

(Implicit) Knowledge, reasons, and semantic understanding

Natalia Waights Hickman 

Worcester College, University of Oxford,
Oxford, UK

Correspondence

Natalia Waights Hickman, Worcester
College, Oxford OX1 2HB, UK.
Email: natalia.waightsickman@
philosophy.ox.ac.uk

This paper exploits recent work on the normative and constitutive roles of knowledge in practical rationality, to put pressure on the idea that speakers could communicate without exploiting linguistic *knowledge*. I defend cognitivism about meaning, the view that speakers have *rationaly accessible* (i.e., implicit rather than tacit) knowledge of semantic facts and principles, and that this knowledge is constitutive of their linguistic competence.

KEYWORDS

knowledge norm, knowledge view of reasons, rational integration criterion, rationality of language, semantic cognitivism, semantic understanding

1 | INTRODUCTION

This paper aims to bring recent work in epistemology to bear on a longstanding debate concerning the role of (largely implicit) knowledge of language in an ordinary speaker's linguistic competence. In particular, I am interested in the role that knowledge of meaning plays in the composition and interpretation of utterances: In the question whether such knowledge mediates communication.

The debate concerns, at the most abstract level, how speakers are related to the languages at their command, and so potentially bears on the conception of linguistic theory and the facts to which it is responsible, as well as on theories of linguistic understanding. Despite therefore being of foundational importance in the metaphysics and epistemology of language, the debate has reached something of a deadlock, and a change of tack is needed to make certain issues more tractable. The connections recently highlighted in epistemology between knowledge, rationality, reasons, and intentional action may help to provide this.

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 she *φs* (e.g., acts, reasons, judges) or by which she is guided.¹ This is not to say that ignorance in every case defeats rationality or every dimension of rationality, but that an agent cannot be rationally responsive to *the facts* if she is ignorant of them.² A rationalising explanation of an agent's acts or beliefs thus imputes to her knowledge of the *explanans*—the fact that is the explanatory reason. If this is correct, there is 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. It is through this lens I want to revisit the question whether an ordinary speaker's use of language should be seen as merely conforming with facts and principles of her language, or as expressing *knowledge* of these.

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) have argued that, *ceteris paribus*, intentional actions that are premised on a reason that is not known *violate a norm* of practical rationality. If this is correct, it matters enormously whether language use, as such, is considered to be intentional action appropriately explained in terms of the agent's reasons. If it is, but the linguistic facts which *would* rationally explain an agent's linguistic actions and interpretations are held not to be epistemically accessible to her, then it will follow that her use of language is not merely *non-rational* but systematically *irrational* along one dimension.

In what follows I appeal to these ideas in putting pressure on the notion that ordinary speech and comprehension could proceed without speakers exploiting *knowledge* of language. Although parallel arguments may be available in connection with syntax, my focus will be on semantic knowledge, or knowledge of linguistic meaning, and I will defend the view known as *semantic cognitivism*³ (I will sometimes refer to this simply as “cognitivism”).

According to semantic cognitivism, speakers possess knowledge of the semantic facts and principles of their language, and this knowledge is (partly) constitutive of their linguistic competence. The constitutivity thesis means that the semantic knowledge possessed by speakers is *non-redundant* to their understanding and use of language; it is something a speaker of a language needs in order to count as such. Thus while opponents of cognitivism take linguistic communication to be knowledge independent, recruiting no knowledge of semantic structure, word or even sentence meaning, cognitivists view knowledge as essential to the communicative process: It is only via the speaker's semantic knowledge that she is able to form sentences apt for communicating what she intends, and it is through such knowledge, again, that her audience comes to understand her.

Moreover, the cognitivist thinks that linguistic knowledge is fully fledged knowledge of the *speaker*, it is not a body of information contained in sub-personal systems of the speaker's mind. Some of the knowledge may be available to consciousness and to some degree articulable by its possessor, most plausibly knowledge of word meaning.⁴ Some, however—knowledge of the

¹Unger (1975), Hornsby (2007), Hawthorne and Stanley (2008), and Hyman (1999, 2015).

²If 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, for example, the conditional *p* → *q* is true. Similarly, if an agent *φs* 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.

³Johnson and Lepore (2004).

⁴I am in fact sceptical that speakers have very much explicit knowledge of lexical meaning: One can be semantically competent with an expression without having any idea where to begin on an explanation of its meaning. It is also not

semantic significance of different modes of combination—will be unavailable in these ways and count as *implicit*. But on a cognitivist view it is in either case genuine propositional knowledge which is accessible to reason. As I discuss further in the next section, cognitivism therefore contrasts both with *tacit* knowledge theses, which postulate sub-personal representational states, and with anti-cognitivist views on which the facts and structures of natural language are not represented or accessed in *any* sense by ordinary speakers. Finally, the semantic knowledge cognitivists postulate is assumed to be of a particular, indefinitely productive sort, namely, knowledge of the meaning of simple expressions (words) and knowledge of principles determining 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 has always been its explanatory power with respect to the productive or creative aspect of linguistic competence, that is, a speaker's ability to compose and understand indefinitely many novel sentences of her language.⁵ 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 (Dummett, 1993). Cognitivism allows one to recognise those semantic facts and 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 based (Heck, 2006; Lepore, 1997; Weiss, 2004, 2010). 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. I thus align myself with Higginbotham (1988, 1989) in viewing the explanation of language use that semantic knowledge offers as rational in nature, and available because semantic facts and principles are in some sense available, at least potentially, to an agent in thought. As I shall explain shortly, in current terminology, this requires taking speakers to know semantic facts and rules at least *implicitly* and not merely *tacitly*,⁶ and in this connection it will be necessary to engage with Gareth Evans's contention that semantic “knowledge” is inferentially isolated, and thus not a genuine propositional attitude.

Section 2 situates the cognitivist position I shall champion within a space of familiar views of semantic competence and highlights some of the issues at stake in its defence. Section 3 lays out the problem for opponents of cognitivism if there is a knowledge norm on practical rationality of the kind postulated by Hawthorne and Stanley (2008). Section 4 argues that the reason-guided nature of language use *as such*—not just what we say but, so to speak, *how* we say it—is evident throughout our linguistic practice, and that semantic facts are among the rationalising *explanantia*. Following the “knowledge view of reasons”, speakers must be credited with knowledge of these rationalising facts; moreover, the cases discussed make clear that the knowledge is *non-redundant* to successful linguistic communication. I conclude in Section 5 with a discussion of Evan's rational integration criterion.

likely that a speaker will describe meaning in the terms used by the semantic theorist. For example, she would not, in advance of training, represent the meaning of “green” as, say, a function from possible worlds to functions from objects to truth values; the concept of a function may be entirely alien to her.

⁵Davidson and his followers sought to reorient semantics as a theory of speaker's understanding with precisely this challenge in mind (see Davidson, 1984, essays 1–4, and Lepore, 1985, on the ‘epistemological turn’ Davidson gave semantics). The primary impetus in this direction came of course from Chomsky—on the creativity of language see, for example, Chomsky (1965, 6, 15, 1966, 57–59, 1972, pp. 100).

⁶See Davies (2015) on notions of implicit and tacit knowledge.

2 | SEMANTIC COMPETENCE AND SEMANTIC KNOWLEDGE THESES

Semantic cognitivism is a view about (part of) what *constitutes* linguistic competence. It is easiest to see what the thesis comes to by focusing on linguistic communication. For the cognitivist, communication is mediated by or reliant upon knowledge of language. For example, when Sally understands Jane's utterance of "the dinner is ready", she relies *inter alia* on knowledge of the meaning of the words that Jane has used, and knowledge of what follows (semantically) from combining the words in that way. It is via such knowledge that she comes to interpret *what Jane says* as well as anything she may be implicating, presupposing, and so forth. The opponent of cognitivism does not, as I have said, allow any such role for semantic knowledge in communication. For what the cognitivist and her opponent fundamentally disagree about is what linguistic competence *is*, and accordingly about what its exercise constitutively involves. Thus, it might be (and often is) granted by non-cognitivists that a speaker typically does possess some semantic knowledge—for example, Sally might, as a *bi-product* of her linguistic competence, have knowledge of the standing meaning of "the dinner is ready". But the non-cognitivist maintains that any such knowledge is not constitutive of linguistic competence and is redundant to its basic exercise: Sally does not understand Jane's utterance of "the dinner is ready" *because* she knows what the relevant linguistic expressions mean; it is not through her knowledge of meaning that comprehension is achieved. As Hornsby comments in a critique of cognitivism, "the fact that we *have* such [semantic] knowledge may be an obstacle to our appreciating that we ordinarily *make no use of it*" (Hornsby & Stanley, 2005, p. 28; emphasis added). My defence of cognitivism in this paper aims at showing, on the contrary, that as speakers we do not only have semantic knowledge, we also ordinarily do and must make use of it in communication: Semantic knowledge is *constitutive* of linguistic competence.

Beyond the minimal claim that knowledge of meaning is necessary to successful language use and understanding, the cognitivist makes no claim about the nature of any "psychological processing" of language. This point has proven extraordinarily easy to miss. Cognitivists are often accused of pedalling confused ideas about symmetry between natural language structures and real-time processing in the mind/brain.⁷ But as has been reiterated many times,⁸ the cognitivist position concerns only (some of) the knowledge a speaker draws upon in the use and understanding of language, and says nothing about *how* such knowledge is drawn upon. The misunderstanding might not be so prevalent were there more emphasis on the fact that we refer to an agent's knowledge in *rationalising* explanations of her actions and attitudes. That is the kind of explanation the cognitivist can tell a story about, she need have nothing to say about the causal relations that the postulated knowledge states enter into at a neurophysiological level.⁹

Much of the knowledge of meaning that cognitivists believe ordinary speakers to possess cannot, in general, be explicit or available to conscious reflection. The full formulation of semantic principles that are rich enough to permit the meaning of any sentence to be derived—principles

⁷See, for example, Levin (1977), Soames (1984, 1989), Devitt and Sterelny (1989), and Devitt (2003, 2006, 2011).

⁸Chomsky already addressed the "continuing misunderstanding" at the beginning of *Aspects of the theory of syntax* (Chomsky, 1965, p. 9, cf. Chomsky, 1965, pp. 140–141, 1986, p. 67; George, 1989). Devitt (2006, pp. 62–65) nevertheless contrives to understand Chomsky's use of "linguistic competence" in a way Chomsky explicitly disavows, that is, as referring to language processing mechanisms, and accuses Chomsky of misunderstanding his own commitments. Collins (2007) issues a sound corrective.

⁹See George (1989).

showing how what speakers do achieve can rationally be achieved—is generally thought to be given by a compositional semantic theory, whose statement requires intellectual resources that ordinary speakers do not possess. In fact, it is doubtful whether there is *any* conception of linguistic facts or structures that would not put them beyond the conceptual capacities of ordinary speakers to articulate. Even a philosopher who maintains that the only “real” linguistic rules are those recorded in traditional grammar books¹⁰ must acknowledge a complexity of linguistic structure well beyond what any ordinary speaker has reflective access to, and the need for concepts and theory in its explicit representation that will be quite alien and unnecessary to ordinary speakers.¹¹ The debate between cognitivists and their opponents is not about the correct representation of meaning. It is also not a debate about how speakers relate to particular *theoretical representations* of meaning, but rather about how speakers relate to *what is then represented*—namely, the semantic facts and structures of languages they speak (George, 1989). Cognitivism is at bottom an epistemological thesis concerning the relationship between these natural language facts and structures—however they are best theorised—and the competence of speakers. Is linguistic competence consistent with ignorance of language? The cognitivist says not. Since linguistic structures are highly complex on *any* account, cognitivism entails the notoriously 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. This is also true, of course, of any knowledge possessed by non-human animals and pre-verbal infants: *Explicit* formulations of such knowledge can in those cases not be recognised as apt or otherwise by the subject, and cannot be cited in explanations or justifications of their actions and attitudes except by third parties.

Many philosophers have taken these violations of standard criteria for (explicit!) knowledge ascription to demonstrate that semantic cognitivism is absurd.¹² Others, in partial sympathy, and perhaps wishing to avoid the charge of absurdity, have taken it to show that semantic “knowledge” is not a genuine or full propositional attitude. Thus, Gareth Evans urged that to assimilate semantic knowledge to ordinary propositional attitudes is a “mysterious and confused position” (Evans, 1981, p. 134), and 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 (pp. 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, 1981, 1987, 1989, 1997).¹³

¹⁰For example, Hanfling (1980).

¹¹For example, the entry under “Subjunctive” in Fowler’s *Modern English usage* spans three pages, and it is registered there that the standard reference work on historical English syntax devotes no less than 156 pages to the subjunctive mood. The author nevertheless finds it natural to say in the preface that the rules articulated in his work are ones that the reader—the ordinary speaker—already knows. As far as the attribution of implicit knowledge to ordinary speakers goes, I see nothing more or less extraordinary about this position than the position of those who think the grammatical and semantic rules of language must be recursive or generative, to account for the productivity of language, and that they are aptly represented using the methods of formal logic and mathematics.

¹²For example, Baker and Hacker (1984), Hanfling (1980), Foster (1976), Devitt (2006, 2011), Schiffer (1987), and Hornsby and Stanley (2005).

¹³A similar view of semantic knowledge as “sub-doxastic” is advanced by Peacocke (1986). See Wright (1986a), Davies (1987), Knowles (2000), and Weiss (2004) for critical discussion of this version of cognitivism.

These latter conceptions of semantic knowledge purport to secure its explanatory credentials with respect to linguistic creativity as well as its amenability to empirical investigation, but give 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. 133–134)

As Wright (1986a) 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” (p. 34). 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, 1986a).

It is tempting to reply by urging that Evans’s 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, pp. 338–339; Weiss, 2004, pp. 78–79). 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’s integration criterion. The position I will defend is that Evan’s criterion *is* met by semantic knowledge, and that Wright’s rhetorical questions concerning the role of desires can in fact receive straightforward answers, although there are of course *some* limitations on the potential for integration of any implicit semantic knowledge (more will be said about this in Section 5). Semantic knowledge is, on my view, at least implicit; accessible to reason. It then differs from explicit knowledge in that its possessor cannot put the knowledge into words. But it *shares* with explicit knowledge the feature of being available as a *reason* on which the agent can draw in conjunction with her other beliefs and desires in the pursuit of diverse ends. In this sense, it sits between explicit knowledge which is both verbally *articulable* and inferentially *integrated*, and tacit “knowledge” which is both linguistically *inexpressible* and inferentially *isolated*.

The idea that speakers possess and exploit any kind of linguistic knowledge in the exercise of their linguistic competence has been much controverted; two popular alternatives are a “functional role” conception of semantic understanding, and the view that linguistic

competence is a species of know-how. According to 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.¹⁴ According to the second popular alternative to cognitivism, semantic competence is knowledge-*how*, assumed not to be propositional knowledge of any kind, but a complex ability or disposition of the agent.¹⁵ Though differing in other important respects,¹⁶ these positions agree in the central anti-cognitivist contention that even if speakers *possess* some knowledge of language—for example, knowledge of what sentences mean—this knowledge plays no role in securing successful linguistic communication (Hornsby 2005, p. 128; cf., p. 8; Schiffer 1987, p. 262). Competence is merely having a capacity to go from what is heard—a token utterance—to what is said. 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. (Schiffer, 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 & Stanley, 2005, pp. 111–112). On both views, whatever linguistic competence consists in, it is *not* constituted by or reliant upon semantic knowledge. It is this thesis—the thesis that semantic knowledge is redundant to linguistic competence—that I shall target below.

3 | LINGUISTIC RATIONALITY AND THE KNOWLEDGE NORM

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. 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” (Hyman, 2015, p. ix). The questions that naturally arise concerning knowledge are now different: As Hyman puts it, “we need to think prospectively not retrospectively, about

¹⁴Schiffer (1987, 1992) and Fodor (1984, 1990).

¹⁵This has long been the default alternative put forward by critics of linguistic cognitivism, for example, Stich (1971), Soames (1984, 1989), and Devitt (2011). It is given a more positive elaboration and defence in Hornsby and Stanley (2005). The position is of course tendentious insofar as the knowing-how/that distinction is now a matter of some controversy.

¹⁶For example, Hornsby would not sanction a mechanistic construal of semantic competence or the appeal to a private language of thought; semantic competence is not, on her view, a sub-personal mechanism but an intelligent, complex ability of the agent, creatively employed in performing intentional actions.

how knowledge is applied, employed, expressed, in the infinitely varied circumstances of human life" (ibid).

This shift to considering the practical efficacy and import of knowledge 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 reasons and practical rationality. According to Hawthorne and Stanley, knowledge provides a *norm* of practical rationality (Hawthorne & Stanley, 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 if and only if she knows that p . The proposal motivates our common folk appraisals of action in epistemic terms (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 know that she had, I would be criticised for *acting* on the premise that Elena missed her flight if I did not know it, for example, if I cancelled her place at dinner on those grounds—even if Elena did in fact miss her flight.

Naturally, there are many circumstances under which an agent can be exculpated for acting in ignorance, for example, cases in which we are forced to choose a course of action, but none is favoured by known premises. But as Stanley and Hawthorne note, drawing an analogy with conditions on the acceptability of assertion, the possibility of exceptions or exemptions is consistent with the operation of a genuinely normative principle.

The conceptual structure, one familiar from the normative realm, explains the suitable appraisal of action in terms of a combination of norms and excuses for failure to comply with them. (Hawthorne & Stanley, 2008, p. 273)

Hawthorne and Stanley 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.¹⁷ I am sceptical that the principle is demanding *enough*, that is, 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.¹⁸ But I believe they are correct that knowledge is normatively relevant to the appraisal of intentional actions; in

¹⁷The central source of resistance to views on which rationality requires knowledge is that epistemological scepticism then entails irrationality—indeed, this conclusion is endorsed by Unger (1975), as a consequence of his arguments for scepticism. I believe knowledge and rationality do come together, and the aim of this paper is to show that in so far as language use is rational, it relies on semantic knowledge. But it should be borne in mind that the argument of Section 3 could be modified: A weaker *belief* norm on practical rationality would still motivate the view that linguistic competence involves propositional attitudes, though they may not constitute *knowledge*. The argument of Section 4 is that speakers must be credited with *knowledge* of semantic facts, however, not merely belief.

¹⁸For 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*, that is, whether or not the agent *prefers* ϕ -ing conditional on whether or not p . By my (non-Humean) lights the former is a criterion of practical rationality, and an agent can fail to satisfy it while 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.

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 (1993) famously maintained that language use is the paradigm of rational activity, and that this imposes a constraint on theories of meaning:

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*. (Dummett, 1993, p. 104)

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 is provided by knowledge of the sentence's meaning. Lepore (1997) likewise argues that non-cognitivist accounts of understanding should be rejected because they leave speakers' interpretations of utterances unrationalized, that is, they fail to identify *reasons* that ground them.¹⁹ But if Hawthorne and Stanley are correct in maintaining that acting intentionally on unknown premises contravenes a norm of practical rationality, then the anti-cognitivist position is more problematic. Given the knowledge norm, an agent is under normal circumstances subject to rational criticism if she intentionally ϕ s without known reasons to ϕ . A speaker will then, 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 of any reason to adopt S for the achievement of her communicative goals. And according to the anti-cognitivist, this is the position that a speaker is in as a matter of course, not in extraordinary and exculpatory circumstances. By Stanley and Hawthorne's lights, even a true belief about the meaning of S would not be enough,²⁰ normally speaking, to make the speaker's use of S fully practically rational: "When someone acts on a belief that does not amount to knowledge, she violates the norm, and is thus subject to criticism" (Hawthorne & Stanley, 2008, p. 574).

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 action²¹ 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. The anti-cognitivist might want to say that the

¹⁹See also Campbell (1982), McDowell (1988), and Weiss (2004).

²⁰Pettit (2002) has argued that knowledge is not necessary to semantic understanding because *belief* formed on unsafe or unreliable testimony will suffice. The choice between attributions of linguistic knowledge versus belief is not of central concern here, but the arguments of this study suggest nothing weaker than semantic *knowledge* will, in general, be adequate to rationally explain and justify language use. If this position is correct, Pettit's examples are better taken as evidence that semantic knowledge, like some knowledge-wh, is not subject to standard epistemic warrant or safety conditions—it does not "gettierize".

²¹For example, Raz (1999, Ch. 2) and Mele (1992, Ch. 6.). The view goes back to Davidson (1963) and Anscombe (1957).

knowledge norm does not apply to ordinary speech and interpretation, because these are things we “just do” *without* grounds. The anti-cognitivist's positive position is, after all, that the exercise of linguistic competence involves no implicit mediation by linguistic premises, so why should she take it as a criticism that on her view the speaker is taken to avail herself of linguistic expressions in communication without knowledge of premises or reasons on which to do so?

There are two ideas that could vindicate such a response: The idea that language use is not intentional, and the idea that language use, though intentional, belongs to that anomalous category of intentional acts legitimately performed on a whim or without special ground. Neither seems to me tenable. There may be certain linguistic acts that are not plainly intentional or purposive, such as mumbling absent-mindedly as 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 or are in some sense arbitrary. It cannot in general, then, be taken to be the sort of activity to which norms of practical rationality, including Stanley and Hawthorne's knowledge-norm, simply do not apply. 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 the arbitrary adoption of (what turns out to be!) means—that is, the adoption of certain means, without being apprised of reasons to think those means appropriate?

But this brings into view an important commitment concerning the *aspect* in which language use can be taken to be intentional, subject to reasons-explanations and assessable in terms of norms of rationality. My contention above was that one's choice of communicative *means*—that is, one's choice of words—as well as whatever one is doing in communicating (testifying that *p*, promising to *φ*, etc.) is intentional in character. The opponent of cognitivism might well grant that uses of language are intentional if described as *sayings*, and that of course a speaker will have reasons which explain and perhaps justify their saying various things, but their knowledge of language need not figure in any such explanations or justifications. This is precisely Hornsby's (2005) position. A speaker's intentionally saying that *p* is in general not even potentially rationalised by semantic facts—it may be rationalised, *inter alia*, by knowledge that *p*, and aspects of the conversational context. So the mere fact that language use is intentional under some description is not enough to build a case against the anti-cognitivist: It must be shown to be intentional *as such*, intentional in an aspect that semantic facts are suited to 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.²²

²²Hornsby claims in the same paper that a conception of semantic competence as know-how can also vindicate the Dummettian idea that language use is distinctively rational, without viewing this as constituted by knowledge of semantic facts and principles. I cannot see the grounds for this claim, however. Hornsby draws attention to a number of important features of linguistic practice which suggest that it is rational, including its complexity, and the fact that linguistic competence is exercised in concert with other abilities and knowledge. She also points out that speech is something over which we exercise rational control, and that in speaking we usually have complex aims—not just to say that *p*, but implicate that *q*, and so on (Hornsby & Stanley, 2005, pp. 125–126). But all of this is, so to speak, the rationality *data* which an account of linguistic competence needs to explain or accommodate, and I do not see how this is to be achieved by an account of competence as a complex ability—an ability independent of knowledge of, and thus rational sensitivity to, semantic facts and principles.

4 | LINGUISTIC RATIONALITY AND THE KNOWLEDGE VIEW OF REASONS

4.1 | Knowledge and reasons-explanations

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 defended by Unger (1975), but developed and argued for at greater length by Hornsby (2007) 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*.²³ 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 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.²⁴

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 Hyman (2015) 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” (p. 152). If we explain behaviour by citing a fact to which an agent is sensitive—a fact that guides or is a reason for their behaviour—the agent must be credited with knowledge of that fact. Related to this is what may be called the transparency of knowledge. We can rationally explain an agent’s actions *either* by referring to their *knowledge that p* or by referring directly to the fact that *p*; for “knowledge ... is transparent: we can look straight through it to the fact” (p. 157).

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:

²³The 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.

²⁴See Hyman (2011, 2015, Ch. 7).

One cannot be guided by a fact one does not know, any more than one can follow a guide one cannot see. If the traveler 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 (2007):

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. (Hornsby, 2007, p. 94)

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.

4.2 | Implicit reasons, rationalisation, and motivation

The reasons that guide or motivate us cannot be simply identified with those we would or could articulate if asked, and not only because we often lack first-person authority, that is, because we may be deceived about why (the reasons for which) we do, feel or think the things we do. In fact, what I shall call implicit reasons, reasons that guide action without being consciously dwelt upon or offered as reasons by the agent, are ubiquitous, and not because implicit knowledge is ubiquitous. Explicit knowledge can serve as an implicit reason; implicit reasons would be common even if there were no such thing as implicit knowledge. To see this one needs only to consider that what we *think about and offer* as reasons for our actions are not reasons-for-action *simpliciter*, but reasons that (purport to) justify or explain some aspect of the action in question.

Hyman suggests that what unifies manifestations of knowledge that *p* as such is that they are doings which are guided by the reason that *p* (e.g., Hyman, 1999, pp. 440–441). But whenever we act, most of the knowledge informing or “guiding” us is not in focus; it is not brought to mind before or as we act, and some of it is so seldom brought to our attention that we might struggle to identify its presence and role in the background to our deeds. As I walk a particular route through the city, relying on knowledge how to navigate the traffic, read the street names, and cross the roads, as well as on the perceptual knowledge I acquire along the way that lets me keep track of my progress, I may think to myself only (or give as a reason if questioned) that I am doing what I am doing because this is the route to the post office. But every step is guided by—would be impossible without—much more knowledge than is conveyed by a statement that this is a way to reach the post office. That such knowledge is not at the forefront of my

mind or the tip of my tongue does not show that it is at this time, so to speak, beyond the reach of my reason; that it is not taken into consideration or account as I proceed on my way to the post office, that the facts on which I rely are not yet *reasons* that *guide me*.

Some might think this “background” knowledge does not figure as a reason but should be treated as a mere enabler. A reason is a consideration that recommends, from the agent’s point of view, ϕ -ing as in some respect preferable to the alternatives; it is a consideration in light of which the action appears a favourable or good one. An enabler or enabling condition, on the other hand, neither speaks for nor against ϕ -ing, it simply renders ϕ -ing possible (Dancy, 2000, pp. 126–127). It might be said then that the knowledge on which we draw when we act, but do not reflect on or cite as a reason for action, should be categorised as an enabler and not as an “implicit reason”. Thus when Mary makes a move in chess, her reason for making that move and not some other—the reason she has in mind, and will offer if asked—will be strategic; whereas the knowledge she has of defining rules of the game merely *enables* her to make that move, but it does not provide her with reasons; the defining rules of chess do not speak in favour of her making *that* move but not some other.

This is unconvincing. As was said above, what becomes salient as an agent’s reason for action depends on what aspect of the action one wants to explain or to justify, or what the agent herself thinks of herself as needing to explain or to justify; likewise, whether and how each piece of knowledge drawn upon “speaks in favour” of the action, or renders it reasonable or right from the agent’s point of view, depends on what stands in need of explanation or justification, or which alternative is to be ruled out. If Mary is asked why, say, she has not moved the 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 and not as a mere enabler. Depending on who asks the question, it might seem absurd or pointless, but it is not unintelligible or illegitimate, and there is an answer—a reason—to be given in reply. In sum: What qualifies as an agent’s reason and what, instead, as an enabler is *question-relative*.²⁵

Reasons, then, can be and often are implicit. But how do we determine that an agent draws on a fact as one of her reasons if she lacks the vocabulary, perhaps the concepts, to make it explicit? No implicit knowledge attribution is uncontroversial, but a pertinent though hypothetical example for comparison with the linguistic case is described by Wright (1986a):

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, 1986b, p. 23)

I share Wright’s conviction that in this case, the children’s chess-playing merits an imputation to them of implicit knowledge, of at least the defining rules of chess.²⁶ But why? First, it is clear

²⁵For further discussion of the question-dependence of explanations in general and reasons-explanations in particular, see Sandis (2012) and Achinstein (1975).

that the rules *rationalise* the children's behaviour, they provide grounds in light of which their behaviour can be seen as reasonable or right. Second, if we do not ascribe this knowledge it is, as Wright says, difficult to see *how* the children could achieve what they do achieve, in particular how they are able to respond "to board configurations which they have never encountered before". The hypothesis of ignorance leaves the children's creative capacities unexplained. Third, while the children cannot cite the rules in words, they do not conform to them, as it were, mechanically, whether they would like to or not: They will deliberately modify play that flouts the rules, systematically vetoing violations and endorsing permitted moves, demonstrating by example what conformity entails. In this sense we can see the rules not merely as potential rationalisers but as motivators, rules to which the children exhibit sensitivity of the right sort: *Because of which* they act. They count as knowing the rules, because they are rationally responsive to them in the way, following the knowledge view of reasons, only knowledge affords; but the knowledge is implicit, for while the children can betray what they know in action, they cannot articulate it—that is, formulate the rules—in language.

This, I believe, is exactly analogous to linguistic cases: Rational sensitivity to semantic information which only knowledge affords is revealed *inter alia* in our recognition and correction of misinterpretations and misuses of terms, the resolution of ambiguities, and deliberation over how best to express a thought given a variety of intentions and desires. I want to urge, focusing on familiar kinds of cases of this sort, that speakers' routine uses and interpretations of language are explained by semantic facts which constitute *reasons* the speaker is sensitive to, and that here as elsewhere we must credit the speaker with *knowledge* of the relevant facts. By focusing precisely on how semantic facts feature as reasons which interact with an agent's other beliefs and desires in securing successful linguistic communication, I aim simultaneously (a) to illustrate the non-redundancy of semantic knowledge in linguistic competence, and (b) to counter the charge that any semantic "knowledge" is inferentially isolated and not a genuine propositional attitude.

4.3 | The intentionality of language use and exploitation of semantic reasons

As I explained above, to avoid the irrationality charge of Section 3, it would be best for the anti-cognitivist to maintain that language use and comprehension are not *as such* intentional—the things we do with words, such as lie, promise, or inform, may be intentional, but my opponent should claim that our doing them *with particular linguistic expressions* is not. Hornsby (2005) makes just this suggestion: 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"—that is, utter a certain linguistic string—"and says that *p* by doing it" (p. 118). If this is correct, then it would seem to follow that facts about meaning are not even suited to be reasons for language use.

At least, it would follow that facts about meaning are not *in general* suited to rationalise language use. Exceptions would arise where meaning is itself the topic of an utterance. Consider the following example from Weiss (2004), designed to show that it is not always explanatorily

²⁶If later the children learned to articulate the rules in language, it would be natural to describe this as akin to learning the name of an already familiar object—they are learning to "name" the rules they already knew beforehand, rather than learning the rules for the first time.

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” will involve one’s belief about that meaning. (Weiss, 2004, p.80).

In connection with Hornsby’s charge, the salient point 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, since in this case what the speaker is attempting to convey is something about linguistic meaning itself. But precisely because of their nature, cases like these cannot do enough to support cognitivism. Turning language back on *itself* in meta-linguistic discourse is arguably a special and derivative form of language use, and any knowledge recruited to it need not be recruited when language is not the topic of an utterance—in ordinary, predicative, uses of “red”, for example. Yes, if a speaker says something *about a linguistic expression*, she might deploy knowledge concerning it, just as if she says something about Keats, photons or the cat in the doorway she will deploy knowledge concerning *them*. But it by no means follows that in the latter cases too, the speaker recruits semantic knowledge, such as knowledge of the meaning of “cat”, which is what the cognitivist maintains. Cases of meta-linguistic discourse provide limited support for cognitivism because they confirm only that *what a speaker says* is subject to reason-explanations, and that where what she says concerns language itself, linguistic reasons will naturally be among the *explanantia*. But what needs to be shown is that when a speaker says something or other about anything, her *using a certain linguistic string to do so* is rationally explained, *inter alia*, by semantic facts.

The cases I focus on below therefore do not involve meta-linguistic discourse. They are intended to reveal the role that semantic knowledge routinely plays in rationally informing the speaker’s choice of linguistic means to say something, and the hearer’s understanding of that intended proposition. They highlight the intentional nature of language use *as such*, and the constitutive role that sensitivity to semantic reasons plays in the successful exercise of linguistic competence.

First, consider the resolution of 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. Suppose Anne utters (a) and Jane is unsure which of two possible propositions Anne means to convey. In order to ascertain this, she queries the meaning of the *one word* “bug”—does Anne mean an insect or a piece of spyware? The question is made possible by Jane’s recognition of the ambiguity of that expression and how that ambiguity infects the sentence. Raising the question of that word’s meaning 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. It thus presupposes more than isolated knowledge of word meaning; more knowledge than the subject will be able to put into words.

Now consider a case of intentional code-switching. One familiar context in which intentional code-switching occurs is in the attempt to exclude a third party from understanding what is said. Thus, consider two bilingual Spanish/English speakers, David and Natalia. Their friends, Jane and Anne, are monolingual English speakers, and are conversing at the same table. David wants to meet Natalia to buy Jane a surprise gift, and he wants to make arrangements in secret. He could conduct the entire conversation in Spanish, but this would seem rude and alert the others to a conspiracy. So instead he decides to code-switch with a single expression, and says to Natalia:

(b) Let us meet at the market tomorrow and look for a regalo.

I have already described many of the facts that have entered David’s practical deliberations: His *reasons* for intentionally code-switching, for example, he wants to communicate such-and-such exclusively to Natalia, without alerting their neighbours to the attempt to exclude them from understanding. But one crucial reason for his linguistic act, is the fact that the Spanish word “regalo” means *gift*. He exploits this fact in uttering a string that Natalia will be able to interpret; Natalia exploits it in the interpretation; both exploit the linguistic ignorance of their friends. Note that David’s semantic knowledge is integrated with his other beliefs and desires in the service of a particular goal. Knowledge that “regalo” means *gift* is useful in this context only in combination with knowledge that Jane and Anne’s understanding of (b) with the switched expression will be compromised by ignorance of the fact that “regalo” has this meaning; knowledge that Natalia’s understanding will not be so compromised; and all of this is relevant only given David’s desire to exclude Jane and Anne from understanding, but without causing offence or suspicion. In this way David’s semantic knowledge interacts straightforwardly with other beliefs and desires.

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, and sensitivity to more than semantic properties of expressions will often 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 misunderstood 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.

Turning back to language interpretation rather than production, structural ambiguities can also make salient a speaker's need to weigh semantic information in order to successfully exercise her linguistic competence. 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 concerning the processes by which 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.²⁷ But 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 not things that "just happen"; they are forms of intentional behaviour for which semantic facts provide implicit grounds. The knowledge of lexical meaning that a speaker then recruits may or may not be partly explicit in the sense that a speaker may—but need not²⁸—be able to offer some provisional explanation of lexical meaning. But the semantic significance of each different parsing of the sentence, which speakers show sensitivity to in such cases and deliberately exploit, is not something they will be able to articulate.

Cases of this general sort pervade linguistic practice at every level. When communication is smooth and unhesitating, alternative assignments of lexical meaning or semantic structure are not salient, and there is no conscious process of deliberation involving semantic knowledge. But the hiccups and dilemmas that make its essential role more perspicuous are mundane. We all recognise slips of the tongue on the part of ourselves and others, and that such slips compromise success in communicating what was intended. Speakers with relatively undeveloped capacities for meta-linguistic reflection will still deliberately adjust their speech for children and non-fluent speakers, simplifying syntax and using more "simple" or common vocabulary to convey their message. We all recognise and correct more stable misunderstandings of the meaning of expressions, and we recognise that *because* a speaker ignores the facts about meaning she uses expressions which do not serve her communicative aims. It is often a very straightforward

²⁷Binder, Duffy and Reiner (2001), and Binder and Morris (1995).

²⁸A speaker can be semantic competent with many expressions without knowing how to explain their meaning, and any explanations she can give will likely be partial at best.

affair to see communication is being botched by linguistic ignorance (a student thinks “refute” means *deny*) and how it would be repaired by knowledge: The lesson is that such knowledge is the general condition of success.

Of course, when communication is smooth and unhesitating the positive role that semantic knowledge plays in rationally guiding it will be phenomenologically invisible. We are generally 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. And here Hornsby’s metaphors of “hearing meaning” and “voicing thoughts” are more likely to resonate. But it does not follow that Hornsby is right in thinking smooth linguistic communication bypasses semantic knowledge. The phenomenology here is not peculiar to an agent’s reliance on *semantic* reasons; our reasons are often not conscious or salient to our attention. If Mary is engaged in playing chess, and her focus is on the relative merit of different strategies, the defining rules of the game are simply taken for granted, they will not be a focus of reflective attention. Her focus is, so to speak, not on the framework of the game but on the moves available within it. If, however, she attempts an illicit move or someone proposes a move that is illicit, then the relevance of the framing rules, and proper grasp of these, becomes immediately apparent. We surely should not think that Mary makes rational use of knowledge of the defining rules of chess *only* where they are called explicitly into question or brought to conscious attention; that her ordinary reflection on and deployment of chess strategies in the normal case could proceed unimpeded if she made no use of this knowledge or simply lacked it (cf., Section 4.2). Likewise, we cannot conclude that semantic knowledge is redundant in cases where the speaker and hearer are focused only on what is being said, and not on *how*, and comes suddenly to be relevant only when some hiccup or dilemma prompts attention to be directed to a speaker’s language as such.

5 | INFERENCE INTEGRATION AND ITS LIMITS

If the foregoing is correct, then at least a great deal of semantic knowledge must be accessible to reason.²⁹ I cannot, from any number of such cases, prove that *all* semantic facts are rationally available and deployed by speakers in order to successfully communicate. But I believe it is clear that some is both rationally available and necessary for communicative success, from which it follows (a) that speakers have semantic knowledge which is non-redundant to, or *constitutive* of, their semantic competence; and (b) that such knowledge is not *tacit* or inferentially isolated.

It is not the case, 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 so 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, for example, aspects of the discourse context, 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

²⁹Cognitivism does not require *all* semantic knowledge to be integrated in reasoning; one can coherently maintain that some semantic knowledge is implicit while some is merely tacit (see Johnson, 2004, 2007).

(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 her ultimate choice. There is simply no limit to the potential information about her surroundings and her 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 she 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 is true that semantic knowledge is only serviceable in the performance of and response to *linguistic* acts, that is, is 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 Chomsky 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" (Knowles, 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. 141–142). For example, we may suppose that an *unconscious* generic belief that women are oversensitive, and a *conscious* suspicion that *this* particular woman's complaints are overblown, are inferentially linked. Thus, implicitness or inaccessibility to consciousness cannot immediately account for a propositional attitude's inferential isolation. But I deny Evans's 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 implicit generic belief about women, a fairer comparison would be something along these lines. Suppose I implicitly believe that women are oversensitive (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 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. But acknowledging *some* limitation to the inferential integration of an attitude should not impugn its status as a genuine propositional, personal-level attitude. For one thing, the inferential integration of many of our explicit beliefs—think of mathematical beliefs—is plausibly limited by our cognitive abilities. Inconsistencies and “blind-spots” can be quite robust, and some inferences require rare deductive skills. Smith may be robustly unable to see and respond to the inferential relations between two propositions she believes no matter how much instruction or priming she receives, and the fact that the relevant beliefs could in principle be fully integrated in some *other* subject's reasoning—her teacher's, for example—does not seem relevant to the question whether Smith herself holds those beliefs. It is a (putative) propositional attitude's potential for integration with others in the thought of an individual subject that matters to the question whether that attitude can be ascribed to *her*.

According to Evans, tacit semantic “knowledge” has no flexibility or versatility *at all* in the agent's rational psychology; the semantic facts are never, on Evans's view, reasons that can be weighed alongside others, and responded to differently depending on the agent's desires in a way that manifests her rational agency. Following the knowledge view of reasons, it is then natural for Evans to refuse such “knowledge” the status of a genuine propositional attitude, for genuine knowledge that *p* is expressed, on this view, in responding to *p* as a reason. Yet as can be seen in the cases discussed above, semantic facts do figure as reasons which a linguistic agent can intentionally respond to differentially, depending on other beliefs and desires; and following the knowledge view of reasons, the agent's responsiveness to these facts *as* reasons is one that only genuine knowledge affords. What Evans does not acknowledge, then, is how rich the information integrated in linguistic activity itself 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, for example, 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.

It is no surprise that much of our semantic knowledge is restricted to use in the comprehension and composition of sentences, for its content earmarks it for use in precisely these tasks, while its unavailability to consciousness precludes its being recruited in an arbitrary range of further projects. Nevertheless, within linguistic practice, semantic knowledge is deeply integrated with reasoning, mediated by desires and intentions, and beliefs about other linguistic properties and “pragmatic” factors, such as the interests, background knowledge and expectations of our audience. In this way it is *essential* to ordinary speech and linguistic comprehension: It is *constitutive of linguistic competence*.

ACKNOWLEDGEMENTS

I am grateful to Anna Drożdżowicz, Herman Cappelen, Martin Davies, Timothy Williamson and Adrian Moore for comments on an early draft of this paper.

ORCID

Natalia Waights Hickman  <https://orcid.org/0000-0003-0762-0635>

REFERENCES

- Achinstein, P. (1975) *The object of explanation*. in S. Körner (ed.), *Explanation*. Oxford: Blackwell.
- Anscombe, G. E. M. (1957) *Intention*. Oxford: Blackwell.
- Alvarez, M. (2010). *Kinds of reasons*. Oxford: Oxford University Press.
- Baker, G. P. & Hacker, P. M. S. (1984). *Language, sense and nonsense*. Oxford: Blackwell.
- Binder, K. S., Duffy, S. A. & Rayner, K. (2001). The effects of thematic fit and discourse context on syntactic ambiguity resolution. *Journal of Memory and Language*, 4(2), 297–324.
- Binder, K. S. & 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, 1186–1196.
- Campbell, J. (1982). Knowledge and understanding. *The Philosophical Quarterly*, 32, 17–34.
- Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.
- Chomsky, N. (1966). *Cartesian linguistics: A chapter in the history of rationalist thought*. New York, NY: Harper & Row.
- Chomsky, N. (1972). *Language and mind* (Enlarged ed.). New York, NY: Harcourt Brace Jovanovich.
- Chomsky, N. (1986). *Knowledge of language*. New York, NY: Praeger.
- Collins, J. (2007). Book review: Ignorance of language by Michael Devitt. *Mind*, 116, 416–423.
- Dancy, J. (2000). *Practical reality*. Oxford: Oxford University Press.
- Davidson, D. (1963). Actions, reasons and causes. *The Journal of Philosophy*, 60(23), 685–700.
- Davidson, D. (1984). *Inquiries into truth and interpretation*. Oxford: Oxford University Press.
- Davies, M. (1981). Meaning, structure and understanding. *Synthese*, 48(1), 135–161.
- Davies, M. (1987). Tacit knowledge and semantic theory: Can a five percent difference matter? *Mind*, 96(384), 441–462.
- Davies, M. (1989). Connectionism, modularity and tacit knowledge. *The British Journal for the Philosophy of Science*, 40(4), 541–555.
- Davies, M. (1997). Meaning and semantic knowledge. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 71, 177–209.
- Davies, M. (2015). Knowledge—Explicit, implicit and tacit: Philosophical aspects. In J. D. Wright (Ed.), *International encyclopedia of social and behavioral sciences* (2nd ed.). Oxford: Elsevier Ltd.
- Devitt, M. (2006). *Ignorance of language*. Oxford: Clarendon Press.
- Devitt, M. (2011). Intuitions in linguistics. *British Journal for the Philosophy of Science*, 57, 481–513.
- Devitt, M. & Sterelny, K. (1989). What's wrong with the “right view”. *Philosophical Perspectives*, 3, 497–531.
- Dummett, M. (1993). *The seas of language*. Oxford: Clarendon Press.
- Evans, G. (1981). Semantic theory and tacit knowledge. In S. H. Holtzman & C. M. Leich (Eds.), *Wittgenstein: To follow a rule*. London: Routledge and Kegan Paul.
- Fodor, J. A. (1990). Review of Stephen Schiffer's *Remnants of meaning*. In *A theory of content and other essays*. Cambridge, MA: MIT Press.
- Foster, J. (1976). Meaning and truth theory. In G. Evans & J. McDowell (Eds.), *Truth and meaning: Essays in semantics*. Oxford: Clarendon Press.
- George, A. (1989). How not to become confused about linguistics. In A. George (Ed.), *Reflections on Chomsky*. Oxford: Blackwell.
- Hanfling, O. (1980). Does language need rules? *The Philosophical Quarterly*, 30(120), 194–205.
- Hawthorne, J. & Stanley, J. (2008). Knowledge and action. *The Journal of Philosophy*, 105, 571–590.
- Heck, R. (2006). Reason and language. In C. Macdonald & G. Macdonald (Eds.), *McDowell and his critics*. Oxford: Blackwell.
- Higginbotham, J. (1988). Is semantics necessary? *Proceedings of the Aristotelian Society*, 88, 219–241.
- Higginbotham, J. (1989). Knowledge of reference. In A. George (Ed.), *Reflections on Chomsky*. Oxford: Blackwell.
- Hornsby, J. (2007). Knowledge, belief, and reasons for acting. In C. Penco, M. Beaney & M. Vignolo (Eds.), *Explaining the mental*. Newcastle: Cambridge Scholars Publishing.

- Hornsby, J. (2005). Semantic knowledge and practical knowledge. *Aristotelian Society Supplementary Volume*, 79, 107–145.
- Hyman, J. (1999). How knowledge works. *The Philosophical Quarterly*, 49, 433–451.
- Hyman, J. (2011). Acting for reasons: Reply to Dancy. *Frontiers of Philosophy in China*, 6, 358–368.
- Hyman, J. (2015). *Action, knowledge and will*. Oxford: Oxford University Press.
- Johnson, K. (2004). Tacit belief, semantics and grammar. *Linguistics and Philosophy*, 27(1), 57–91.
- Johnson, K. (2007). Tacit and accessible understanding of language. *Synthese*, 156(2), 253–279.
- Johnson, K. & Lepore, E. (2004). Knowledge and semantic competence. In I. Niiniluoto, M. Sintonen & J. Wolenski (Eds.), *Handbook of epistemology*. Dordrecht: Kluwer Academic Publishers.
- Knowles, J. (2000). Knowledge of grammar as a propositional attitude. *Philosophical Psychology*, 13(3), 325–353.
- Lepore, E. (1985). In defense of Davidson. *Linguistics and Philosophy*, 5, 277–294.
- Lepore, E. (1997). Conditions on understanding a language. *Proceedings of the Aristotelian Society, New Series*, 97, 41–60.
- Levin, M. (1977). Explanation and prediction in grammar (and semantics). *Midwest Studies in Philosophy*, 2(1), 128–137.
- McDowell, J. (1988). In defense of modesty. In B. Taylor (Ed.), *Michael Dummett: Contributions to philosophy*. Dordrecht: Martinus Nijhoff.
- Mele, A. R. (1992). *Springs of action: Understanding intentional agency*. Oxford: Oxford University Press.
- Peacocke, C. (1986). Explanation in computational psychology: Language, perception and level 1.5. *Mind & Language*, 1, 101–123.
- Pettit, D. (2002). Why knowledge is unnecessary for understanding language. *Mind*, 111(443), 519–550.
- Rattan, G. (2002). Tacit knowledge of grammar: Reply to Knowles. *Philosophical Psychology*, 15(2), 135–152.
- Raz, J. (1999). *Engaging reason*. Oxford: Oxford University Press.
- Sandis, C. (2012). Can action-explanations ever be non-factive? In D. Bakhurst, M. O. Little & B. Hooker (Eds.), *Thinking about reasons: Themes from the philosophy of Jonathan Dancy*. Oxford: Oxford University Press.
- Schiffer, S. (1987) *Remnants of Meaning*. Cambridge, MA: MIT Press.
- Soames, S. (1984). Linguistics and psychology. *Linguistics and Philosophy*, 7, 155–179.
- Soames, S. (1989). Semantics and semantic competence. *Philosophical Perspectives*, 3, 575–596.
- Stich, S. P. (1971). What every speaker knows. *The Philosophical Review*, 80(4), 476–496.
- Stich, S. P. (1978). Beliefs and subdoxastic states. *Philosophy of Science*, 45, 499–518.
- Unger, P. (1975). *Ignorance: A case for scepticism*. Oxford: Oxford University Press.
- Weiss, B. (2004). Knowledge of meaning. *Proceedings of the Aristotelian Society*, 104, 75–94.
- Weiss, B. (2010). *How to understand language: A philosophical inquiry*. Durham: Acumen.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Wright, C. (1986a). How can the theory of meaning be a philosophical project? *Mind & Language*, 1, 31–44.
- Wright, C. (1986b). Theories of meaning and speakers' knowledge. In C. Wright (Ed.), *Realism, meaning and truth* (pp. 204–238). Oxford: Blackwell.

How to cite this article: Waights Hickman N. (Implicit) Knowledge, reasons, and semantic understanding. *Mind & Language*. 2021;36:707–728. <https://doi.org/10.1111/mila.12286>