird, insect and machine behaviour to draw lessons about human language mastery is misleading because our pre-theoretical intuition in these cases is that the organism or system's production of an "output" is non-intentional and thus not subject to a rationalising explanation. Pre-theoretically the competent agents do not strike us, in Dummett's words, as 'creatures to whom can be ascribed *intention* and *purpose*.' That human linguistic behaviour is *disanalogous* to Devitt's favourite cases in this respect is absolutely critical. Behaviour that does not strike us as intentional calls for no rationalising explanation, i.e. no explanation of the kind that knowledge and belief are suited to provide, and for this very reason Devitt's contention that the dancing bee and the pocket calculator are *ignorant* of rules governing their outputs seems innocent. Language use is not like that: *prima facie* a rationalising explanation is called for, an explanation that the knowledge imputed by cognitivists is suited to provide. And so in this case we *are*, *pace* Devitt, under independent pressure to see the sophistication of the "output" as a reflection of the agent's knowledge or belief states.

Conclusion

Many commentators have pointed out that an agent's conformity with a system of rules does not *entail* that the conformity issues from knowledge, a point with which linguistic cognitivists will readily agree, advancing cognitivism as an explanatory hypothesis accounting for a speaker's ability to understand indefinitely many novel sentences. Devitt and Soames put forward what is ostensibly a more serious objection, challenging the cogency of the cognitivist's explanatory aspiration, arguing that it is based on a conflation of the domains of language and linguistic understanding, and the respective theories thereof.

I have argued that their form of objection does not hold water. First, the separation of the domains of external language and linguistic competence is not one that cognitivists have ignored. From the perspective of Chomsky's internalism, linguistics is concerned exclusively with linguistic competence, linguistic properties having no genuine mind-independent existence; and on a view on which syntactic and semantic properties independently belong to external linguistic symbols, the cognitivist can hold nevertheless that a speaker who is competent with the language must know the rules in virtue of which linguistic expressions possess those properties. Second, the supposed challenge to cognitivism is in fact directed against a conception on which grammars and semantic theories purport to recapitulate unconscious linguistic processing, a view quite distinct from cognitivism—which sees these as the contents or objects of speakers' *knowledge*—and one explicitly eschewed by Chomsky.

Finally, I have argued that the examples of output-sophistication consistent with ignorance to which Devitt appeals as a model for thinking about language production are not suitable for that purpose, as they are not models of intentional behaviour on the part of rational agents. The rational nature of language use provides independent motivation for cognitivism, given that its is *rationally* that reference to (speakers' knowledge of) linguistic rules purports to explain their linguistic behaviour. Explanation referring to a speaker's knowledge is a reasons-explanation, and this is a form of explanation that is called for where the behaviour is, as in the case of language use, intentional and purposive.

(Implicit) Knowledge, Reasons and Semantic Understanding

Introduction

This paper exploits recent work in epistemology in support of semantic cognitivism, the view that linguistic competence is partly constituted by (implicit, propositional) semantic knowledge. On the one hand, there is now significant agreement that only knowledge of a fact could permit an agent to rationally respond to it: only if an agent knows a fact can it be a *reason on which* he ϕs (acts, reasons, judges) or by which he is guided. This is not to say that ignorance in every case defeats rationality, but that an agent cannot be rationally responsive to *the facts* if he is ignorant of them[1]. If this is correct, we have a criterion for knowledge that is in principle clear and independent of the criterion of assertion or explicit statement; in other words, we have the basis of a test for ascriptions of implicit knowledge, knowledge that the subject cannot (in advance of instruction or priming) put into words. On the other hand, there is growing support for the view that knowledge plays a normative role in various domains, including intentional action. In particular, Hawthorne and Stanley (2008) maintain that *ceteris paribus*, intentional actions that are premised on a reason that is not known *violate a norm* of practical rationality.

In what follows, I argue from these recently highlighted normative and constitutive roles of knowledge in practical rationality, to the conclusion that speakers possess semantic

[1] If it is allowed that reasons need not always be facts (even if successful reasons-*explanations* are always factive), then one can rationally operate with reasons as premises even when the premises are false or unknown (Hyman 2011). Thus the inference from *p* to *q* may be thought blameless *qua* inference even if the agent falsely believes that *p*, provided that e.g. the conditional $p \rightarrow q$ is true. Similarly if an agent ϕ-s on grounds that *p*, taking himself to know that *p*, he may be thought (partly) practically rational even if *p* is false, if *p* would have, had it been true, provided grounds for ϕ-ing.

knowledge and deploy this knowledge in language use and comprehension. In particular, I will argue that:

(1) Non-cognitivism has the consequence that language use is systematically irrational.

(2) Non-cognitivism is inconsistent with the rational sensitivity to semantic facts routinely exhibited by speakers in speech and understanding.

1. Linguistic competence and implicit knowledge theses

Cognitivism is the view that linguistic competence partly constituted by propositional knowledge of the syntactic and semantic rules of the languages at one's command. On this view natural language grammars and semantic theories purport to make explicit the knowledge (implicitly or tacitly) possessed and drawn upon in language use and comprehension. My focus in this paper will be on *semantic* cognitivism, the view that speakers (a) possess knowledge of semantic facts or rules and (b) exploit this knowledge in speech and comprehension. This knowledge is assumed to be of a particular, indefinitely productive sort, *viz*. knowledge of the meaning of simple expressions (words) and knowledge of the semantic significance of combining those expressions in different ways. Such knowledge enables a speaker to know, for any sentence S of the language with meaning *M*, that S means *M*.

The primary motivation for cognitivism is its explanatory power with respect to the creative aspect of linguistic competence, i.e. a speaker's ability to compose and understand indefinitely many novel sentences of her language. But the more relevant

virtue of cognitivism for this paper is its vindication of a conception of language use as distinctively and intrinsically *rational* activity: cognitivism sees those semantic truths or principles which do in fact license the interpretation of a sentence S as meaning *M*, as premises or *reasons* on which a competent speaker's interpretations are implicitly based (Heck 2006; Lepore 1997; Weiss 2004). Claiming this advantage for cognitivism requires that the linguistic knowledge ascribed to speakers be of a certain sort, or rather, that it be possessed in a certain way: so as to be accessible to reason. As I shall explain shortly, in current terminology, this requires taking speakers to know semantic rules *implicitly* and not merely *tacitly*.

The knowledge of meaning that ordinary speakers possess cannot, in general, be explicit. The full formulation of semantic principles that are rich enough to permit the meaning of any sentence to be derived—principles showing how what speakers do achieve can rationally be achieved—is given by a compositional semantic theory, and requires intellectual resources that ordinary speakers do not possess. Cognitivism therefore entails the contentious claim that speakers possess knowledge whose content they are not only typically unable to formulate themselves, but which they may not recognise an apt formulation of. In this way the postulated semantic knowledge differs from the familiar phenomenon of a propositional (or other) attitude that is implicit only in the sense that one does not recognise having it, e.g. an implicit sexist belief which one can well articulate or recognise an apt articulation of, but disclaims having. Nevertheless, it should be borne in mind that it is precisely on grounds of how one acts and not on grounds of a verbal articulation that such implicit attitudes are routinely and un-problematically ascribed.

Because the semantic knowledge postulated by cognitivism is implicit in the deeper sense of being not merely unarticulated but (in advance of formal instruction) un-articulable by its possessor, cognitivism ascribes knowledge of semantic rules to speakers even though:

- Ordinary speakers may not consciously understand or even recognise formulations of these rules.
- If possession of a concept requires mastery of the corresponding term, then ordinary speakers lack many of the concepts employed in articulation of the rules.
- The (formulations of the) rules play no pedagogical or explicit normative role in the teaching and moderation of linguistic practice.
- The rules are not cited in explicit explanation or justification of linguistic practice.

This is exactly the profile we should expect for states of implicit knowledge. For example, it is evident that if pre-verbal children and animals can be credited with some implicit knowledge, there will be the same pressure to resit crediting them with the relevant concepts, and explicit statements of the knowledge will be similarly irrelevant to guiding their behaviour. Just as with implicit semantic knowledge, explicit statements of what is known will play a role in justifying or explaining behaviour only from the third-person standpoint. Nevertheless, many philosophers have taken these violations of standard criteria for (explicit!) knowledge ascription to demonstrate that semantic cognitivism is absurd[2]. Others, in partial sympathy and perhaps wishing to avoid the

[2] E.g. Baker and Hacker 1984; Hanfling 1980; Foster 1975; Devitt 2006; Schiffer 1987; Hornsby 2005.

charge of absurdity, have taken it to show that semantic "knowledge" is not a genuine or full propositional attitude. Gareth Evans urged that to assimilate semantic knowledge to ordinary propositional attitudes is a "mysterious and confused position" (Evans 1981, p. 134). Instead, he proposed to understand this knowledge, which he called *tacit* knowledge, as a set of dispositional states of the agent corresponding to axioms of a semantic theory, states that play a causal-explanatory role in an account of how semantic competence is acquired, lost and revised (*ibid*, 122-130). Developing on Evans, Davies defines tacit knowledge of a theory T as a causal explanatory-structure that mirrors the derivational structure of T, consisting not in a set of dispositions but in the categorical basis of those dispositions (Davies 1982, 1987, 1989)[3].

This conception of semantic knowledge purports to secure its explanatory credentials with respect to linguistic creativity as well as its amenability to empirical investigation, but it gives no support to a rationalistic account of linguistic competence and its exercise (Weiss 2004). According to Evans, tacit knowledge differs crucially from a genuine propositional attitude in that the latter but not the former is 'inferentially integrated' in reasoning with other beliefs, and apt to combine with the agent's desires in the service of indefinitely many projects (Evans 1981, p. 132; Cf. Stich 1978). Thus a genuine belief (a genuine propositional attitude) is apt to be expressed in indefinitely heterogeneous ways; states of tacit knowledge, on the other hand, are deployed and manifest in only one-pattern processes:

> Possession of tacit knowledge is exclusively manifested in speaking and understanding a language; the information is not even potentially at the service of any other project of the agent, nor can it interact with any other beliefs of the agent… There is no question of regarding the information being brought by the subject to bear upon speech and interpretation in rational processes of thought. (Evans 1981, pp. 133-4)

[3] A similar view of semantic knowledge as 'sub-doxastic' is advanced by Peacocke (1986). See Wright 1986, Davies 1987, Knowles 2000 and Weiss 2004 for critical discussion of this version of cognitivism.

As Crispin Wright explains, semantic knowledge is not a full propositional attitude by Evans's criterion, and counts as merely tacit, because the postulated "knowledge of a meaning theoretic axiom would seem to be harnessed to the single project of forming beliefs about the proper content of sentences which contain the expression, or exemplify the mode of construction, which it concerns." Although the beliefs that might be formed are indefinitely numerous they are always of the same kind—beliefs about sentence meaning. Moreover, Wright asks rhetorically, "what is supposed to be the role of desire? What is the (implicit) desire which explains why the subject puts his semantic axiomatic beliefs to just this use, and what are the different uses to which they might be put if his desires were different?" (Wright 1986, p. 34).

It is tempting to reply by urging that Evans' criterion is not a genuine condition on an attitude qualifying as intentional or propositional, and that it is inappropriate to cases of implicit knowledge. One might think that it is precisely in virtue of being explicit that a propositional attitude becomes serviceable in indefinitely many projects, and that it is unsurprising that knowledge or belief states that are unarticulated and unavailable to consciousness should not integrate with other mental states in the way explicit attitudes do (Knowles 2000, p. 338-9; Weiss 2004, pp.78-9.). But it is precisely the connection between reason and knowledge that I wish to exploit in support of cognitivism here, and so it would be incongruous to reject Evans' integration criterion. The arguments of this paper shall show that the criterion is simply met by semantic knowledge—though there are of course *some* limitations on the potential for integration of semantic knowledge with other aspects of our mental life, arising from its being implicit rather than explicit; more will be said about this in due course. The arguments below are consistent with the view that Evans's criterion is a constitutive condition of genuine propositional attitudes

(Rattan 2002), for I shall be arguing that semantic knowledge is not merely tacit but *implicit,* being accessible to reasoning and implicated in *rationalising* explanations of linguistic behaviour.

Cognitivism in any form has been much controverted; two popular alternatives are a 'functional role' conception of semantic understanding and the view that semantic competence is a species of know-how. On the first view, linguistic competence is whatever (sub-personal) state of the agent reliably issues in correct mentalese translations in response to the tokenings of sentences of the public language[4]. According to the second view, semantic understanding is a complex ability or disposition of the agent, assumed to have no propositional content[5]. Both positions deny that semantic knowledge of the sort described is exploited in the production or understanding of language. Even if speakers have some semantic knowledge (e.g. knowledge of what sentences of their language mean) it is superfluous to language use and comprehension; speech and comprehension are on these views direct achievements, unmediated by inferences involving semantic knowledge. As Schiffer puts it:

> When we understand the utterance of a sentence we do not first come to the belief that it means such-and-such and then have that as our basis for thinking that the utterer was saying such-and-such.(1987, p. 262).

Hornsby, defending the idea that semantic competence is knowledge-*how*, concurs: speech is on her view a matter of directly "voicing one's thoughts," and understanding is a matter of directly "hear[ing] the meaning in the words" (Hornsby 2005, p. 111-112).

2. Linguistic Rationality and the Knowledge Norm

[4] Fodor 1989, Schiffer 1987.

[5] This has always been the default alternative put forward by critics of linguistic cognitivism, e.g. Stich 1971; Devitt 2006. It is given a more positive elaboration and defence in Hornsby 2005.

The last two decades of work in epistemology have seen a gradual paradigm change. Williamson (2000) has persuaded many philosophers that a reductive analysis of knowledge is not to be had, and that knowledge is better treated as the primary notion in terms of which other questions in epistemology can be illuminated; he also recognises the centrality of knowledge in the explanation of intentional action, and argues that it plays a normative role in relation to assertion and belief. John Hyman (1999; 2015) has likewise argued that reductive conceptions of knowledge as belief with a special accreditation or guarantee are doomed to failure; on his account, this is because knowledge is an ability, "so if we want to understand what knowledge is, and why we value it, we need to ask what it is an ability to do, instead of how it can be certified or acquired" (2015, preface ix). The questions that naturally arise concerning knowledge are now different; as Hyman puts it, "we need to think prospectively not retrospectively, about how knowledge is applied, employed, expressed, in the infinitely varied circumstances of human life." (*ibid*).

This shift is congenial to the investigation of implicit knowledge generally; especially relevant for those interested in implicit semantic knowledge is the growing interest in how knowledge relates to rationality. According to Hawthorne and Stanley, knowledge provides a *norm* of practical rationality (Stanley and Hawthorne 2008, p. 577), one constraining the adoption of reasons for action. On their view, wherever an agent's preferred action in a choice situation is dependent on whether or not *p*, it is appropriate for her to treat *p* as a reason to ϕ iff she knows that *p*. The proposal motivates our common folk appraisals of action in epistemic terms (*ibid*, p.572). For example, just as I would be criticised for believing or asserting that Elena will arrive late at the conference

on grounds that she missed her flight if I did not that she had, I would be criticised for *acting* on the premise that Elena missed her flight if I did not know it, e.g., if I cancelled her place at dinner on those grounds—even if Elena did in fact miss her flight.

As with any norm, there are exceptions or exculpations for violations, for example there may be situations in which one is forced to choose between ϕ-ing and φ-ing, though there are no known premises favouring one course of action over the other. Stanley and Hawthorne discuss and respond to concerns that the norm is too demanding at some length, but it may be that ultimately one needs to be optimistic—as I am—about how much knowledge we have and employ in daily life in order to find their proposal persuasive. I am sceptical that the principle is demanding *enough*, i.e. that knowledge of a proposition is *sufficient* to render it practically rational (or 'appropriate by the standards of practical rationality') to adopt as one's reason for acting, as Hawthorne and Stanley contend[6]. But I believe they are correct that knowledge is normatively relevant to the appraisal of intentional actions; in particular, that *ceteris paribus*, an intentional action performed on an unknown premise is subject to rational criticism on those grounds.

Now consider linguistic behaviour. The capacity to use and understand language is commonly regarded as the hallmark of our rational nature, if not its condition. Dummett famously maintained that language use is the paradigm of rational activity, and that this imposes a constraint on theories of meaning (Dummett 1993, p.104):

[6] For their account does not recognise sensitivity to the *objective relevance* of a candidate reason *p* to the question whether or not to ϕ as a criterion of practical rationality; the only kind of relevance their proposal incorporates is *subjective*, i.e. whether or not the agent *prefers* ϕ-ing conditional on whether of not *p*. By my (non-Humean) lights the former *is* a criterion of practical rationality, and an agent can fail to satisfy it whilst obeying the knowledge norm. Just as it cannot be theoretically rational to infer any proposition from *p* as long as one knows that *p*, it cannot be practically rational to ϕ on grounds that *p*, for any value of ϕ, as long as one knows that *p*. For example, knowledge that it is Tuesday does not make it practically rational to attempt the destruction of the whole world rather than scratching my finger, even if I prefer the former to the latter conditional on it being Tuesday.

> Any adequate philosophical account of language must describe it as a rational activity on the part of creatures to whom can be ascribed *intention* and *purpose*. The use of language is, indeed, the primary manifestation of our rationality: it is *the* rational activity *par excellence.*

Philosophers who share Dummett's conviction have seen a connection between the rationality of language use, and linguistic knowledge. Heck (2006) maintains that, because the utterance of a sentence is an intentional act, it is subject to a reasons-explanation, and this explanation can be given only in terms of knowledge of the sentence's meaning. Lepore (1997) likewise argues that non-cognitivist accounts of understanding must be rejected because they leave speakers' interpretations of utterances unrationalized, that is, they fail to identify *reasons* that ground them[7]. But if Hawthorne and Stanley are correct in maintaining that acting intentionally on unknown premises contravenes a norm of practical rationality, then non-cognitivism has a greater cost than leaving language use unrationalized, i.e. failing to vindicate the idea that mastery of a language is a distinctively and intrinsically rational competence. Language use will not turn out to be simply non-rational. Instead, it must be systematically *irrational* along some dimension, for it systematically *violates* the knowledge norm. Given the norm, one will under normal circumstances be subject to rational criticism for intentionally uttering S in order to assert that *p* despite ignorance of the meaning of S, and thus ignorance concerning the suitability of S to achieve one's communicative goal. One may well *believe* that S is appropriate for this purpose, and believe truly, but by the lights of Hawthorne and Stanley's norm, this will not suffice to render the utterance rational: "When someone acts on a belief that does not amount to knowledge, she violates the norm, and is thus subject to criticism."(Hawthorne and Stanley 2008, p.574).

[7] See also Campbell 1982, McDowell 1988 and Weiss 2004.

The intentional character of linguistic behaviour is a critical motivation for entering this objection to non-cognitivism. For a popular view in the philosophy of action is that an action counts as intentional just in case it is reason-guided, with some arguably uninteresting exceptions: voluntary actions that are non-purposive, but merely expressive of emotion, such as kicking the door in a fit of temper or throwing up one's arms in frustration. So if we combine this idea with the normative claim that agents ought (in general) only to act on reasons that are *known* facts or truths, the upshot is that intentionally ϕ-ing without any known reason to ϕ is rationally criticisable. There may be certain linguistic acts that are not plainly intentional or purposive, such as mumbling absent-mindedly as one one searches for the car keys, and here there is no clear demand for a rationalising explanation of the agent's utterances. But in the normal case, language use is both intentional and purposive, and cannot be assimilated to those intentional acts that merely express emotion. Language use must therefore be reason-guided, and subject to Stanley and Hawthorne's norm. We use language in order to say things (and thereby, e.g. make promises, issue reproaches, raise doubts, etc.)—it is a means to an end, and what could be more suspect by the lights of practical rationality than employing some means with an end in mind, without being apprised of reasons to think that means appropriate?

This brings into view an important commitment behind our charge, namely the idea that in using language, one's choice of communicative *means*—i.e. one's choice of words—as well as one's communicative act, is intentional in character. Jennifer Hornsby (2005) has denied this, and argued that a speaker's intentionally saying that *p* is not, accordingly, even potentially rationalised by semantic facts. (It may be rationalised, *inter alia,* by knowledge that *p,* and aspects of the conversational context). So the

intentional character of language use does not offer immediate motivation for semantic cognitivism: it must be shown that it is intentional in an aspect that semantic facts can rationalise. I take up this challenge below, in the course of defending the contention that speakers exhibit rational sensitivity to semantic facts that only knowledge affords.

3. Linguistic Rationality and the Knowledge View of Reasons

The 'knowledge view' of reasons is a descriptive thesis concerning the connection between reasons and knowledge in the *ex-post* explanation of intentional action, not its justification. It is the view that an action-explanation that refers to the agent's reason is an explanation that refers to *knowledge* upon which the agent acts; a claim originally argued for by Unger (1975) but developed in more detail by Hornsby (2007a, 2007b) and Hyman (1999, 2011, 2015). Indeed, it might be better labelled the 'knowledge view of reasons explanation'. For arguably, one can have a proposition *p* as a reason or ground for action without knowing that *p*; *p* may even be known or believed to be false yet treated as a reason or premise, as when one argues for *reductio*[8]. However, these are not cases in which the relevant proposition can *explain* action. If Emma goes to a restaurant on the ground or premise that it serves vegan food, but it turns out that the restaurant does not serve vegan food, then *that the restaurant serves vegan food* cannot possibly explain Emma's action. What explains Emma's behaviour, rather, is the fact that *she believed that* the restaurant served vegan food. This reflects a general point about explanations, namely that explanations are factive: only a fact or truth can be

[8] The import of these cases is discussed in more detail in Hyman (2011). Some philosophers, however, hold that reasons must be facts or truths (Alvarez 2010, Raz 1999), and accordingly treat cases where the agent acts on (putative) grounds that *p* although *p* is false as cases of acting on an *apparent* or *putative* reason; that is, cases in which the agent takes herself to have a reason but does not in fact have one. I am sympathetic to this stronger view, but the argument of this paper does not depend upon excluding the possibility that reasons can sometimes be false beliefs or premises. The important commitment here is to the idea that the reasons figuring in reasons-*explanations* of action must be known facts or truths.

explanatory, and in the case of Emma's false belief, it is *the fact that she held it* that explains her behaviour, not the (false) belief itself. Reasons explanations, explanations in which the explanans is the agent's reason or ground, are no exception to factivity: if the agent's ground or reason is itself to serve as explanans, it cannot be a falsehood.

The knowledge view of reasons does not, however, amount to the fairly uncontroversial claim that an explanatory reason must be a fact or truth: it is the claim that it must be a fact or truth *that the agent knows*. On this view, knowledge that *p* is a condition on the fact that *p* guiding an agent, and accordingly explaining the action *ex-post*. And so, as Hymans contends, "a standard because-explanation of [intentional] action, thought or feeling, i.e. one in which the explanatory clause expresses or purports to express the agent's ground, implies that the agent knows the explanans." (2015, p.152). If we explain behaviour by citing a fact to which an agent is sensitive—a fact that is a reason or ground for their behaviour—the agent must be credited with knowledge of that fact.

Proponents of the knowledge view of reasons have primarily been interested in demonstrating that only knowledge allows an agent to ϕ for reasons that are facts or truths, where the default alternative is *true belief*; what matters for my purposes here is that *nothing weaker than knowledge* is sufficient. Hyman illustrates this by drawing an analogy between responsiveness to facts and responsiveness to things:

> One cannot be guided by a fact one does not know, any more than one can follow a guide one cannot see. If the traveller sees the guide taking the left fork and follows him, then he is guided by the guide; but if he hallucinates him taking the left fork, and takes it himself for that reason, then he is not guided by the guide, even if the hallucination happens to be true. (Hyman 2015, pp.208-209).

Only if we see the guide can we follow *him,* and only if we know the facts can we respond to *them*. The contrast between true belief and knowledge with respect to

reasons-guidance can also be seen in cases held to demonstrate that knowledge cannot be *reduced to* justified true belief. To give just one example from Hornsby (2007a, p.94):

> Edmund has just asked his mother whether the ice on the pond is thick enough to skate. His mother tells him that it is too thin, and indeed it is too thin. However, unbeknownst to Edmund, his mother believes that the ice is thick enough to skate, but is trying to trick Edmund into thinking that the ice is too thin. On the basis of his mother's testimony, Edmund believes that the ice is too thin. And because he believes this he stays off the ice.

The fact that the ice is too thin is not the reason for Edmund's staying off the ice; it is not a fact by which he was guided, and which can in turn explain his behaviour. For intuitively, Edmund cannot come to know that fact through the insincere testimony of his mother, and because he lacks knowledge, that fact cannot be one he responds to.

Consider again linguistic behaviour. Are speakers 'guided by' semantic facts? Do they exhibit rational *sensitivity* to these, of the kind that only knowledge can afford? I submit that they do. Of course, the greatest challenge where a reason is implicit in our sense—where it is a reason that the agent cannot articulate consciously or put into words—is showing that reasons are not merely available as rationalisers, but do also *motivate* the behaviour in question, and thus explain it. That is, there must be evidence that an available reason was *the agent's* reason. Behavioural evidence for reasons-guidance by semantic facts must therefore go beyond conformity with those facts: we need signs that agents do in some sense think about, take into consideration, and deliberately respond to semantic facts. I believe the cases discussed below show that these signs are in evidence, and that Hornsby's charge from §1 can be answered.

Hornsby (2005) maintains that speech is intentional under a description in which linguistic meaning is *assumed*; an agent performs an intentional act when she *says* that *p*,

but "there need be nothing such that she intentionally does it and says that *p* by doing it". (*ibid*, p. 118). That is, it is the speech-act that is intentional, not the adoption of linguistic expressions to perform it. If this is correct, then it would seem to follow that facts about meaning are not even suited to be reasons for language use. I do not think this conclusion follows, but more importantly, Hornsby is certainly wrong about the (non)intentional character of language use as such. I defend these points in turn.

The irrelevance of semantic facts to speech-rationalisation does not follow from the assumption that only *sayings* are intentional, for meaning itself can be the topic of an utterance. Consider the following example from Weiss (2004), designed to show that it is not always explanatorily sufficient to credit speakers with knowledge of *sentence* meaning. A speaker utters the sentence 'the apple is red,' in order to explain the meaning of the word 'red'. Ascribing knowledge of what the sentence means will not suffice to make rational sense of the utterance; knowledge of what the *word* 'red' means must combine with awareness of the context to yield the judgement that this would be a suitable sentence, under the circumstances, to provide an explanation of meaning. As Weiss explains,

> One might think that stating obvious truths by using the word ['red'] would do, but what has to be obvious is not the truth of the sentence but which truth the sentence asserts. So, for instance, no amount of utterances of obvious truths in a context involving only red discs and blue squares will convey that 'red' means *red* rather than, e.g., *circular*. And recognising that such a context is inappropriate for explaining the meaning of 'red' ill involve one's belief about that meaning. (Weiss 2004, p.80).

In connection with Hornsby's charge, the salient point about Weiss' case is that knowledge of word meaning is recruited in a rationalising explanation, even when we take the explanandum to be the speaker's *saying* that such-and-such (rather than her uttering a certain string in order to say that such-and-such). The reason is that, in this

case, what the speaker is attempting to convey is something about linguistic meaning itself.

But it is not only when language is used for its own illumination that rational sensitivity to meaning-facts is evidenced by speakers. I believe this sensitivity underlies *all* language use and comprehension, and would be impossible without it. Before turning to other cases in which speakers' rational sensitivity to semantic facts is salient, it is worth considering what language use would really be like if there were no role for such sensitivity to play in accounting for linguistic behaviour; that is, if Hornsby were correct that, generally speaking, language use is intentional only under what Heck (2006) calls a "propositional description"—described as *saying* that such-and-such—and not under a "verbal description", i.e. described as the utterance of a certain linguistic string. As Heck argues, the claim is not false *a priori*. We can imagine something like linguistic communication of which the claim is true. But it fails to describe *our actual* language use. For on the hypothesis that language use is non-intentional under a verbal description, it must be construed as a kind of noisy telepathy:

> What would that be like, to be a 'speaker' for whom 'speaking' was intentional under propositional descriptions, but not under verbal ones? Well, such an agent would be able to form an intention, say, to tell his wife that he loved her, and then that would just be something he did. Of course, he might make some sort of sound—perhaps a sound speakers of English would hear as "I love you"—but it would be no part of what he (intentionally) did that he uttered that sentence. Indeed, such an agent need not even be consciously aware of his uttering a sentence. A sentence would just be produced, much as our own lips and tongues just move when we speak. ... His wife need not consciously recognize the sentence that has been uttered to be able to recognize her husband as having … (for lack of a better word) said that he loves her. Communication between such agents would be like telepathy. (Heck 2006, p.13)

As Heck goes on to say, this plainly is not how things are for us. That language use is generally intentional also under a *verbal* description can be seen from a range of familiar features of our linguistic practice. For example, we recognize that we or someone else

has *failed* to say what was intended because the wrong word was used, whether due to semantic misunderstanding (e.g. taking 'livid' to mean *angry*), or to a slip of the tongue. Likewise, we know that our understanding of what another has said can be compromised if we are unfamiliar with one of the words they have used. We also deliberate not only about what we should say but how we should say it, e.g. whether to use the active or passive voice (Heck 2006, pp. 13-14).

Such aspects of our linguistic practice as Heck draws attention do not only demonstrate the intentionality of speech under a verbal description. They are alien to the quasi-telepaths precisely because for them there is no *rational connection* between the utterance of a certain string and the communication of a given thought. And this rational dimension of speech is not only best accounted for by assuming semantic cognitivism; in the cases considered below, the role of semantic information in practical reasoning is, I maintain, simply evident on inspection. And if speakers exhibit rational sensitivity to semantic information, it must be information that they *know*.

Where speech and language comprehension is unhesitating and smooth, our sensitivity to and exploitation of semantic facts is liable to be phenomenologically invisible, making our experience of language use a little more like that of Heck's quasi-telepaths. For ordinarily we do not stop to consciously reflect on our own words or those of our interlocutors; we are interested in *what is being said*, and in reasons that may motivate what is being said; we are not focused on reasons for taking someone to have said what we think they have, or on reasons for thinking one has oneself succeeded in saying what one intended. Here Hornsby's metaphors of "hearing meaning" and "voicing thoughts" are likely to resonate. It is worth dwelling more on cases, then, in which the use and

comprehension of language is less straightforward, and where accordingly our choice of, and reflection upon, linguistic expressions is in its own right more conspicuously intentional.

Consider first a textbook case of lexical ambiguity:

(a) There's a bug in the room.

Resolution of this and other lexical ambiguities is plausibly influenced by the subject's estimate of the likelihood of the proposition expressed by the sentence on each reading, in a way that does not immediately indicate a role for semantic facts about either word or sentence meaning in inferences: the competing representations of sentence meaning might just as well be delivered by a sub-personal mechanism. But of course, there is a linguistic reason, which the subject can indicate, for her deriving two different interpretations of the meaning of the whole sentence, namely the fact that the word 'bug' is ambiguous. Indeed, if the hearer is unsure about the meaning of the *sentence* (a) uttered in some context, she need only ask the speaker for clarification concerning the *one word* 'bug'—does she mean an insect or a piece of spyware? The question is made possible by the subject's recognition of ambiguity, and raising it with her interlocutor makes rational sense as a means to settle *sentence* meaning only because she is able to derive the meaning of the sentence from the meaning of its constituent expressions and their mode of combination. Why, once the intended interpretation of 'bug' is settled, does the hearer now reach one interpretation of (a) rather than another? —part of the (rationalising) *explanation* is in turn given by the interpretation assigned to the ambiguous word. And that the hearer can rightly conceive this as a reason for her interpretation

shows her sensitivity not just to the meaning of 'bug', but to the semantic contribution made by the word when combined with other expressions in that particular sentence.

Now consider a familiar kind of intentional code-switching, where two bilingual speakers attempt to exclude a third party from understanding what is said. Suppose two bilingual Spanish/English speakers, A and B are out with their friends, C and D; C and D are monolingual English speakers, conversing at the same table. A wants to meet B to buy C a surprise gift, and wants to make arrangements in secret. She could conduct the entire conversation in Spanish, but this would seem rude and alert C and D to a conspiracy. So instead she decides to code-switch with a single expression, and says to B:

(b) Let's meet at the market tomorrow, and look for a regalo.

I have already described many of the facts that have entered A's practical deliberations: her *reasons* for intentionally code-switching, e.g. she wants to communicate such-and-such exclusively to B, without alerting their neighbours to the attempt to exclude them from understanding. But one crucial reason for her linguistic act, is the fact that the Spanish word 'regalo' means *gift*. A exploits this fact in uttering a string that B will be able to interpret; B exploits it in the interpretation; both exploit C and D's ignorance. Note that A's semantic knowledge is integrated with her other beliefs and desires in the service of a particular goal. The knowledge that 'regalo' means *gift* is avowedly put to use only in a linguistic act, but the reasoning behind it is quite complex. Knowledge that 'regalo' means *gift* is useful in this context only in combination with knowledge that completing the English sentence with an expression with this semantic value will produce a sentence meaning *let's meet at the market tomorrow, and look for a gift*; knowledge

that C and D's understanding of (b) with the switched expression will be compromised by ignorance of the fact that 'regalo' has this meaning; knowledge that B's understanding will not be so compromised. And all of this is relevant only given A's desire to exclude C and D from understanding, but without causing offence or suspicion.

Deliberations about which expressions to adopt for one's purposes are of course familiar enough even within one language, especially (but not only) if it is not our mother tongue. Naturally, sensitivity to more than semantic properties of expressions will typically be relevant to our reasoning. For one may be wondering how best to say that *p*, but also to implicate that *q*, or to achieve or avoid some further perlocutionary effect. Thus a speaker can choose the more familiar of two semantically close expressions (e.g. 'off-putting'/ 'rebarbative'; 'wander'/ 'peregrinate') in order to make their statement accessible to a wider audience; or they may choose the less common term in order to impress or exclude. Also familiar is deliberate adjustment of one's language use in some context in order to minimise offence or other negative perlocutionary effects. Thus a professor may take care to refer to female undergraduates as 'young women' rather than 'girls', because of the different social connotations of the two terms. Similarly, one may adopt 'miserly' in preference to 'niggardly' because the latter may be misinterpreted as a racial pejorative. In such cases the background reasoning makes crucial use of semantic knowledge, but the use of that knowledge is mediated by other beliefs about other properties of the candidate expressions, as well as beliefs concerning the context of utterance, and one's interlocutors. The semantic knowledge is also, to answer Wright's rhetorical question from §1, mediated by *desires*, for example the desire not to offend.

Finally, consider the disambiguation of attachment ambiguities, such as:

(c) John saw the man with binoculars.

(d) Young boys and girls like cats.

In (c), 'saw' can attach either to 'the man' or to 'with binoculars'; in (d) 'young' can attach to 'boys' or 'boys and girls'. There is no consensus yet concerning how attachment ambiguities are resolved in normal linguistic comprehension, for example concerning the nature and extent of the impact of discourse context on the favoured interpretation[9]. The important point for my purposes is that ordinary speakers recognise the ambiguity and can trace it to different ways of parsing the sentence, manifesting sensitivity to the semantic consequences of different ways of "attaching" the constituent expressions in (a) and (b). This is also reflected in the fact that a speaker will often make deliberate use of linguistic or pragmatic devices such as a comma, stress or pause in order to induce the intended interpretation in her interlocutor. These are all forms of intentional behaviour for which semantic facts provide rationalising explanations.

4. Linguistic Competence *vs* Meta-Linguistic Competence

Many more examples like those above can be generated easily; they are cases in which the speaker/hearer is naturally described as responding to reasons that are semantic facts, either in reflection on linguistic expressions in order to recover what is being said, or in deliberations concerning the expressions with which to convey something herself.

[9] Binder, Duffy and Reiner 2001; Binder and Morris 1995.

Following the knowledge view of reasons, this rational sensitivity to semantic facts is secured by knowledge alone. Cognitivism is then not simply recommended by its explanatory success concerning the creative aspect of linguistic competence, or by its potential to vindicate a basic intuition that language use is, as Dummett said, "rational activity *par excellence*". For I do not think that the rational exploitation of semantic information by speakers is an explanatory hypothesis: it is a familiar phenomenon. The cases of the sort I have described are part of our everyday life with words, they are not anomalous or contrived examples of reflection upon language that is only engaged in by linguists and philosophers. They exemplify the routine reliance of ordinary speakers on semantic reasons or premises, on semantic information which they take into consideration or account. And if one recognises this, one must recognise that, *contra* the anti-cognitivist, speakers do possess *and must regularly exploit* semantic knowledge in language use and comprehension.

My opponent can always reply that such cases do not and cannot demonstrate that *all* semantic knowledge is available to reasoning or that semantic knowledge is exploited in *all* linguistic communication. This is certainly correct. As far as the general accessibility to reason of semantic knowledge is concerned, cognitivism *per se* does not require all semantic knowledge to be integrated in reasoning; that is, one can coherently maintain that some semantic knowledge is implicit while some is merely tacit (Johnson 2007, 2004). As for the claim that semantic knowledge is rationally exploited in *all* linguistic communication, I concede that an argument from instances can never prove a general claim. But as has been said, the cases from which I argue are not fringe or anomalous examples of communicative language use, they are quite mundane. My own view is that the incorporation of semantic reasons in communication is more salient to us in these

cases and less in others only because some slight impediment or dilemma has slowed language processing down, drawing conscious attention to linguistic *means* and not just ends. If human computation in these cases were faster, the phenomenology of reasoning in light of semantic information would quickly disappear. This is really just to rehearse the familiar point that awareness of a conscious, deliberative process is not the condition of behaviour being genuinely reason-guided; most rational action involves no anterior conscious deliberation. And so I believe that smooth and unhesitating linguistic interactions too depend upon reasoning with semantic information, regardless of the phenomenology. One might test this claim by considering counterfactuals. For any normal case of unhesitating linguistic interaction in which the parties' choice of expressions is not in question or in focus, we can construct a counterfactual scenario in which it is brought into question or focus e.g. by one party's mishearing or being unfamiliar with one of the other's words.

Finally, my opponent might argue that in cases of linguistic comprehension, e.g. the resolution of ambiguities, reasoning about the semantic properties of expressions is only part of *meta-linguistic* competence, e.g. the *ex-post* judgement concerning the source of some sentence ambiguity, but that it is extraneous and irrelevant to the exercise of first-order linguistic competence, i.e. that it is irrelevant to language *comprehension*. The question whether meta-linguistic judgements, e.g. judgments concerning the ambiguity and grammaticality, and communicative language uses of language, issue from different cognitive sources is controversial, thus the importance of the distinction between linguistic competence and meta-linguistic competence for cognitivist arguments is unclear[10]. More importantly, in the ambiguity cases described above, semantic

[10] See Chomsky 1986, Devitt 2006b, 2010; Rey 2006; Collins 2008.

information does *not* only inform meta-linguistic judgements, but is essential to arriving at an interpretation. Thus in the first example, the hearer first queries the intended interpretation of 'bug'—motivated, if one likes, by the *meta*-linguistic judgment that 'bug' is ambiguous—but the information acquired then enables her to derive the meaning of the whole sentence, and here knowledge of the intended meaning of 'bug' is playing an essential role in comprehension. Finally, there can be no recourse to this line of argument where cases of language-*production* are concerned. Where a speaker reasons in light of semantic information in intentional code-switching, or in deliberating about the most accessible or least offensive way of conveying a message, semantic facts figure essentially in *ex-ante* deliberations concerning *how to say* what she wants to say, and cannot be treated as extraneous or supplementary to successful linguistic communication.

5. Inferential Integration and its Limits

If the foregoing is correct, then (at least a great deal of) semantic knowledge must be implicit, and accessible to reason. It is not, as Wright contends, merely that (putative, tacit) knowledge of meaning is brought to bear in forming beliefs about what *sentences* mean, and that *these* beliefs about sentence meaning are accessible to reason and at the service of whatever projects one may pursue by linguistic means. A speaker who deliberates about which of two expressions to adopt, or makes use of a stress or pause in order to forestall misunderstanding of a structurally ambiguous sentence, is displaying rational sensitivity to semantic information that goes deeper than sentence meaning. In linguistic communication, such information is one input, alongside e.g. facts concerning aspects of the spacio-temporal and discourse contexts, informing a speaker's choice of

expressions, given her intentions and desires; and likewise to the hearer's recovery of what her interlocutor is saying (implicating, etc.) by her words. Thus, for example, a speaker's choice to adopt the more esoteric of two synonymous expressions is informed by knowledge of meaning, but that knowledge is also mediated by beliefs about the audience's knowledge, expectations and interests, and combines with the speaker's intentions and desires with respect to that audience (e.g. to impress), in determining his ultimate choice. There is simply no limit to the potential information about his surroundings and his interlocutor that a speaker may take into account in determining not just what to say but how best to say it, and no limit to the kinds of intentions and desires he may be concerned to satisfy in doing so. Thus semantic knowledge is certainly not inferentially isolated, but is "at the service of many distinct projects, and …its influence on any project [is] mediated by other beliefs". (Evans 1981, p.134 quoted above).

Of course, it can be replied that semantic knowledge is only serviceable in the performance of and response to *linguistic* acts, i.e. is only only serviceable in whatever projects may be pursued by means of such acts, and thus that it does not exhibit full versatility and aptness for integration in the rest of our mental lives. In particular, this knowledge cannot be brought to bear by speakers in assessing or correcting a semantic theory—what cognitivists contend is no more than an explicit formulation of what the speaker already knows. As Stich (1978) objects, a speaker who (allegedly) implicitly knows a linguistic rule R will not infer that Chomsky is mistaken, given her explicit belief that if R, then Chomsky is mistaken. But *this* kind of limitation on the inferences that a speaker can draw from semantic knowledge and the tasks she can exploit it in should be utterly unsurprising, given that her knowledge of R is implicit and her belief

that if R, then Chomsky is mistaken, is explicit. In fact, the speaker's explicit belief must be that *if R obtains or is a rule of language*, then Chomsky is mistaken. And as Chomksy says, knowing that R is quite different from knowing that R obtains, or is a rule of language:

> [If John knows that R] we cannot assume that John knows that R holds, obtains, is a rule of his language. John quite probably does not know this, although some linguist may. In other words, there is no legitimate "semantic ascent" from "John knows that R" to "John knows that R holds". (Chomsky 1986, p.266).

Why not? Well, for the ascent to be legitimate the speaker's knowledge that R would need to be conscious and explicit. In this case, Knowles is surely correct that a "straightforward reason for denying that inferential integration takes place is just that linguistic knowledge is, in a rather deep way, inaccessible to consciousness." (Knowledge 2000, p.339). In a reply to Knowles, Rattan objects that inaccessibility to consciousness cannot explain (full) inferential isolation, since ordinary unconscious attitudes, such as an implicit sexist belief, do exhibit some inferential integration (Rattan 2002, pp. 141-2). For example, we may suppose that an unconscious belief that women in general are constitutionally hysterical, and a conscious suspicion that *this* particular woman's complaints about sexual harassment are overblown and paranoid, are inferentially linked. Thus implicitness or inaccessibility to consciousness cannot immediately account for a propositional attitude's inferential isolation. But I deny Evans' claim that semantic knowledge is fully inferentially isolated. My contention is only that the limitation on inferential integration *about which Stich complains* is a consequence of the knowledge being implicit and unconscious. In the case of the generic belief about women's hysteria, a fairer comparison would be something like this. Suppose I implicitly believe that women are constitutionally hysterical (but men are not), and my analyst has suggested that I have this implicit belief. I consciously recognise that if I do have the implicit belief, then

my analyst is right, but I will not, of course, draw the conclusion that my analyst is right as long as my belief about women's hysteria remains unconscious.

There is, as I have said, a more general limitation on the integration of semantic knowledge in the rest of our mental lives. For semantic knowledge is only recruited in the use and understanding of language. Evans is right about this much; what he does not acknowledge is how rich the information integrated in linguistic activity is, and how versatile, within that activity, our knowledge of the meaning of sub-sentential expressions and of the significance of semantic structure can be. Again, the activity-restriction should be unsurprising when we consider the content of the knowledge together with the fact that it is implicit. As Weiss points out, knowledge of the best chess strategies for mating when one has only two bishops will not be manifest in projects other than those aimed at check-mating a chess opponent, because that is what the knowledge is *for* (Weiss 2004, p. 78). Of course, if the knowledge is explicit it will be potentially serviceable in other pursuits, e.g. one might reel off chess strategies to impress someone. But as long as the knowledge is implicit, and given that its content earmarks it for use in chess games, it is precisely there we should expect to see the knowledge exploited and manifest. The analogy between language use and chess is well-worn, but it is inessential to this particular point about the restrictions on integration implied by an attitude's *content*. For consider the belief that women are prone to hysteria. This belief will be rationally integrated only with other attitudes which share some part of the content, e.g. the desire not to be hysterical (which could combine with the belief to make one eschew femininity), or the belief that hysterical people are volatile, which could lead one to think women are more likely to be volatile. Of course, a long enough chain of inferences could lead the belief to be (rather indirectly) manifested in just about

any form of behaviour. But the same is true of semantic knowledge. When you utter a sentence, my semantic knowledge is part of what informs inferences about what you say, what you intend, what background beliefs and expectations you have. For example, if I know that you have recourse to the term 'off-putting' but employ 'rebarbative' in a context in which half of your audience will not understand you, I may infer either that you are trying to exclude them or make them uncomfortable, or that you are ignorant of your audience; on that basis I may form certain hypotheses about your character, which go on to inform my own behaviour towards you in all kinds of ways.

Conclusion

Linguistic behaviour is paradigmatically intentional: the use of language is an essential means to achieving perhaps the majority of our goals. As such it demands a rationalising explanation, one that cannot be given if we suppose speakers to be ignorant of semantic and syntactic information about their language. Worse, if Hawthorne and Stanley are correct in thinking that intentional action proceeding from no known premise violates a norm of practical rationality, the denial of cognitivism suggests that language use is systematically *irrational* insofar as it violates the knowledge norm.

In fact, our everyday linguistic transactions can be seen to rely upon our common rational sensitivity to semantic information, most conspicuously when some hitch or obstacle forces us to hesitate and reflect upon how to best convey something by linguistic means, or how most reasonably to interpret our interlocutor's remarks. In all cases, a wealth of other information is brought to bear in these deliberations, and our guidance by semantic facts is mediated by sensitivity to other factors, and by our intentions and

desires. It follows not only that we have genuine semantic *knowledge*, for absent such knowledge, we could not show the rational appreciation of semantic information which we do show; this knowledge is rationally integrated with other attitudes, it is not merely *tacit*.

BIBLIOGRAPHY

Achinstein, P. (1975) 'The Object of Explanation', in S. Körner (ed.), *Explanation*. Oxford: Blackwell.

Alvarez, M. (2010) *Kinds of Reasons*. Oxford: Oxford University Press.

Antony, L.(2008) 'Meta-Linguistics: Methodology and Ontology in Devitt's *Ignorance of Language*' *Australasian Journal of Philosophy,* vol 86 (4), pp. 643-656.

— (2003) 'Rabbit Pots and Supernovas: On the Relevance of Psychological Data to Linguistic Theory', in Berber, A (ed.) *Epistemology of Language*. Oxford: Oxford University Press.

— (1997) 'Meaning and Semantic Knowledge' *Proceedings of the Aristotelian Society*, Supplementary Volumes, Vol. 71, pp. 177- 209.

Aquinas, T. (1953[1475]) *Quaestiones Disputatae. Volumen I: De Veritate*, Torino: Marietti.

Aristotle (2006) *Metaphysics Book Θ*, Makin, S. Trans. Oxford: Clarendon Press.

Baker, G.P. and Hacker, P.M.S. (1984), *Language, Sense and Nonsense,* Oxford: Blackwell.

Bengson, J. and Moffett (2011), M. Introduction to *Knowing How: Essays on Knowledge, Mind, and Action* (Bengson & Moffett eds.), New York: Oxford University Press.

Brown, J. (2013) 'Knowing-how: linguistics and cognitive science' *Analysis* 73 (2): 220-227.

Binder, K.S., Duffy, S.A. and Rayner, K. (2001) 'The Effects of Thematic Fit and Discourse Context on Syntactic Ambiguity Resolution.' *Journal of Memory and Language*, Vol 4 (2), pp.297-324.

Binder, K.S. and Morris, R.K. (1995), 'Eye movements and lexical ambiguity resolution: Effects of prior encounter and discourse topic' *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, pp. 1186–1196.

Campbell, J. (1982) 'Knowledge and Understanding', *Philosophical Quarterly* 32, 17-34.

Cath, Y. (2011) 'Knowing How without Knowing That' Knowledge' in Bengson & Moffett (eds.) *Knowing How: Essays on Knowledge, Mind, and Action*, New York: Oxford University Press.

Collins, J. (2008) 'Knowlegde of Language Redux' *Croatian Journal of Philosophy*, vol. 22, pp 3-43.

— (2007) 'Book Review: Ignorance of Language by Michael Devitt' *Mind*, pp. 416-423.

— (2004) 'Faculty Disputes' *Mind and Language*, vol. 19 (5), pp. 503-533

Chomsky, N. (1957) *Syntactic Structures*. New York: Mouton de Gruyter.

— (1965) *Aspects of the Theory of Syntax*, Cambridge, MA: MIT Press

— (1966) *Cartesian Linguistics: A Chapter in the History of Rationalist Thought*. New York: Harper & Row.

— (1972) *Language and Mind*, Enlarged edn. New York: Harcourt Brace Jovanovich.

— (1980) *Rules and Representations*. New York: Columbia University Press.

— (1986) *Knowledge of Language*

— (2000) *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.

Dancy, J. (2000) *Practical Reality*. Oxford: Oxford University Press.

— (2011) 'Acting in Ignorance' *Frontiers of Philosophy in China*, vol. 6, pp. 345-357.

Davidson (2001[1984]) *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press.

Davies, M. (1981) 'Meaning, Structure and Understanding', *Synthese* Vol 48, 1, pp. 135-161.

— (1986) 'Tacit Knowledge and the Structure of Thought and Language' in Travis, C. S. (ed.), *Meaning and Interpretation*. Oxford: Blackwell

— (1987) 'Tacit Knowledge and Semantic Theory: Can a Five Percent Difference Matter?' *Mind*, vol. 96 (384), pp. 441-462.

— (1989) 'Tacit Knowledge and Subdoxastic States' in George, A. (ed.) *Reflections on Chomsky*, Oxford: Blackwell.

— (2015) 'Knowledge – Explicit, implicit and tacit: Philosophical aspects' in J.D. Wright (ed.), *International Encyclopedia of Social and Behavioral Sciences* (Second Edition). Oxford: Elsevier Ltd.

Devitt, M. (2006a) *Ignorance of Language.* Oxford: Clarendon Press.

— (2006b) 'Intuitions in Linguistics' *British Journal for the Philosophy of Science* 57, pp. 481–513

— (2010) 'Linguistic Intuitions Revisited' *British Journal for the Philosophy of Science*, vol. 61 (4), pp. 833 - 865.

— (2011a) 'Linguistic Knowledge' in Bengson & Moffett (eds.)

Knowing How: Essays on Knowledge, Mind, and Action, New York: Oxford University Press.

— (2011b) 'Methodology and the nature of knowing how' *Journal of Philosophy* 108: 205–18.

Devitt, M. and Sterelny, K. (1989) 'What's Wrong with the 'Right View'', *Philosophical Perspectives,* vol. 3, pp. 497-531.

— (1987), *Language and Reality: An Introduction to the Philosophy of Language*, Cambridge: MIT Press.

Dorr, C. (2014) 'Transparency and the Context-Sensitivity of Attitude Reports' in García-Carpintero, M. & Martí, G. *Empty representations: Reference and Non-Existence*, Oxford: Oxford University Press.

Dummett, M. (1993) *The Seas of Language*, Clarendon Press: Oxford.

Evans, G. (1981) 'Semantic Theory and Tacit Knowledge', in Holtzman, S.H. & Leich C.M. (eds.), *Wittgenstein: To Follow a Rule*, Routledge and Kegan Paul: London

Fodor, J.A., (1990) 'Review of Stephen Schiffers's *Remnants of Meaning*', in *A Theory of Content and Other Essays* MIT Press: Cambridge, MA.

— (1989) 'Why Must the Mind be Modular', in George, A. *Reflections on Chomsky,* Blackwell: Oxford.

— (1983) *The Modularity of Mind.* Cambridge, MA: MIT Press.

— (1975) *The Language of Thought.* New York: Crowell.

Foster, J., (1976) 'Meaning and Truth Theory', in Evans, G. and McDowell J., *Truth and Meaning: Essays in Semantics*, Oxford Clarendon Press: Oxford.

Geurts, B. and Rubio-Fernández, P. (2015) 'Pragmatics and Processing' *Ratio* (new series) XXVIII, 4, pp. 446-469.

Fricker, E. (1983), 'Semantic Structure and Speaker Understanding', *Proceedings of the Aristotelian Society.*

Ginet , C. (1975) *Knowledge, Perception, and Memory,* Dordrecht: Reidel.

Grandy, R.E. (1990) 'Understanding and the Principle of Compositionality' *Philosophical Perspectives* 4, pp. 557-572.

George, A. (1989) 'How not to become confused about Linguistics' in George (ed.), *Reflections on Chomsky*, Oxford: Blackwell.

Harman, G. (1975) 'Language, Thought and Communication', in Gunderson, K. (ed.), *Minnesota Studies in the Philosophy of Science*, vol. 7, Minneapolis: University of Minnesota Press.

Hawley, K. (2003) 'Success and Knowing How' *American Philosophical Quarterly*, 40.1 (2003), 19-31.

Hawthorne, J. and Stanley, J. (2008). 'Knowledge and Action'. *The Journal of Philosophy* 105, pp. 571 – 590.

Heck, R. (2006a) 'Reason and Language', in Macdonald, C. & Macdonald, G. (eds.), *McDowell and His Critics*. Oxford: Blackwell.

— (2006b) 'Idiolects' In Thomson, J.J. & Byrne, A. (eds.), *Content and Modality: Themes From the Philosophy of Robert Stalnaker*. Oxford: Oxford University Press

— (2007) 'Meaning and Truth Conditions' In Greimann, D. & Siegwart, G. (eds.), *Truth and Speech Acts: Studies in the Philosophy of Language*. Routledge.

Hornsby, J. (2007a). 'Knowledge, Belief, and Reasons for Acting' in *Explaining the Mental.* Penco, C., Beaney, M., & Vignolo, M. (eds). Cambridge Scholars Publishing.

— (2007b). 'A Disjunctive Conception of Acting for Reasons' in Haddock, A. and MacPherson, F. (eds), *Disjunctivism: Perception, Action, Knowledge*. London: Birkbeck ePrints.

— (2011) 'Ryle's Knowing-How and Knowing How to Act', in Bengson & Moffett (eds.) *Knowing How: Essays on Knowledge, Mind, and Action*, New York: Oxford University Press.

Hornsby, J. & Stanley, J. (2005) 'Semantic Knowledge and Practical Knowledge' *Aristotelian Society Supplementary Volume*, 79:107-145.

Higginbotham, J. (1988) 'Is Semantics Necessary?' *Proceedings of the Aristotelian Society*, 87, pp. 219-241.

— (1995) 'The Place of Natural Language' in Leonardi, P. and Santambrogio, M. (eds.) *On Quine*. Cambridge: Cambridge University Press.

Hyman, J. (2015) *Action, Knowledge and Will,* Oxford: Oxford University Press

— (2011). 'Acting for Reasons: Reply to Dancy'. *Frontiers of Philosophy in China*, 6, pp. 358 – 368

— (2006) 'Knowledge and Evidence'. *Mind*, 115 (460), pp. 891-916.

— (1999) 'How Knowledge Works'. *The Philosophical Quarterly* 49, pp. 433 – 451

Johnson, K. and Lepore, E. (2004) 'Knowledge and Semantic Competence' in Niniluoto, I., Sintonen, M. and Wolenski, J. (eds.) *Handbook of Epistemology*, Dordrecht: Kluwer Academic Publishers.

Johnson, K.(2007) 'Tacit and Accessible Understanding of Language' *Synthese*, vol.156 (2), pp. 253-279.

— (2004) 'Tacit Belief, Semantics and Grammar' *Linguistics and Philosophy*, Vol. 27, No. 1, pp. 57-91.

Kenny, A. (1989) *The Metaphysics of Mind,* Oxford: Oxford University Press.

Kneale, W. C. & Kneale, M. ([1962] 1984)) *The Development of Logic*, Oxford: Clarendon Press.

Kripke, S. (1982) *Wittgenstein on Rules and Private Language,* Oxford: Blackwell.

Larson and Segal (1995) *Knowledge of Meaning: an Introduction to Semantic Theory*. Cambridge, MA: MIT Press.

New Series Vol 97, pp. 41-60

— (1985) 'In Defense of Davidson', *Linguistics and Philosophy* 5, pp. 277- 294.

Levin, M. (1977) 'Explanation and Prediction in Grammar (and Semantics)' *Midwest Studies in Philosophy,* 2 (1), pp. 128-137.

Lewis, D. (1975) 'Languages and Language' In Gunderson, K. (ed.) *Language, Mind and Knowledge*. Minneapolis: University of Minnesota Press.

Longworth, G. 'Conflicting Grammatical Appearances' *Croatian Journal of Philosophy*, vol. 21 (3), pp. 403-426.

Ludlow, P. (2003) 'Externalism, Logical Form and Linguistic Intentions' in Berber, A. (ed.) *Epistemology of Language*. Oxford: Oxford University Press.

— (2011) *The Philosophy of Generative Linguistics,* Oxford: Oxford University Press.

Matthews (2003) 'Does Linguistic Competence Require Knowledge of Language?' in Berber, A. (ed.) *Epistemology of Language*. Oxford: Oxford University Press.

— (2006a) 'Could Competent Speakers Really be Ignorant of Their Language?' *Croatian Journal of Philosophy* Vol. 6 (3), pp. 457-467

— (2006b) 'Knowledge of Language and Linguistic Competence' *Philosophical Issues* vol. 16, pp. 201-220.

McDowell, J.(1988) 'In Defense of Modesty', in Barry Taylor, ed., *Michael Dummett:*

Mele, A.R. (1992) *Springs of Action: Uniderstanding Intentional Agency*. Oxford: Oxford University Press.

— (2003) *Motivation and Agency*. Oxford: Oxford University Press.

Moore, A. (1997) *Points of View,* Oxford: Clarendon Press.

— (2003) 'Ineffability and Nonsense', *Proceedings of the Aristotelian Society, Supplementary Volumes*, Vol. 77 , pp. 169-193.

Noë, A. (2005) 'Against intellectualism' *Analysis* 65: 278–90.

Partee, B. (1975) 'Montague Grammar and Transformational Grammar.' *Linguistic Inquiry*, vol. 6, pp. 203-300.

— (1979) 'Semantics: Mathematics or Psychology?' In Bäuerle, R. and Egli, E (eds.) *Semantics from Different Points of View*, pp. 1-14.

Partee, B. H. (2015) 'The Garden of Eden period for deep structure and semantics' in *50 Years Later: Reflections on Chomsky's* Aspects*,* Gallego, A.J. and Ott, D. (eds.) Cambridge, MA: MITWPL

Pavese, C. (2015) 'Practical Senses' *Philosopher's Imprint*, 15 (29).

Peacocke, C. (1986) 'Explanation in Computational Psychology: Language, Perception and Level 1.5' *Mind and Language* vol. 1, pp. 101-123

— (1989) 'When is a Grammar Psychologically Real?' in George, A. (ed.) *Reflections on Chomsky*, Oxford: Blackwell.

Petit, D. (2002) 'Why Knowledge is Unnecessary for Understanding Language' *Mind*, vol. 111 (443), pp. 519-550.

Poston, T. (2009) 'Know How to be Gettiered?' *Philosophy and Phenomenological Research*, vol. 179 (3), pp. 743-747.

Prior, A.N, Geach, P.T. and Kenny, A.J. (1971) *Objects of Thought*, Oxford: Oxford University Press.

Quine, W.V.O. (1972) 'Methodological Reflections on Current Linguistic Theory' In Harman, G. and Davidson, D. (eds.) *Semantics of Natural Language*. New York: Humanities Press.

Raz, J. (1990) *Practical Reason and Norms*. Princeton, NJ: Princeton University Press.

— (1999) *Engaging Reason.* Oxford: Oxford University Press.

Rey, G. (2006) 'Conventions, Intuitions and Linguistic Inexistents: A Reply to Devitt' *Croatian Journal of Philosophy,* vol.1 (18) pp. 549-569

Rumfitt, I. (2011), 'Inference, Deduction, Logic', in Bengson & Moffett (eds.) *Knowing How: Essays on Knowledge, Mind, and Action*, New York: Oxford University Press.

— (2003) 'Savoir Faire', *Journal of Philosophy,* 100, pp. 158-166.

Ryle, G. (1946) 'Knowing How and Knowing That: The Presidential Address' *Proceedings of The Aristotelian Society,* 46, pp. 1-16.

Ryle, G. (2009[1949]) *The Concept of Mind,* Oxford: Routledge.

Sandis, C. (2012) 'Can Action-Explanations Ever Be Non-Factive?' in Bakhurst, D., Little, M.O., & Hooker, B. (eds.), *Thinking About Reasons: Themes From the Philosophy of Jonathan Dancy*. Oxford: Oxford University Press.

Setiya, K. (2015) 'Causality in Action' *Analysis* 3 (3), pp. 501-512.

Schroeder, S. (2006) *Wittgenstein: The Way out of The Fly Bottle.* Cambridge: Polity Press.

Snowdon, P. (2003) 'The Presidential Address: Knowing How and Knowing That: a Distinction Reconsidered', *Proceedings of The Aristotelian Society,* 104, pp. 1-29.

Soames, S. (1989) 'Semantics and Semantic Competence' *Philosophical Perspectives,* Vol. 3, pp. 575–596.

— (1984) 'Linguistics and Psychology' *Linguistics and Philosophy* 7, pp. 155-179.

Stanley, J. (2011) *KnowHow*. Oxford: OUP

— (2009) 'Knowing (how).' *Nous* 45.2: 207-238.

Stanley, J. and Williamson, T.(2001) 'Knowing How', *The Journal of Philosophy,* 98 (8), pp. 411-444.

Stich, S.P. (1971) 'What Every Speaker Knows' *The Philosophical Review*, Vol. 80, no. 4, pp. 476-496.

Stich, S. P. (1978) 'Beliefs and Subdoxastic States' *Philosophy of Science*, vol. 45, pp. 499-518.

Stroud, B. (1991) 'The Background of Thought' in Lepore & Van Gulick (eds.), *John Searle and his Critics,* Oxford: Blackwell.

Toribio, J. "How do we know how?" *Philosophical Explorations* 11 (1), 39-52.

Unger, P. *Ignorance: A Case for Skepticism*, Oxford: Oxford University Press.

Waights Hickman (unpublished a) 'Knowing in the "executive way": knowing how, rules, methods, principles and criteria.'

— (unpublished b) 'Novelty and Knowledge (How)'.

— (unpublished c) 'Devitt on Linguistic Competence'

— (unpublished d) '(Implicit) Knowledge, Reasons and Semantic Understanding.

Weatherson, B. (forthcoming) 'Intellectual skill and the Rylean regress' *Philosophical Quarterly*.

Weiss, B. (2010) *How to Understand Language: A Philosophical Inquiry*, Durham: Acumen.

— (2004), 'Knowledge of Meaning' *Proceedings of the Aristotelian Society,* Vol. 104, pp. 75-94.

White, A.(1982) *The Nature of Knowledge*. New Jersey: Rowman and Littlefield.

Wiggins, D. (2012) 'Practical Knowledge: Knowing How To and Knowing That', *Mind,* Vol. 121, no. 481

— (2009) 'Knowing How To and Knowing That' in *Wittgenstein and Analytic Philosophy: Essays for P.M.S. Hacker* (Glock & Hyman eds.). Oxford: Oxford University Press.

— (1997) 'Languages as Social Objects' *Philosophy*, Vol. 72 (282), pp. 499-524.

Williamson, T. (2000) *Knowledge and its Limits,* Oxford: Oxford University Press

Wilson, D. and Sperber, D. (2012) *Meaning and Relevance.* Cambridge: Cambridge University Press.

Wittgenstein, L. (2001[1953]). *Philosophical Investigations.* Oxford: Blackwell

— (1981[1967]) *Zettel.* Oxford: Blackwell

— (2007[1958]) *The Blue and Brown Books*. Oxford: Blackwell

Wright, C. (1986) 'Theories of Meaning and Speakers' Knowledge', in *Realism, Meaning and Truth*, Oxford: Blackwell.

— (2001) *The Rails to Infinity*. Harvard University Press

— (2007) 'Rule following without reasons: Wittgenstein's quietism and the constitutive question' *Ratio*, 20(4):481 – 502.