

Idiolects

Richard G. Heck, Jr.

Brown University

1 Language

Students of language should presumably study language, but there is room for disagreement about how the ordinary notion of a language should be refined for the not so ordinary purposes of scientific or philosophical inquiry. One central question is to what extent, or in what ways, the appropriate object of study is socially or communally constituted. And one answer, of course, is that it isn't. On this view, the primary objects of study are *idiolects*. An idiolect, in this sense, belongs to a single individual, in the sense that one's idiolect reflects one's own linguistic capabilities and, in that sense, is fully determined by facts about oneself. An alternative is to hold that the primary objects of study should be *common languages*, so called because they are the common property of a community of individuals. It need, of course, be no part of this view that common languages are individuated as we individuate languages in ordinary discourse: Perhaps so, but perhaps not. The important claim is that such languages are in some way socially constituted, in the sense that their properties—most importantly, for present purposes, their semantic properties—are fixed not by the linguistic capabilities of a single individual but only by relations among the linguistic capabilities of an entire community of individuals.

Noam Chomsky has argued, to my mind convincingly, that what he calls 'E-languages' are of no concern to theoretical linguistics (as practiced in the generative tradition).¹ It is rather with 'I-languages'

¹ Arguments for this conclusion can be found in some of Chomsky's earliest writings, for example, in *Aspects* (1965). There is a more recent exposition in *Knowledge of Language* (Chomsky, 1986).

that linguistics is concerned. Now, despite the suggestive prefixes,² the distinction between E-language and I-language is not the distinction between 'external' and 'internal' languages—between common languages and idiolects. Rather, the 'E' and 'I' stand for 'extensional' and 'intensional', in the sense in which those terms are sometimes used in the philosophy of mathematics: E-language is to I-language as product is to process.³ The notions of E-language and I-language are thus like the two senses of the technical term 'theory': In one of these, a theory is a set of sentences, the theorems of the theory; in the other, a theory is a set of principles, of axioms and rules of inference. 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.⁴ And what Chomsky argues—or just observes, really—is that E-languages in this sense play no significant role in linguistic theory: Syntacticians are not trying to specify the set of grammatical sentences of a given language but rather to specify the *syntactic principles* that determine whether a sentence is grammatical.

² And despite Chomsky's own occasional sloppiness about this point: He frequently writes as if the commitment to the primacy of I-languages is *ipso facto* a commitment to a pretty radical sort of interalism (Chomsky, 1986, pp. 44ff.).

³ The terminology used in *Aspects* (Chomsky, 1965) seems more appropriate, ultimately: There I-languages are called *grammars* and E-languages are called *languages*. If one so speaks, however, one then has to say that linguistics is not concerned with languages, which sounds paradoxical. It is in part for that reason, it seems, that Chomsky changed terminology.

⁴ In fact, or so Chomsky has argued, there is no particular reason to think that grammars do, in any straightforward sense, generate E-languages, since there appears to be no hard-and-fast distinction between grammatical and ungrammatical sentences—none of significant interest to theoretical linguistics, anyway. But I shall waive this point here.

Nor should semantics, if practiced within the generative tradition, attempt to specify a mapping from sound to meaning. Rather, semantics should attempt to describe the principles—the composition rules and semantic axioms—that generate such a mapping.

Chomsky is, nonetheless, thoroughly dismissive of the notion of a common language. Part of his reason is that the notion has, he thinks, never been adequately characterized. He is clearly right on this score. Even if we do not presume that English, German, and the like should be the primary objects of study, but say only that the important question should be whether the objects of study should *in some way* be the property of a linguistic community, no one, so far as I know, has ever managed to say with any precision how the boundaries of linguistic communities should be drawn. But maybe that problem isn't as serious as it seems, or maybe it can be solved. The more serious problem is that it is hard to know what common languages are if they are not E-languages. English, for example—or the dialect thereof spoken in Cambridge, or the language spoken at First Church in Cambridge, or whatever might turn out to be the relevant notion—is an E-language. For what might one mean by *the* set of principles that generates the grammatical sentences of such a 'language'? As Quine long ago pointed out (Quine, 1970), if there is one such set of principles, there are many. Even if we assume that each speaker must use some set of principles to decide whether a given expression is a grammatical sentence and, if so, what it means, different speakers may use different such sets. If E-languages are what matter to us, then such sets of principles are mere means to ends, the ends of determining whether a sentence is grammatical and what it means, and sets of principles—that is, I-languages—are the property of individual speakers, not of linguistic communities. That is why so many philosophers who take E-languages to be the proper objects of philosophical investigation—Michael Dummett, for example—are inclined to dismiss as 'merely empirical' the questions studied in generative linguistics.

Chomsky takes the attitude of such philosophers to be based upon prejudice, and he may to some extent be right.⁵ But it is not just based upon prejudice. To understand Dummett's position, for example, one must understand that he is concerned with the *use* of language—something that is, Chomsky readily admits, in a sense alien to linguistic

⁵ This prejudice Chomsky sees as an heir of behaviorism. He characterizes it as "methodological dualism", which he opposes to "methodological naturalism" (Chomsky, 2000c,b).

theory. Dummett is, as Frege was, particularly impressed by the fact that speakers are able to communicate with one another, and he takes the central problem of the philosophy of language to be to explain how linguistic communication is possible. For that, he thinks, it will be enough if speakers agree about what utterances of sentences do or would mean: It is not required that they arrive at their beliefs about what sentences mean in the same way. Even if Martian psychology is radically different from human psychology, that need not prevent them from communicating with us. Why should it matter what principles Martians use to arrive their judgements about what English utterances mean? If they agree with us, say, that "Snow is white" means that snow is white, isn't that enough? Dummett thinks it is. That is, he thinks that whether speakers share an I-language cannot matter so far as the *use* of language is concerned: What matters is that speakers share an E-language.⁶

Now, as said, E-languages are not common languages: One can draw the distinction between I-language and E-language even for the languages of individuals. But, Dummett claims, the fact that a group of speakers *do* share an E-language can only be explained if what they share is a *common* language, in the sense explained above. He writes:

[O]n Frege's theory, the basic notion is . . . that of an idiolect, and a language can only be explained as the common overlap of many idiolects. This Putnam is quite right to say is wrong. . . . An English speaker both holds himself responsible to, and exploits the existence of, means of determining the application of terms which are either generally agreed among the speakers of English, or else are generally acknowledged by them as correct. . . . [T]his has to be so if words are to be used for communication between individuals. (Dummett, 1978, p. 425)

Dummett's reasons for this claim are similar to Putnam's, which derive from the now famous (but then recent) elm-beech example on which Dummett is here commenting. After rehearsing that example, Dummett concludes that ". . . there is no describing any individual's employment of his words without account being taken of his willingness to subordinate his use to that generally agreed as correct" (Dummett, 1978, p. 425).

⁶ This reasoning, I believe, is extremely common. Versions of it can be found in Dummett (1993, pp. 176–78), Donald Davidson (1984), and David Lewis (1985, pp. 177–8).

So, as said, Dummett does have reasons to think that common languages should be the focus of the philosophical study of language. Let me summarize his train of thought. First, the central preoccupation of the philosophical study of language is the phenomenon of linguistic communication. Second, successful communication requires only that individuals share E-languages, not that they share I-languages. And third, the E-languages individuals share cannot be characterized simply as overlapping idiolects, that is, in terms that ignore the existence of other speakers. Rather, "... one cannot so much as explain what an idiolect is without invoking the notion of a language considered as a social phenomenon" (Dummett, 1978, p. 425). In particular, idiolects should be understood as comprising partial, and partially incorrect, collections of beliefs about the common language.

Note that the crucial question here is one of explanatory priority. Dummett's claim is that common languages are the proper objects of philosophical study because common languages are, in an important sense, more fundamental than idiolects, at least from the point of view of philosophers interested in the use of language for communication. Dummett need not deny—and I do not think he would deny—that the empirical study of language, as exemplified by generative linguistics, might most profitably focus its attention on idiolects. Nor need he deny that, as a matter of empirical fact, actual human speakers' beliefs about what their words mean are the result of cognitive though sub-conscious processes of just the sort Chomsky says linguists are attempting to characterize. Dummett's claim, rather, is that the sorts of facts that are centrally relevant to the general theory of communication—most centrally, facts about what lexical primitives mean—can only be understood properly if common languages are taken as fundamental, because these sorts of facts are socially constituted. That they are is supposed to be shown by the fact that there are generally acknowledged linguistic norms—norms of the common language—to which ordinary speakers hold themselves accountable and without which communication as we know it would be impossible.

I am not endorsing Dummett's argument, though I am prepared to accept the first claim. *My* central preoccupation, *qua* philosopher of language, is indeed the use of language in communication. Even if it were not, I think it would remain an important question whether the general theory of communication—the theory of language-use—can operate with the notion of an idiolect or instead needs the notion of a

common language. The second claim—that successful communication requires only E-languages to be shared—is, I think, extremely tempting, though ultimately mistaken, but I shall not argue that point here.⁷ It is the third claim that will be my focus. What I want to argue is that the existence of the sorts of normative phenomena at which Dummett gestures—in particular, speakers' willingness to hold their usage responsible to that of others—does not imply that idiolects cannot be explained without invoking common languages.

Before I even begin to address this issue, however, let me emphasize that there are several senses in which human language might, in some interesting sense, be 'social' that are not at issue here. One, mentioned to me by Justin Broackes, concerns cultural identity, an important aspect of which may be how one speaks: what words one uses, and with which meanings, how one pronounces them, and the like may all be parts of one's cultural identity. That is true not just among teenagers but among various sorts of ethnic groups (say, Irish Americans) and even among groups whose identity is defined more by geography than by heredity (such as Liverpoolians). Cultural norms may enforce conformity along semantic and phonological dimensions, as well as others. I do not deny that. Nor would I—or, I think, Chomsky—wish to deny that there may be some broadly scientific study of the history and evolution of languages in the everyday sense in which English and German are languages, and perhaps such a study should operate not with idiolects but with languages characterized as socially determined objects that persist through change. As David Wiggins points out (1997),⁸ there is much to be said, and much that has been said, about the development of Italian and the role Dante played in that development. But it does not follow that, within the general theory of communication, common languages must be treated as more basic than idiolects, and that is what is at issue here.

Perhaps most importantly, I do not deny that, given that there are such things as dictionaries, speakers may treat dictionaries as in some

⁷ The argument would have two parts. First, I would argue that this second claim requires that there be no 'residue', so to speak, at the conscious level, of the sub-conscious processes by means of which speakers arrive at their beliefs about what utterances mean. Second, I would argue that there is such a residue (Heck, 2004).

⁸ Since I am going to be disagreeing with almost everything Wiggins says, let me say now that, in my opinion, Wiggins's paper is absolutely terrific. I recommend it to everyone interested in these issues as containing the clearest exposition I know of the position I am opposing.

sense authoritative. There may even be good reasons of various sorts that one ought to hold one's usage responsible to what some dictionary says—perhaps different reasons, and different dictionaries, for different speakers. The issue, however, is not whether one should use the dictionary. What is at issue here does not depend upon the actual existence of dictionaries. Dummett's argument is meant to apply not just to the languages of the industrialized West but also to dying languages that have no written form and only a few hundred speakers, and it is important to keep this fact firmly in view. What Dummett is claiming is that the very possibility of communication among speakers of *any* language demands not only that there must be facts of the sort a correct dictionary for the language might record, and not only that speakers of the language must regard these facts as normative standards governing their own usage, but that an individual speaker's own understanding of the language cannot be characterized independently of these norms.⁹ That is what I deny.

The remainder of the paper is structured as follows. Section 2 introduces a conception of communication—or, better, of communicative exchanges—that guides the discussion. In the discussion that follows, I argue that what we need to explain communicative success is, in the first instance, the notion of what an expression means *to a speaker* rather than the notion of what an expression means *in a language*. There is still room for Dummett to argue that common languages must be invoked in the explanation of successful communication, but I shall argue that, in the sort of explanation Dummett envisages, the reference to common languages is explanatorily idle. Section 3 offers an explanation of the normative phenomena that so impress Dummett, first considering an earlier attempt by Alexander George. Section 4 returns to the question how communicative success should be explained, offering an account that makes no reference to common languages and identifying what I take to be the crucial issue here.

⁹ Compare the quotation from Dummett (1978, p. 425), above. Dummett's argument is thus—to mangle an old joke of Kripke's—a transcendental deduction of the existence of Platonic dictionaries.

2 Communication and Common Languages

Why does Dummett think common languages are implicated in the explanation of successful communication?

Consider the following simple example of a communicative exchange. Say I'm standing outside my office and a student enters the alcove, heading toward Charles Parsons's door. Knowing that Charles is presently teaching, and not wanting the student to waste her time waiting, I say to her, "Prof. Parsons is teaching". She responds, "Oh, thanks," and goes on her way.

Let me call attention to several features of this example:

- (1) The student, whom we may call Janet, acquires a belief as a result of our communicative exchange, namely, the belief that Prof. Parsons is teaching. It is because she acquires this belief, and infers that he is therefore not in his office, that she walks away without knocking on his door. Moreover, under the right circumstances, Janet's newly acquired belief will constitute knowledge.¹⁰
- (2) The belief Janet acquires is one I myself possess, and it is because I wish to convey the information that Prof. Parsons is teaching to Janet that I say that Prof. Parsons is teaching. Moreover, my saying that Prof. Parsons is teaching is an intentional action on my part, one that is under my rational control.
- (3) Janet recognizes the intentional character of my speech. In particular, she recognizes that it is with the intention of conveying to her the information that Prof. Parsons is teaching that I said that Prof. Parsons was teaching. It is for this reason that she thanked me.

That speech is a form of rational action is something I regard as extremely important. It will come to the fore only in section 6, but it will constantly be in the background and so is worth emphasizing.¹¹

What has been said to this point leaves out an important feature of the exchange, namely, that it was a *linguistic* exchange. We can imagine

¹⁰ What the right circumstances are and why her belief would constitute knowledge are, of course, much debated issues. It is my firm belief that nothing here turns upon how they are resolved, but of course I could be wrong.

¹¹ I have discussed the rationality of language-use in some depth elsewhere (Heck, 2005).

variants of the example that would not have had this feature. As Janet approached Prof. Parsons's door, I might have looked at her and shook my head. She might then have acquired the belief that Prof. Parsons was not in his office, thanked me, and gone her way. My shaking my head would yet be an intentional act performed because I wanted to convey to Janet that Prof. Parsons was not in his office. A fourth feature of the example, then, is that I use language in my attempt to convey information to Janet:

(4) I say that Prof. Parsons is teaching by uttering the sentence "Prof. Parsons is teaching".

Like my saying that Prof. Parsons is teaching, my uttering the sentence "Prof. Parsons is teaching" is an act that is under my rational control.

This sort of case, then, involves three apparently related events: (i) my uttering the sentence "Prof. Parsons is teaching"; (ii) my saying that Prof. Parsons is teaching; and (iii) Janet's coming to have the belief that Prof. Parsons is teaching. How precisely are these events related? Why does my uttering *this* sentence lead to Janet's forming *this* belief instead of some other belief, or none at all? It will be worth my saying a few words about why I think this question is the right one to ask.

Before I do so, however, let me emphasize that, while I take the example I am discussing to be reasonably representative of ordinary linguistic communication, I am not claiming, either on Dummett's behalf or my own, that everything we would regard as linguistic communication is, in one way or another, a variation on this theme. What is distinctive of this particular sort of case is that, as noted at (2), a speaker expresses a belief in language and, as a result of her doing so, her audience comes to hold what we would ordinarily regard as the very same belief.¹² Not all cases of successful linguistic communication are like that: There are cases involving implicatures and the like, and there may be other sorts of cases, too. Nonetheless, I think that cases like the one I am treating as exemplary are in some important sense fundamental. I cannot defend that claim in detail at this time, but I will return to the matter from time to time below.

¹² Whether we really should regard it as the same belief—that is, as a belief with the very same content—is another matter. In fact, on my view, the belief does not really have to have the same content for communication to have succeeded. It is a difficult question exactly how the beliefs have to be related, an issue I have addressed in relation to proper names and demonstratives elsewhere (Heck, 1995, 2002). That issue is a subsidiary one here, but it is what led me to the present discussion of idiolects.

H. P. Grice, at one point in his career, anyway, held that successful communication in this case would consist not in Janet's coming to believe that Prof. Parsons is teaching, but rather in her coming to believe that I, the speaker, believe that Prof. Parsons is teaching (Grice, 1989b). If one were to hold that view, then the question I want to raise could be raised by asking why Janet comes to believe that I believe that Prof. Parsons is teaching, rather than that I believe that pigs dance in Peru, or nothing at all. Now, Grice is of course right that communicative success in this case does not require that Janet should in fact form the belief that Prof. Parsons is teaching: It is no failure of the communicative process if Janet refuses to trust me. (Perhaps it is so important that she see Prof. Parsons that she isn't going to trust anyone.) But I do not think the right way to accommodate this observation is to take communicative success to require merely that she come to believe that I believe that Prof. Parsons is teaching.¹³ It seems clear that, struck by a bout of paranoia, Janet might refuse even to believe that I am speaking sincerely and yet that there should be no failure of the communicative process. Janet might understand perfectly what I have said, namely, that Prof. Parsons is teaching. It might, indeed, be only because she understands what I have said that she will neither believe me nor even believe that I am speaking sincerely. Had my remark instead concerned, say, a slippery patch on the floor, perhaps she'd have been happy to take evasive action.

The obvious thing to say, then, would be that communicative success, in the sense in which I am interested in it, requires only that Janet should know what I have said—in this case, that she should know that I have said that Prof. Parsons is teaching. But I would prefer not to raise the question in that familiar form, in large part because this direct approach seems to me unlikely to deliver the illumination we seek. Not only are my intuitions about when someone knows what someone else has said weak and unreliable, one would expect them easily to be corrupted by my theoretical commitments regarding the nature of communication, the notion of what is said, and the like. My intuitions about whether someone could acquire knowledge as a result of a particular communicative exchange are not only stronger and more reliable but are, at least to some degree, independent of the sorts of theoretical issues just mentioned and so are less vulnerable to corruption. So it

¹³ For detailed criticism of this sort of approach, see McDowell (1998c) and Rumfitt (1995). Michael Rescorla extended these arguments in his unpublished PhD dissertation.

seems to me that we need another approach, and my suggestion, in effect, is that we should study the cause by studying its effects. Why, after all, does it so much as matter whether one knows the meaning of an utterance someone else has made? Well, in a certain familiar sort of communicative exchange, it matters because what one takes an utterance to mean will determine the content of the belief one acquires if one takes the speaker at her literal word. If one misidentifies what the speaker has said—if one thinks that her utterance of “Prof. Parsons is teaching” meant that pigs dance in Peru, say—then one risks acquiring a belief that has no particular likelihood of being true, even if one’s informant has spoken knowledgeably. So it seems a reasonable and potentially fruitful strategy to study knowledge of meaning by considering the role it plays in the communicative exchanges it makes possible.¹⁴

The question I should like to discuss is therefore this one: What is the relation between my uttering the sentence “Prof. Parsons is teaching” and Janet’s coming to believe that Prof. Parsons is teaching? In particular, why does my uttering this sentence lead her to acquire this belief? The question is not why she *acquires* the belief, rather than declining my offer of information, but why, if she accepts my offer of information, she will acquire *this* belief instead of some other; and the question does not merely seek a causal explanation but a sense of why it is rational for Janet to form this belief (a fact that is related to its potentially constituting knowledge). Of course, there are many different beliefs Janet might reasonably form on the basis of my having uttered the sentence “Prof. Parsons is teaching”. Some of these depend upon what else she believes. But I take it that, among these beliefs, there is one—namely, the belief that Prof. Parsons is teaching—that is the source of the others, which arise from it by various inferences. Or again, Janet may form quite different beliefs if she takes me, say, to have flouted conversational maxims: She may, in particular, come to believe something I have implicated, and it may be my intention that she should

¹⁴ Attempting to justify a similar approach, I once suggested (Heck, 1995) that, although it is certainly true that the success of a particular communicative exchange does not require (say) that Janet come to believe that Prof. Parsons is teaching, nor even that I believe he is, it is nonetheless true that such a transfer of information is the normal goal of many ordinary communicative exchanges. Strawson (1971) expresses a related concern, which he puts by saying that it is difficult to know how we might characterize the act of *saying* without invoking something like communicative intentions. I would like to think I no longer need to hold that sort of view, but I admit to still finding it attractive. Much thanks to Brett Sherman for help here.

do so. Here again, however, I take it that Janet needs to recognize what I have said—namely, that Prof. Parsons is teaching—if she is to be able to determine what I have implicated. So, to put it as precisely as I can put it, the question I want to raise is why the belief that Prof. Parsons is teaching is the one Janet would acquire if she were to accept my offer of information—the one I made when, intending to be understood literally, I uttered “Prof. Parsons is teaching”—entirely because she accepted that offer, literally understood. But I shall speak in terms of the simpler formulation.

An obvious suggestion is that what explains Janet’s acquiring the belief that Prof. Parsons is teaching when I utter “Prof. Parsons is teaching” is that, as a sentence of English, this sentence *means* that Prof. Parsons is teaching. If so, then facts about what sentences mean in our common language, English, are implicated in the explanation of successful communication, just as Dummett thinks. This obvious suggestion, however, cannot be right. What the sentence in fact means cannot explain why Janet forms the belief she does. Had Janet believed that “Prof. Parsons is teaching” meant that Benny Parsons was racing, she would have formed the belief that Benny Parsons was racing, rather than the belief that Prof. Parsons was teaching. What explains the belief Janet forms is not what the sentence means but what she *takes* it to mean.

Similar remarks can be made about the speaker—me, in our example. What Janet comes to believe does not, of course, depend upon what I take my words to mean. But if I am successfully to convey my belief that Prof. Parsons is teaching to Janet, I must utter a sentence that expresses that belief. Had I intended to inform Janet that Benny Parsons was racing, uttering the sentence “Prof. Parsons is teaching” because I believed that was what it meant, then, while Janet might still have formed the belief that Prof. Parsons was teaching, she would not have acquired the belief that I had intended to communicate to her. Moreover, her new belief would not have constituted knowledge, even if I myself knew that Prof. Parsons was teaching, for it would not have been formed in the right sort of way. So again, what a sentence actually means, as a sentence of English, cannot on its own explain communicative success. What the sentence means to the speaker and hearer must also be involved. But, of course, it does not follow that what the sentence means in English does no explanatory work, and there is reason to think that what sentences mean to individual speakers cannot wholly explain

communicative success, either.

Suppose that, by some weird accident, Janet and I are under the same misimpression: By some coincidence, we both think that “Prof. Parsons is teaching” means that Benny Parsons is racing. Then while Janet will, in this case, acquire the belief I intend to communicate, that belief still would not constitute knowledge, even if I know that Benny Parsons is racing. Something still seems wrong about how she acquires the belief. It does not, therefore, seem enough to explain communicative success simply to note that Janet and I both *believe* that “Prof. Parsons is teaching” means that Prof. Parsons is teaching. Successful communication requires something more than mere agreement about what the uttered sentence means. It requires more than mere coincidence of belief—or, as we might put it, following Dummett, a mere overlap of idiolects.

What, then, distinguishes these two cases? Well, the obvious suggestion this time is that the crucial difference lies in what the sentence actually means, as a sentence of English: In the one case, our shared belief about what the sentence means is *true*, whereas in the other case it is not. But the obvious suggestion is once again mistaken.

Suppose Janet utters the sentence, “I bought a goat today”. Suppose, however, that I do not hear the first sound of the word “goat” clearly, so that I am unsure whether she has uttered that sentence or instead “I bought a coat today”. Knowing that Janet has been wanting a goat for a long time, and hearing the joy in her voice, I might come to believe that she has uttered the sentence “I bought a goat today”, and my belief may be true, even justified. But it still need not be knowledge. If not, then I do not *know* that she said she bought a goat (rather than a coat) and so cannot come to *know*, though I might reasonably come to believe, that she bought a goat. But then, if Janet does not *know* that “Prof. Parsons is teaching” means that Prof. Parsons is teaching, but only truly believes that it does, then she will not *know* that I have said that Prof. Parsons is teaching, but only truly believe that I have done so, and so she cannot come to *know* that Prof. Parsons is teaching on the basis of my telling her so. Successful communication—at least in so far as it involves the potential transfer of knowledge—therefore seems to require not just that the communicants have *true beliefs* about what the uttered sentence means, in that context, but that they *know* what the uttered sentence means.

Tyler Burge, in his discussions of testimony, seems to deny an important presupposition of the preceding discussion (Burge, 1993, 1999). I

am supposing here that beliefs acquired by testimony—that is, beliefs that are based upon what someone has said—are, in part, justified by beliefs about what words were uttered and what those words meant. In particular, on my view, Janet’s belief that I said that Prof. Parsons is teaching is (in an epistemologically relevant sense) based upon her beliefs that I have uttered the sentence “Prof. Parsons is teaching” and that this sentence means that Prof. Parsons is teaching. According to Burge, however, beliefs about utterances and about the meanings of uttered sentences, though perhaps *causally* necessary if one is to know what has been said, play no justificatory role with respect to beliefs acquired through testimony. Rather, the perception of utterances plays a merely *enabling* role, much as the perception of a written proof, though it may be necessary if one is to think through the proof, is no part of what justifies one’s belief in the proof’s conclusion.

Now, I accept the traditional distinction upon which Burge is relying, but I do not think that it applies to the case of language comprehension. I have discussed this issue elsewhere (Heck, 2005) and so will be brief here. It is true, of course, that we rarely notice the words other speakers have uttered. But Burge’s point is not phenomenological; it is epistemological. Consider again, then, the example just mentioned of Janet’s utterance of “I bought a goat”. Given the context, it may be clear enough that she must have meant to say that she bought a goat, not a coat. And frequently, that would be the end of the matter, since it need not matter what sentence Janet uttered nor, for that matter, what she actually said. But sometimes it does. In that case, reason to think she meant to *say* that she bought a goat would be good reason to think that she meant to *utter* “I bought a goat”. That, in turn, may be good reason to suppose she did utter “I bought a goat” and so did say that she bought a goat. But what Janet *did* say depends upon which sentence she actually uttered, and ordinary speakers know as much. If I am given reason to suppose that Janet actually uttered the sentence “I bought a coat”, then I am thereby given reason to suppose she said that she bought a coat, since, if she did utter that sentence, then that is what she said, whether she meant to say it or not.¹⁵ In so far as we want to distinguish the context of discovery from the context of justification, then, beliefs about the words someone has uttered seem more at home in the context of justification than in the context of discovery. As said, we

¹⁵ I am assuming here again that Janet has no unusual beliefs about what the relevant words mean. The importance of this assumption will emerge later.

often hardly notice the words our communicative partners have uttered, since our interest is rarely in the words themselves. It does not follow, however, that one's belief about what has been said is not, in an epistemologically relevant sense, based upon one's belief about—or perhaps better, one's perception of—the words uttered. Ordinary speakers are aware that what someone has said is determined by the words she has uttered. When the question is seriously raised what someone has said, we know that, while quite diverse evidence, including evidence about the speaker's intentions, might be brought to bear upon the question, all such evidence must ultimately bear upon what sentence was uttered, for it is the sentence uttered that determines what was said.

The foregoing gives us some reason, then, to suppose that successful communication depends upon both the speaker's and her audience's knowing what the uttered sentence means in the context in which it is uttered. Now, knowing that the meaning of the uttered sentence is (in our example) that Prof. Parsons is teaching is, of course, more than truly believing that it means that Prof. Parsons is teaching. But one can only know if one truly believes. And if such beliefs are to be so much as *capable* of being true, then it would certainly appear as if there must be something the sentence really does mean: Without that, there would be nothing for speakers to be right or wrong about. And so it would seem that what Janet and I both need to know, if we are to communicate successfully, is what the sentence I uttered means in the language we both speak—in English, or in the dialect thereof spoken in Cambridge, or whatever the right 'common language' might turn out to be. Say it's English. Then for me and Janet to know the meaning of the uttered sentence is for us to know what it means in English.¹⁶ If so, then, once again, reference to common languages and what sentences mean in them is essential to the explanation of successful communication.

Although, so far as I know, Dummett never argues in quite these terms, I think the foregoing constitutes a reasonable interpretation of his thought on these matters. I shall now argue, however, that the reference to English—or, more generally, to common languages—is explanatorily idle.

Consider the following well-worn example. I am one of those unfor-

¹⁶ When one puts what is at issue in these terms—whether common languages are the objects of linguistic beliefs—it seems as if the view I am attacking may also be held by Jim Higginbotham (1989). But I am not sure Higginbotham does hold any position my arguments here would undermine.

tunate people who used to believe that the word 'livid' meant *flushed* rather than *bluish-gray*. So I might have uttered the sentence "Jones was livid" meaning to convey the belief that Jones was red-faced. Now, certainly, what I took the word to mean is not what a dictionary will tell you it means. And in that sense, of course, I was mistaken about what it means. On the other hand, however, as the familiarity of the example suggests, I was hardly alone in this misconception. (Perhaps you too shared this misconception until very recently.) Imagine, then, that most of the people with whom I regularly communicated were under the same misimpression as I was. And suppose I wished to convey to my friend Steve my belief that Jones was, as I would then have put it, 'livid', that is, flushed. Steve shares my belief that "Jones was livid" means that Jones was flushed. So if I say to him, "Jones was livid", and he takes me at my word, he will form the belief that Jones was, as he would put it, 'livid', that is, flushed. Is there a failure of communication here? I do not see why we must say there was. But if not, then our joint failure to know what the word 'livid' means in English—and so what the sentence "Jones was livid" means in English—did not frustrate our attempt to communicate. If not, then knowing what the uttered sentence means in the 'common language' cannot be a condition of successful communication and knowledge of meaning is not what explains successful communication.

Now, I myself just argued that mere agreement about meaning is not enough to explain successful communication. In particular, while Steve and I agree about what the word 'livid' means, in the sense that we both think it means *flushed*, mere agreement of this sort is not enough to explain our communicative success. So one might wonder how this example differs from the one discussed earlier, in which both Janet and I were under the misimpression that "Prof. Parsons is teaching" meant that Benny Parsons is racing. The difference is that, in that case, I stipulated that it was an *accident* that the two of us agreed about the meaning of the sentence. And in this case, too, if one has the intuition that it is just an accident that Steve and I agree about what the word 'livid' means, then one should, I think, also have the intuition that we have not really succeeded in communicating with one another. But it need not be an accident that Steve and I agree about what the word 'livid' means. Of course it might be. But there is no reason to suppose it must be. It does not, in particular, follow from the fact that 'livid' does not mean *flushed* in English that it is an accident that Steve and I

agree that ‘livid’ means *flushed*. The belief that ‘livid’ means *flushed* is sufficiently widespread that the third edition of *Webster’s New World Dictionary*, published in 1988, actually notes that it is sometimes used to mean *red*.¹⁷ We can imagine, then, that we stand in the middle of a process of language-change. The word ‘livid’ might be on its way to meaning *flushed*, rather than *bluish-gray*, at least in some of its uses, as a result of people’s using it to mean just that. How else would one imagine such a change occurring? And such changes do, of course, occur. In the midst of such a change, or even at the beginning of such a change, one would expect to find groups, some perhaps quite isolated, in which the word is already consistently used with what will one day become one of its meanings in the ‘common language’. Within such groups, communication will proceed quite normally, divergence from ‘standard usage’ notwithstanding.

One might suggest, however, that Steve and I do know what the word ‘livid’ means in our common language. It is just that our common language is not English but rather the language spoken by some smaller community of which we are part. But how should this smaller community be specified? The point is not just that it is hard to say. It is that once we start shrinking the linguistic community, there will be no unique community to which both Steve and I belong. We both belong to many linguistic communities, and which one is relevant may depend upon which words are being used. Although Steve and I agree about what the word ‘livid’ means, we may disagree about what some other words mean, in much the same way we disagree with other speakers about what ‘livid’ means. Maybe Steve applies ‘fish’ to all aquatic creatures (so that Steve will say ‘Whales are fish’),¹⁸ though I apply it only to aquatic creatures that share biological characteristics I cannot myself enumerate. Maybe Steve is part of a linguistic community all (or most) of whose members use the word ‘fish’ as he does and so communicate perfectly well using the word in that way. If so, then I see no reason not to say about Steve’s use of ‘fish’ within that community what I just said about our joint use of ‘livid’. If so, then, as I claimed, there is no single linguistic community to which both he and I belong, any more than there is a single non-linguistic community to which we both be-

¹⁷ Many people with whom I’ve discussed these issues have confessed that they thought the word ‘livid’ meant *angry*.

¹⁸ The third edition of *Webster’s New World Dictionary* also lists this usage as a ‘loose’ one.

long. Moreover, we do, of course, talk to others who do not share our belief about what ‘livid’ means, say, Janet, who agrees with *Webster’s* that it means *bluish-gray*. So when I say to Janet, “Jones was livid”, communication will fail. In this case, of course, the explanation seems simple enough: We just do not agree about what ‘livid’ means. But is one of us wrong about what it means in our ‘common language’? Which language—whose language—would that be?

I do not deny, of course, that there are intuitions here that must be respected: I do not deny, in particular, feeling a strong pull to say that *I* am the one who is wrong about what ‘livid’ means. The next section offers an account of these intuitions. I shall return, in section 4, to the question how communicative success should be explained.

3 Semantic Norms

I used to think that ‘livid’ meant *flushed*. Had I then said to Janet, “Jones is livid”, she, not sharing my misconception, would have formed the belief that Jones was livid, that is, bluish-gray (or pale), rather than the belief I was trying to communicate to her, namely, that Jones was red-faced. My communicative intention would thus have been frustrated. Indeed, Janet, upon discovering that Jones was not pale, and that I had known he wasn’t, might have accused me of lying. Now, I wasn’t lying. I made every effort to speak the truth. Nonetheless, there is a strong intuition that it is my fault that Janet acquired a false belief. I told her that Jones was livid, and he was not.

The intuition that facts about what sentences mean have normative force is deeply ingrained. I am at fault, not Janet, because I *misunderstood* the word ‘livid’ and so used it *wrongly*. But if there is no ‘common language’ that we both speak—if there is no such thing as what the word ‘livid’ means in that language—what content is there to the claim that I was using it wrongly? If, moreover, I am not wrong about what the word ‘livid’ means, why should I feel compelled to change my usage? But we often do feel so compelled. Upon learning of my disagreement with the dictionary, I did not continue to use ‘livid’ as I had. Nor would I have regarded it as reasonable or rational for me to do so. Similarly, consider Bert, who insists that the word ‘arthritis’ applies to all rheumatoid ailments.¹⁹ Bert says “I have arthritis in my thigh”, believing that what

¹⁹ Let me emphasize that it is important to my treatment of this example that Bert

he says means that he has a rheumatoid ailment in his thigh, which he may well have. But as the literature's response to such examples makes plain, there is a strong intuition that what Bert has said is false, whether or not he has a rheumatoid ailment in his thigh, because the word 'arthritis' in fact applies only to rheumatoid ailments of the joints. If that is not what the word means in our common language—because there is no common language and hence nothing that the word means in it—how can we say that Bert speaks falsely?

Chomsky, for his part, accepts these consequences almost gleefully. He does not, of course, say that there are no linguistic norms. He would insist, for example, that I can indeed be wrong about whether a particular sentence is grammatical. Whether the sentence is grammatical is not determined by whether I think it is but by the grammatical principles I tacitly know. There is no reason to suppose that what those principles imply will always be transparent to me. (Garden path sentences, such as "The horse raced past the barn fell", are a typical sort of example.) For similar reasons, I could be wrong about what some sentence means. If I think that "No eye-injury is too trivial to ignore" is a sensible thing to put on a sign in a hospital—as did the people who apparently put it on just such a sign—then I am again just wrong.²⁰ These possibilities of error derive, however, from the familiar gap between competence and performance: That these beliefs are erroneous can be explained without invoking anything beyond my own linguistic competence. What Chomsky means to deny is that there are any sources of linguistic norms—any sources, that is, of what can rightly be called 'error'—that lie outside a speaker's own linguistic competence. Perhaps better: Any such norms as there may be are merely hypothetical, depending for their force upon, say, the speaker's desire to conform her usage to that of others. That is not, of course, to say that such a desire cannot be, or even typically is not, rational, but only that it is not required by anything that is specifically linguistic.

But does rejection of the notion of a common language really lead to such radical conclusions? I am going to argue that it does not.

actually thinks the word applies to all rheumatoid ailments, rather than being ignorant about what its extension is. So the case is a Gödel-Schmidt case not a Feynman case ?, pp. 81ff. Feynman cases raise different issues, which I hope to address elsewhere.

²⁰ As a little thought will show, the sentence actually means that no eye-injury is so trivial that one would not be justified in ignoring it: That is, it's OK to ignore all eye injuries, even the trivial ones. (I learned of the example from George Boolos, who told me he had heard it from Jonathan Bennett, but I have not located an original source.)

The following passage is reasonably representative of Chomsky's many discussions of this matter:

If Bert complains of arthritis in his ankle and thigh, and is told by a doctor that he is wrong about both, but in different ways, he may (or may not) choose to modify his usage to that of the doctor's. . . . If my neighbor Bert tells me about his arthritis, my initial posit is that he is identical to me in this usage. I will introduce modifications to interpret him as circumstances require; reference to a presumed 'public language' with an 'actual content' for *arthritis* sheds no further light on what is happening between us, even if some sense can be given to the tacitly assumed notions. Chomsky (2000a, p. 32)

Chomsky here denies that Bert means by "I have arthritis in my thigh" what the doctor means. If not, then there is, presumably, no peculiarly linguistic reason Bert need modify his usage to conform with the doctor's. Of course, there may be reasons he should do so, even quite compelling ones, but, if there are, these reasons derive not from the nature of Bert's language but from something external to it. It is, moreover, important to see that Chomsky's main point is not that no sense can be made of the notion of a public language: It is that, even if some sense can be made of it—which he, of course, doubts—appeal to such a notion would be unexplanatory.

In his intriguing defense of the primacy of idiolects, Alexander George argues that we should distinguish what a word means in a person's idiolect from what she *thinks* the word means in her idiolect. If there is such a distinction, then we get much of the benefit of the appeal to common languages without having to invoke them. In particular, we can make sense of the phenomenology of linguistic error, which George describes thus:²¹

When I came to believe that 'livid' does not mean *red* but rather *bluish-gray*, I then took myself to have been in error, to have made a mistake about the meaning of the word 'livid' in my language. When I thought back to times I had used the word, I felt I was discovering what I had actually said, as opposed to what I thought I had. On [Chomsky's] view, however, such reactions are out of place. George (1990, p. 289)

²¹ George calls the view in question the 'no-error view'. I've interpolated the reference to Chomsky, which is implicit in the paper.

The intuition here is that there is a difference between what my language is and what I take it to be. But we can have that difference, the thought is, even if there no such thing as a common language. The gap George wants us to acknowledge is not, it should be noted, the one generated by the distinction between competence and performance. On the contrary, what he wants is a distinction that would allow me to be wrong about what the word ‘livid’ means in my idiolect not because of a performance error but rather because of a competence error: To do so, he thinks, we need to find a way to say that what I think the word ‘livid’ to mean need not be what it means, *even in my own idiolect*.

The question immediately arises, of course, what determines what ‘livid’ means in my idiolect, if it is determined neither by what I (tacitly) take it to mean nor by what it means in a communal language I allegedly speak. George quickly admits that he does not know the answer to this question. He suggests, however, that the meaning of a word “is sensitive to considered changes to one’s linguistic beliefs that one would make as a result of communication with others or observation of them” George (1990, p. 292). As George notes, there is considerable unclarity in a notion so characterized, largely due to unclarity surrounding the relevant counterfactuals. But let that pass. The more significant worry, it seems to me, is that a notion so characterized is of no explanatory use, even for the limited purposes for which George wants it. George wants to be able to say that I am mistaken about the meaning of the word ‘livid’ if, and only if, what I believe it means is not what it in fact means, even in my own idiolect. But when unpacked, that turns out, according to George, to mean (very roughly) that I am mistaken just in case I would change my view about what ‘livid’ means if I had the right sorts of interactions with other speakers. If so, however, then we cannot say, as George would apparently also like to be able to say, that I should change my view *because it is wrong*. Rather, it is wrong only if I am (and in suitably similar circumstances would be) prepared to change it.

George comes close to acknowledging this point, and it may just not bother him. The reason is that, like Chomsky, George does not think that idiolects, in his sense, are of any interest to linguistics George (1990, p. 295). His purpose is not to rehabilitate a notion of common language—or some ersatz—that will be suitable to figure in scientific investigations of language. He is simply trying to gain some purchase on normativity from within the perspective on language that informs contemporary linguistic theory. Nonetheless, George does seem to want

his notion of an idiolect to do some explanatory work. Thus, he concludes his paper as follows:²²

... [C]ertain features of grammars over which speakers might have some control confront a seemingly objective reality that determines the correctness of the relevant linguistic beliefs. This reality partly consists in actual and potential considered transformations of individual grammars consequent upon communicative interactions among their bearers. That we take ourselves to be constrained in the interpretations we can place upon expressions if we are to proceed correctly is [an] important aspect of the perceived objectivity and independence of language. George (1990, p. 297)

But why should what I *would* do if had certain sorts of interactions with others so much as seem to constitute an “objective reality”? How can I regard what I would do in such circumstances as correct if what is correct is determined by what I would do? Here—to borrow from Wittgenstein—it really does look as if what seems right is right.

In discussing the passage from Chomsky above, I remarked that he is not claiming, of course, that one cannot have reason to change one’s usage. His claim, rather, is that the reasons one can and often does have are not *linguistic*: They have nothing particular to do with the nature of language but are reactions to external forces, for example, peer pressure. This observation gives us another reason to be dissatisfied with George’s response to the problem of linguistic error. One can have a variety of reasons to change one’s linguistic beliefs as a result of interaction with others. Changes one would make because of peer pressure, however, do not seem appropriate constraints on the identity of one’s idiolect. Now, George seems aware of this point, too, writing that such changes must be “considered”, by which he means that they should “result from our reflection on the correctness of our [linguistic] beliefs” George (1990, p. 289). But the problem, once again, is that it is difficult to see how to understand such reflection if correctness is what George says it is. If my beliefs are only correct if I would not make considered changes to

²² The quotation is consistent with the suggestion that George thinks that there really is no correct way to proceed and that such constraints are ultimately illusory, though we do “take ourselves” to be so constrained and do suppose that there is a right and wrong about how we use our words. The overall tone of the paper suggests to me, however, that this is not George’s view.

them as a result of interaction with others, and if what I am trying to decide, when reflecting on the correctness of my beliefs, is whether to make just such changes, I seem to be going in a circle.

The problem to which we keep returning derives from George's characterizing the semantic properties of idiolects in purely dispositional terms. That these properties are explanatorily impotent should thus be no surprise: An object's possessing a property that is characterized in wholly dispositional terms cannot explain its exercise of the very dispositions in terms of which that property is characterized: That a pill is dormative cannot explain why it makes you sleepy if its being dormative *just is* its tending to make people who take it sleepy.²³

Consider again the example of Bert and his doctor. Suppose this time not that Bert says to the doctor, "I have arthritis in my thigh", but that the doctor says to Bert, "You do not have arthritis in your thigh". Bert thinks to himself, "Man, what a relief!" In fact, Bert has absolutely nothing about which to be relieved, for the doctor continues, "You do, however, have a very serious rheumatoid ailment in your thigh".

Chomsky would have us say that Bert is, if he wishes, free to continue using the word 'arthritis' as he does. And, in some sense, presumably, he is. But my intuitions are somewhat different in this case than they were in the earlier one. When it is Bert who is speaking, there is at least some temptation to say, with Humpty Dumpty, that the question is who should be master: If Bert hardheadedly insists that he's always used the sentence "I have arthritis in my thigh" to mean that he has a rheumatoid ailment in his thigh and he isn't going to stop now—that may be silly, but it doesn't seem incoherent. But if, on the other hand, Bert insists that he's always understood people who've told him "You do not have arthritis in your thigh" to mean that he does not have a rheumatoid ailment in his thigh, and he's gonna keep right on understanding them that way—then I think he's not just being silly and is being incoherent. As the example makes plain, Bert's failure to understand the doctor makes him liable to form false beliefs not just about what the doctor has said but about whatever may happen to be the subject of conversation: health, sports, politics, or the weather. The intuition that Bert misunderstands the word 'arthritis' is thus, pretty much as George suggested, the intuition that Bert would have good

²³ Of course, the pill's being dormative would explain its making you sleepy if dormativity had a categorical basis, as fragility familiarly does. But it is definitive of the position George is defending to deny that they do.

reason to change how he understands the word if he had certain sorts of knowledge about other speakers:²⁴ Understanding 'arthritis' as he does makes the beliefs Bert forms on the basis of what others say to him using that word liable to be false—or, at least, deprives those beliefs of any likelihood of truth, even when conditions are otherwise ideal (for example, his informant is speaking knowledgeably).

If the goal were to rehabilitate a notion of common language that will do serious explanatory work, then we would, I am happy to admit, remain far short of it: The question how the boundaries of linguistic communities are to be drawn remains, as do the various sorts of relativity about which I complained earlier—relativity to conversational partners and even to the words one is using. But the goal was simply to account for a certain intuition—the intuition that Bert misunderstands the word 'arthritis'—without appealing to common languages, and I say I have done that. The intuition has its source in the fact that, when language is used for the purpose of exchanging information, how one understands an utterance determines the content of the belief one acquires if one accepts it as true and so whether the belief so formed is likely to be true, even under otherwise ideal conditions.

It should be no surprise that this account of semantic norms makes reference to communication: There is obviously no question of conforming one's usage to that of other speakers unless there are other speakers to whose usage one might conform, and there will be no need for conformity unless one is interacting with those speakers linguistically. And assuming that Bert is communicating with his doctor does not trivialize the account I have offered: Even if we do so assume, it is not immediately obvious why Bert should then be subject to any semantic norm. If my view were that Bert is subject to semantic norms only insofar as he wants to understand his doctor and be understood by him, then that, I would be happy to agree, would be entirely empty. But that is not my view. My view is that Bert is subject to semantic norms insofar as (i) Bert is exchanging information with his doctor and (ii) Bert wants to have true beliefs. If someone wanted to insist that (i) and (ii) are conditions whose satisfaction is not required by linguistic competence and so that semantic norms are not purely *linguistic* norms, I would not disagree. I do not see why it should matter, one way or the other.²⁵

²⁴ Of course, Bert might be so hard-headed that he won't change how he uses any word for any reason, but the intuition is that there are certain sorts of semantic norms.

²⁵ Thanks to Tad Brennan for some pressure here that took a long time to take effect.

One might now object, however, that no reason has been given to suppose that Bert uses the word ‘arthritis’ *wrongly*. When Bert discusses his health with his doctor, they misunderstand *each other*. Why think the doctor is right and Bert is wrong? It seems to me, however, that the strength of the intuition that Bert is wrong varies with the degree of his idiosyncrasy. It is because Bert disagrees not just with the doctor but with just about everyone else, too, that we have such a strong sense that he is mistaken. But if one thinks about a case like that of the word ‘livid’, intuitions aren’t nearly so strong. At least, mine aren’t. In fact, I’m not at all sure, as I look back on my earlier usage of the word, that I feel as if I know now what I was actually saying then: If my conversational partners also took the word ‘livid’ to mean *flushed*, then they interpreted me as saying just what I then intended to say.

4 Communicative Success

In section 2, I argued that communicative success is not to be explained by communicative partners’ knowing what uttered sentences mean in some common language they all speak. I argued that this proposal will not do, because there is no independent standard to which to hold beliefs about meaning. At least, if there is one, it is not the meaning of the words in some common language. Mere agreement about meaning, on the other hand, is insufficient to support successful communication, at least if communicative success is supposed to make it possible for one to come to know what one is being told. But where is the middle ground? It is not enough for speakers to agree about what a given sentence means, because mere agreement can be accidental: And if it is a mere accident that Steve and I agree about what “Jones was livid” means, then the truth of Steve’s newly acquired belief will itself be accidental, even if I know that Jones was livid. Requiring that both Steve and I must *know* what “Jones is livid” means will serve to prevent such accidents, but

For what it’s worth, I am inclined to think that the desire that one’s beliefs be true—if one should call it that—is a precondition of rationality. And, in any event, semantic norms are no worse off than epistemic norms, or even logical norms, if they are conditioned by that desire. The more interesting question concerns the status of the assumption that Bert is exchanging information with his doctor. As it happens, I am inclined to think that communication most fundamentally is a means of exchanging information. But even if that is not so, it doesn’t matter for present purposes: Absent the assumption that Bert is exchanging information with his doctor, I simply do not have the intuition I have set out to explain, namely, that Bert ought to use his words in a certain way.

that proposal requires there to be something the sentence really does mean in our common language, and there is no common language that we share. But there is actually no need to appeal to common languages to implement the main idea behind this proposal: It is enough to require that Steve and I *knowingly agree* about what the sentence means. If so, it will be no accident that Steve’s newly acquired belief is true if mine is.²⁶

Note that I did not say that Steve and I must *know that we agree* about what “Jones is livid” means to communicate successfully.²⁷ That would suggest, surely wrongly, that one must have the concept of *agreement* to be able to communicate. Of course, one might well want to know what this notion of *knowing agreement* is supposed to be, and I confess that I am not entirely sure. Part of what is wanted is that I should know what Steve means by “Jones is livid”, and he should know what I mean. But, for reasons I shall not rehearse, it is likely also necessary that I should know that Steve knows what I mean, that he should know that I know what he means, and so forth Grice (1989a): What I should probably say, then, is that *common knowledge* between me and Steve each of us means that Jones is flushed by “Jones is livid”.²⁸ But I do not want to pursue that issue here. I therefore leave my view a bit underspecified: Communicative success depends upon and should be explained in terms of communicative partners’ knowingly agreeing about what uttered sentences mean.

One might object that this view makes no distinction between ordinary cases of communication and the following sort of case. Suppose Bob and Patrick have agreed that today is opposite day: They are always to say the opposite of what they really mean. In particular, when Bob says “Gary is not home”, he knows that Patrick will take him to mean

²⁶ There are various reasons to think knowledge may not be quite the right notion here. For one thing, as I’ll note below, one’s apprehension of meaning is most immediately perceptual, and perceptual states are not the sort of thing that can be classed as knowledgable or otherwise. But if it turns out that we need some analogous notion that applies to perceptual states, one should be available.

²⁷ Brett Sherman has urged me to say what it might have seemed obvious I was about to say, namely, that Steve and I must *non-accidentally* agree about what the sentence means, rather than to say we must know we agree and then derive the non-accidental character of our agreement from that fact. The problem, however, is that it’s not obvious that non-accidental agreement about meaning is enough to secure knowledge in ordinary communicative exchanges.

²⁸ See ? for the classic discussion of common knowledge. The notion also figures importantly in Grice’s work.

that Gary *is* home. If so, then it would seem that Bob and Patrick knowingly agree that Bob's utterance of "Gary is not home" means that Gary is home. Now, in a sense, that's fine with me: Communication may indeed succeed in such cases; one wants to be able to explain why it succeeds; and that seems a pretty good explanation. But in another sense, something does seem to have been left out, namely, what distinguishes this case from the normal case. One might suspect that what distinguishes them is that "Gary is not home" just doesn't *mean* that Gary is home, however Bob and Patrick might have agreed to interpret it, and when one talks about what the sentence means, one is talking about what it means in English, or whatever Bob and Patrick's common language might be. As David Wiggins puts it:

Normally, it may be said, [a speaker] aims to be understood. But that isn't quite right. A speaker aims to be understood, but not as saying just anything. He aims to get across a certain thing, but he aims, if possible, to say this thing *and* to be understood as saying it. . . . If there were no such thing as a [common] language, what would be the difference between someone's simply being understood. . . by an audience as wanting to communicate that such and such and someone's successfully. . . *saying* this or that? Wiggins (1997, p. 504)

Without the notion of a common language, Wiggins suspects, we cannot distinguish merely managing to get it across that Gary is home from getting it across by *saying* that Gary is home. This distinction certainly needs respecting, but I do not think we need to invoke common languages to respect it.

As John McDowell has emphasized McDowell (1998a,b), when we hear another speak, we do not hear just her words, adding an interpretation in thought. We literally *hear* what she says in her words. To hear what someone has said—in the semantic as well as the phonetic sense—is to perceive, and like other forms of perception, this one is in some ways similar to belief. When I see a scene before me, it looks as if objects really are arranged a certain way in the space around me. Visual perception, that is to say, is representational not only in the sense that it has truth-evaluable content but also in the sense that it represents the world as actually *being* a certain way.²⁹ Perception of speech in a language one

understands is, in this respect, similar. If my wife tells me, "There are moths in the closet", my perceptual system represents my wife not just as having uttered these words but as having *said* that there are moths in the closet. Similarly, when the doctor says to Bert, "You do not have arthritis in your thigh", Bert's perceptual system represents the doctor as having *said* that he does not have a rheumatoid ailment in his thigh. Bert's perceptual system thus misleads him in this case, for the doctor has said no such thing.

I can't change how a sentence sounds to me simply by deciding that it should sound some other way, any more than I can make myself see a scene otherwise just by deciding that things ought to look differently—and that is as true when 'sounds' has semantic import as it is when it just has phonological import. In that sense, judgements of meaning have their source outside our conscious minds. Their source, however, is not outside our minds altogether, in facts about what our expressions mean in some common language. Their source is, rather, our *unconscious* minds: what we tacitly know about meaning. I hear the sentence "Gary is not home" as meaning that Gary is not home, and use it to say that Gary is not home, because that is what my I-language determines that this sentence means to me. What distinguishes Bob and Patrick's conversations on opposite day from their conversations on other days is thus the source of the judgements about meaning on which those conversations depend. In normal cases of linguistic communication, the source of these judgements is one's linguistic competence. In abnormal cases, as on opposite day, these judgements have some other, partially non-linguistic, source.³⁰

This last point is actually overstated. Our judgements about what utterances mean almost never have a purely linguistic source. The reason is that almost every sentence we utter contains one or another expression whose interpretation depends upon the context in which it

³⁰ Similar remarks serve to distinguish literal communication from that involving implicatures. If Prof. Smith writes in his letter of recommendation that Jones is punctual and has good penmanship, and conspicuously declines to say anything else, then both Smith and his reader may know that Smith's utterance in some sense meant that Jones is not well-qualified for graduate study. But there is, as Grice taught us, an important sense in which Smith's utterance does not mean that, and this difference should be registered in our account of communicative success. Here again, I would say that, while Smith and his reader may well know that Smith's utterance in some sense means that Jones is a poor student, the source of this knowledge lies not in their linguistic competence but at least partially elsewhere.

²⁹ For more on this matter, see Heck (2000).

is uttered. So to be able to interpret almost any utterance, one needs to know relevant facts about the context in which the utterance is made. What we ought to say, then, is that what distinguishes normal communication from communication on opposite day is that, in the normal case, judgements of meaning depend as little as possible upon non-linguistic sources. A couple of assumptions are being made here: First, that opposite day does not itself present us with an example of context-dependence; and Second, that context-dependence can be brought under sufficient theoretical control that the qualifier as ‘as little as possible’ can be given a clear meaning. The first assumption, though worth noting, should be uncontroversial. The second is not at all uncontroversial. A first stab at explaining what ‘as little as possible’ is supposed to mean might be the following: only in so far as is necessary if the utterance is to be interpreted at all. But that constitutes only a very little bit of progress. To defend—or even to explain—this second assumption in any detail, one would have to take a stand on the very large and very difficult question how context affects what is said. I have no such stand to take. But if one were to adopt Jason Stanley’s view that context affects what is said only in so far as it fixes the assignment of values to free variables present in the logical form of the uttered sentence Stanley (2000), then ‘as little as possible’ could be taken to mean: only in so far as is necessary to assign values to variables present at LF. I do not, I should emphasize, think this second assumption requires any view as strong as Stanley’s to be correct. But it does require that it should be possible to isolate the contribution context makes to determining what is said—lest too many ordinary cases turn out to be too much like opposite day—and Stanley’s view illustrates what sort of requirement that is.

So, to a slightly better approximation, my suggestion is as follows: Communicative success depends upon and should be explained in terms of communicative partners’ knowingly agreeing about what uttered sentences mean; what distinguishes the ordinary case, in which one is not only understood as saying something but actually says it, is that speakers’ judgements about meaning derive entirely from their linguistic competence, except in so far as these judgements depend upon such knowledge of contextual features as is necessary if the utterance to be assigned any meaning at all.

The following might seem like a counterexample.³¹ Suppose my four-

³¹ Thanks to Stephen Schiffer for this example, which he mentioned as I struggled

year-old friend Sophie says to me, “I have a brother”. If I believe her, and if she does have a brother and knows that she does, then one would naturally suppose I might thereby come to know that Sophie has a brother. It is far from clear, however, that Sophie and I now have the same belief, because it is far from clear that we mean the same thing by ‘brother’. For me, *brother* is a biological notion: Someone is your brother only if the two of you have the same biological parents. Sophie knows nothing of biological relationships. Her concept of *brother* is more a social one. So our beliefs have different contents, and that makes it natural to suppose that our words do, as well. If so, then we do not even agree about what Sophie’s utterance of “I have a brother” means, let alone knowingly agree, but we nonetheless seem to have communicated, at least in the sense that I have acquired knowledge as a result of what Sophie said to me. If so, then successful communication does not even require agreement about meaning, let alone known agreement.

This example undeniably has some appeal. But it seems to me that it derives its appeal from the fact that the word ‘brother’ *isn’t* always used to express a purely biological notion, even by adults. When I say that my niece Julie has a brother, or my brother calls Julie his daughter and Alex his son, or I say that Don and Shannon are their parents, that’s all true, even though both Alex and Julie were adopted. There are perfectly good adult uses of such words as ‘brother’ and ‘parent’ on which they denote certain sorts of social relationships: Hence the term ‘birth-mother’.

In the last couple decades, developmental psychologists have constructed an increasingly compelling case that the concepts young children associate with certain words are importantly different from the concepts adults associate with those same words. And, as it happens, the concepts such children associate with words with which well-educated adults associate biological notions are very often social or functional.³²

to understand a question someone else had asked at NYU. I do not know the original questioner’s name but thank her anyway for asking the question.

³² See Carey (1985) for an extensive discussion of young children’s understanding of biological notions. A similar account of young children’s understanding of psychological notions can be found in Wellman (1990). One could presumably construct examples similar to Schiffer’s with many of the concepts Carey and Wellman discuss. Carey and Wellman’s work is not, of course, uncontroversial, and some philosophers have raised doubts about how they individuate concepts: See, for example, Fodor (1998). The really hard questions about how concepts are to be individuated don’t affect the discussion here, however, since the concepts in question differ extensionally.

In some of those cases, the meaning with which the child uses the word does not survive into adult language, and so examples constructed using such words would not suffer from the flaw I just claimed to find. So let's set aside the ambiguity of the word 'brother' and assume that it expresses a social notion for Sophie and a biological one for me. If so, then I don't think I *can* come to know that Sophie has a brother from her telling me that she has (as she would put it) a 'brother'. By hypothesis, Sophie would truly have believed that she had a 'brother' even if Sebastian were adopted, and I presume she would still have said so, too. But then it is only because Sebastian happens not to be adopted that my belief is true if hers is, and I'm just lucky not to be wrong. Granted, it's not sheer luck—as it would be had I had taken Sophie's utterance to mean that she had a Barbie—but the truth of my belief nonetheless seems too accidental to count as knowledge. If so, then communication wasn't successful after all, and there's nothing that needs explaining.

The matter could undoubtedly use further discussion, but it won't get it here.³³

5 Talking to Strangers

I communicate on a regular basis with all sorts of people I've never before encountered, and communication with such people frequently succeeds: When the police officer says to me, "Your car is parked in a fire lane", I can come to know that my car is parked in a fire lane and take appropriate action. If so, then, on my view, the officer and I must knowingly agree about what her utterance means and so, minimally, I must know what her utterance means. But on what basis might I claim such knowledge?³⁴

There are several options here. One would be to insist that I have inductive evidence that what sound like sentences of my language are

³³ One quick remark, however. First, suppose I knew that Sophie's parents had strange religious beliefs that barred them from raising any children but their own biological offspring. Then perhaps I could come to know that Sophie had a biological brother, since that rules out his being adopted. I mention the point because there may be similar cases in which the intuition that I know is strong enough to survive the challenges in the last paragraph. My knowing might then be explained along these sorts of lines, by appeal to relevant background knowledge.

³⁴ Let me register two very large debts here: One is to Chris Peacocke, who forced me to address this question directly; the other is to Kyle Stanford, who helped me understand what my view was.

used with the same meanings with which I would use them. Another option is to say that I have a default entitlement to suppose that what sound like sentences of my language mean what I would mean by them: The question would then arise why it is rational for me to rely upon this entitlement Burge (1993, 1999). Yet another option is externalist: One might deny that I have any warrant—even an entitlement—for such beliefs and claim that they constitute knowledge because they are reliable, or what have you. Fortunately, we need not choose among these options here.³⁵ All of them, in one way or another, require that, typically, what sound like sentences of my language are used with the meanings with which I would use them. To put the point differently: Any reasonable answer to this question must presume that different speakers tend to use the same sentences the same way. And one might well suppose that it is here that the friend of idiolects has a problem: If there is no such thing as a common language that both the officer and I strive to speak, why should there be any convergence at all? That, I suggest, is what is worrying Dummett when he writes:

An English speaker both holds himself responsible to, and exploits the existence of, means of determining the application of terms which are either generally agreed among the speakers of English, or else are generally acknowledged by them as correct. . . . [T]his has to be so if words are to be used for communication between individuals. Dummett (1978, p. 425)

Now, Dummett is well aware, of course, that people do not always understand or use their words the same way. But he is suggesting that it is only because speakers of English recognize shared norms governing their usage that we avoid the chaos of Babel.

I will now argue for two claims: First, that the appeal to common languages cannot do the work Dummett seems to think it can do; and Second, that the convergence of usage that is necessary if we are to be able to communicate successfully with strangers can be explained without any appeal to common languages.

To take the first point first, people who speak what anyone would have to regard as different languages frequently communicate successfully. Speakers of different dialects often communicate successfully, and between these dialects there may be significant syntactic and phonological

³⁵ As it happens, I don't much care for the first option, but I am no epistemologist.

differences, as well as semantic differences at the lexical level: If I may borrow an example from Gabriel Segal, in British English, something is rightly called a ‘pie’ only if it has a pastry top, whereas in American English, a ‘pie’ need not have a top at all. That does not imply that Bostonians can’t communicate successfully with Londoners. Communication will succeed so long as the parties to a given exchange knowingly agree about what the various sentences uttered during that exchange mean, and such agreement will be common so long as the speakers’ dialects do not differ too much in the areas that are typically used for communication between them. But then speakers’ knowingly agreeing about what a particular sentence means cannot be explained in terms of their speaking a common language: There need be no common language they speak, for their dialects can differ arbitrarily outside the parts they typically use to communicate.

The situation is no different, in principle, when I speak to the proverbial man on the street. In order for us to communicate successfully, we do not need to speak a common language. We need only agree about what the sentences that occur in our particular communicative exchange mean. And that’s a good thing, since there is, so far as I can see, no reason to suppose that I and all of my everyday conversational partners actually do speak the same language, in any reasonable sense: For all I know, the officer thinks ‘livid’ means *flushed* or, as may even be more common, *angry*. One could, of course, insist that one of us must be wrong, but I’ve already examined and dismissed that claim. One could try isolating some ‘core language’ and insist that successful everyday communication depends upon our all speaking *it*. But it should be obvious by now that the part of the language on which there must be agreement if communication is to be successful in a particular case will vary from episode to episode: No stable ‘core’ need exist, and there is no obvious reason to suppose one does.

The most that can be demanded, then, is that the officer and I speak a ‘common language’ as regards the sentences used in communication between us: We must—to modify Dummett’s remark—‘hold ourselves responsible to means of determining the application of terms we actually use, means that are either generally agreed among us or else are generally acknowledged by us as correct’. But what does that add to the claim that we must know that we mean the same thing by those sentences?

The appeal to common languages will not, therefore, help us under-

stand why speakers’ usage converges in the way it must if successful communication is to be possible. I shall now sketch an alternative explanation.

In a passage quoted earlier, Chomsky writes that, once Bert is informed that he and the doctor use the word ‘arthritis’ differently, Bert “may (or may not) choose to modify his usage to that of the doctor’s” Chomsky (2000a, p. 32). That strongly suggests that what Bert means by the word is under Bert’s control, and I doubt that is true.³⁶ When I was very young, I had a small piece of purple fuzz on my arm. Pointing at it, I asked my mother what it was, and she said, “That’s a muscle”. So for some time I went around using the word ‘muscle’ to apply to small pieces of fuzz. Later, I found out what had happened and learned that most people use the word ‘muscle’ to apply (as I’d now put it) to muscles, not to fuzz. That information did not, however, lead to an immediate change in what the word meant to me. Over time, of course, it did, but for quite a while, as I recall, I continued to *hear* the word ‘muscle’ as meaning *piece of fuzz*: I might have had reason to say to myself, “I know that Dad didn’t just say that working out has given him bigger pieces of fuzz, but that’s nonetheless how it sounded”. I was thus the victim of a kind of semantic illusion, one driven by features of my semantic competence that were not entirely within my control. And as with visual illusions, these sorts are at least somewhat persistent: Just as knowing that the lines in the Müller-Lyer illusion are the same length doesn’t make them look the same length, believing that ‘muscle’ means *muscle* doesn’t make it sound like it does.

Changing what the word ‘muscle’ meant to be was not something I could accomplish simply by force of will. Fortunately, however, I was not doomed to suffer from illusion forever and so be forced to make do by consciously reinterpreting other speakers. Repeated exposure to other speakers who used the word ‘muscle’ to mean *muscle* eventually led me to hear the word ‘muscle’ as meaning *muscle*, too. And, importantly for our present purposes, it is not obvious that I could have *prevented* this change simply by force of will, either. It seems to me that repeated exposure to speakers who meant *muscle* by ‘muscle’ would have led to such a change by itself: Whether I wanted to do so or not, I eventually would have begun to hear the word ‘muscle’ as meaning *muscle*.

³⁶ As it happens, George speaks of “features of grammars over which speakers might have *some* control” George (1990, p. 297, my emphasis), strongly suggesting that we do not have total control over what our words mean to us, but he makes nothing of the point.

Much the same thing can be said about how my own usage of the word ‘muscle’—that is, how I used it in my own speech, as opposed to how I heard it when it was used by others. Even after I had discovered that I was different from other boys, if I had wanted to inform someone about pieces of fuzz, it was only with some effort that I could stop myself from using the word ‘muscle’. But, eventually, that changed. The reason is very simple: There is not even the possibility of a gap between what one means by a word when one utters it oneself and how one understands it when it is uttered by someone else;³⁷ the semantic and other linguistic information upon which one relies when one speaks is the *same* information upon which one relies when one interprets the speech of others. The change in how I heard other speakers’ utterances of the word ‘muscle’ and the change in my own usage of it—what I meant by it when I uttered it myself—were, in fact, a single change.³⁸

The mentioned features of human linguistic competence thus together encourage, and may even enforce, semantic conformity. If Bert speaks regularly enough with other people about his ‘arthritis’, and if it becomes sufficiently apparent to him that he does not mean by it what they do, then how he understands and uses the word will change. Of course, whether, and how quickly, Bert’s usage of the term changes will depend upon many factors: How idiosyncratic his usage is, how often the word is used, by and with whom, and so forth. But surely that is as it should be: One does not want an ‘explanation’ of why the community of ‘English’ speakers is semantically homogeneous. What one wants is an account of why there is sufficient semantic homogeneity that it is possible (if not always advisable) to talk to strangers. Dummett’s answer to this question is that speakers have a common goal: They strive to speak a common language, recognizing shared norms as governing their usage. Dummett’s view is thus teleological in a way that mine is not. On my view, as a consequence of how human linguistic capacities function in communication, one’s own understanding of a term tends to track that of those with whom one regularly uses it in conversation. In virtue of connections between one’s own regular conversational partners and other speakers, one’s usage tends also to track theirs, the strength of

the tendency being a function of both the length and the nature of the connection.³⁹ More commonly used terms can be expected to be used more uniformly within a wider population, with less commonly used terms more likely to be understood differently within conversationally isolated groups. And so forth. It’s all very messy, but that, again, is surely as it should be.

6 Closing

The question whether idiolects or common languages are more fundamental is sometimes framed as the question whether language is ultimately *mine* or *ours*. I have argued that idiolects are necessarily involved in the explanation of successful communication and that the phenomena that make the appeal to common languages seem necessary have their source in general facts about human linguistic capacities and how they are deployed in communication. These same facts allow us to preserve a sense in which my language is always *our* language, despite the fact that there is no explanatory use to be made of common languages.

As noted earlier, speech is a form of rational action. This form of action is rationalized in part by one’s linguistic—and, in particular, semantic—beliefs.⁴⁰ So my uttering the sentence “Prof. Parsons is teaching”, for example, may be an intentional act, and my reasons for performing it may be something like the following:

- (5) I want Janet to believe that Prof. Parsons is teaching.
- (6) Janet believes that, in the present context, an utterance of “Prof. Parsons is teaching” means that Prof. Parsons is teaching, and she will deploy that belief in interpreting my speech.
- (7) If I utter “Prof. Parsons is teaching” in the present context, Janet will take me to have spoken the literal truth.
- (8) If I utter “Prof. Parsons is teaching” in the present context, Janet will (be in a position to) come to believe that Prof. Parsons

³⁷ *Modulo* context-dependence, of course. The possibility of the sort of reinterpretation illustrated by opposite day, or by Chomsky’s talk of ‘modifications’ to our ‘initial posits’ Chomsky (2000a, p. 32), or by Davidson’s discussion of malapropism Davidson (1986), are not counterexamples, since they do not concern how one *hears* what is said.

³⁸ A cognate point is made in passing by Chomsky Chomsky (1986, p. 14, note 10)

³⁹ The presence of mass-media, of course, greatly increases the number of one’s regular conversational partners.

⁴⁰ Let me say again that it is possible, indeed likely, that the relevant cognitive states are not all or always beliefs. They may be perceptual or, perhaps, of some other kind. What matters, for present purposes, is that they are conscious, in the sense that they are available for the explanation of intentional action, as here.

is teaching. [From (6) and (7)]

(9) I shall utter “Prof. Parsons is teaching”. [From (5) and (8)]

The important point about this rationalization is that my reasons for uttering “Prof. Parsons is teaching” make explicit reference not to what I think this sentence means but, rather, to what I think Janet takes it to mean. The same, of course, is true of my interpretation of her speech: If I want to learn from Janet, I need to consider what she takes her utterance to mean.

This rationalization suffers from a problem we encountered earlier: Nothing distinguishes it from rationalizations we might give of utterances made on opposite day. On opposite day, I would instead utter “Prof. Parsons is not teaching”, and part of my reason for doing so would be that Janet will believe that an utterance of it means that Prof. Parsons is teaching. What distinguishes normal and abnormal cases is again the source of one’s judgements about what one’s conversational partners will take utterances to mean. In the normal case, the source of the judgement is one’s linguistic competence;⁴¹ in abnormal cases, one’s judgement relies upon other sources as well. On opposite day, for example, my judgement about what Janet will take my utterances to mean is a product of my linguistic competence together with my knowledge of our agreement to say the opposite of what we mean.

What expressions mean to me is thus determined by what I tacitly know about the meanings of those expressions. But when this information is deployed, it is used, in the most familiar cases, to arrive at decisions about how one’s conversational partners understand utterances, both one’s own and theirs. And, as emphasized earlier, how I understand an utterance and how I unreflectively take my conversational partners to understand it are one and the same. From my point of view, then, my language is indeed your language: My language, as I understand it, just is your language, as I take you to understand it.⁴²

⁴¹ As we saw earlier, the source is not *solely* one’s linguistic competence, but similar remarks can be made here as were made there.

⁴² Extracts from this paper were read at New York University, in January 2003; at the University of Illinois at Chicago, in February 2003; at Brown University, in April 2003; at Yale University, in September 2003; at the University of St Andrews, in February 2004; and at the University of California at Irvine, in June 2004. The opportunity to present this material and receive the audiences’ comments greatly improved the paper. Questions from Justin Broackes, Bill Hart, Robert May, Chris Peacocke, Jim Pryor, Agustin Rayo, Stephen Schiffer, Kyle Stanford, Dan Sutherland, Brian Weatherston, Crispin Wright, and

References

- Burge, T. (1993). ‘Content preservation’, *Philosophical Review* 102: 457–88.
- (1999). ‘Comprehension and interpretation’, in L. Hahn (ed.), *The Philosophy of Donald Davidson*. Chicago, Open Court, 229–250.
- Carey, S. (1985). *Conceptual Change in Childhood*. Cambridge MA, MIT Press.
- Chomsky, N. (1965). *Aspects of the Theory of Syntax*. Cambridge MA, MIT Press.
- (1986). *Knowledge of Language: Its Origin, Nature, and Use*. Westport CT, Praeger.
- (2000a). ‘Explaining language use’, in Chomsky 2000d, 19–45.
- (2000b). ‘Language from an internalist perspective’, in Chomsky 2000d, 134–163.
- (2000c). ‘Naturalism and dualism in the study of language and mind’, in Chomsky 2000d, 75–105.
- (2000d). *New Horizons in the Study of Language and Mind*. New York, Cambridge University Press.
- Davidson, D. (1984). ‘Reality without reference’, in *Inquiries Into Truth and Interpretation*. Oxford, Clarendon Press, 215–225.
- (1986). ‘A nice derangement of epitaphs’, in E. Lepore (ed.), *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*. Oxford, Basil Blackwell, 433–446.
- Dummett, M. (1978). ‘The social character of meaning’, in *Truth and Other Enigmas*. London, Duckworth, 420–430.

a graduate student at NYU whose name I do not know were especially helpful. Thanks, too, to everyone who has attended the various courses and seminars at Harvard University during which I have discussed my evolving views on these issues, especially Bernard Nickel. Finally, conversations with Michael Rescorla and Jason Stanley did much to shape the paper, and I thank them for their on-going support of my work.

- (1993). 'Language and communication', in *The Seas of Language*. Oxford, Clarendon Press, 166–187.
- Fodor, J. (1998). *Concepts: Where Cognitive Science Went Wrong*. New York, Oxford University Press.
- George, A. (1990). 'Whose language is it anyway? some notes on idi-
olects', *Philosophical Quarterly* 40: 275–298.
- Grice, H. P. (1989a). 'Meaning', in Grice 1989c, 213–223.
- (1989b). 'Utterer's meaning and intentions', in Grice 1989c, 86–116.
- Grice, H. (1989c). *Studies in the Ways of Words*. Cambridge MA, Harvard University Press.
- Heck, R. G. (1995). 'The sense of communication', *Mind* 104: 79–106.
- (2000). 'Non-conceptual content and the "space of reasons"', *Philosophical Review* 109: 483–523.
- (2002). 'Do demonstratives have senses?', *Philosophers' Imprint*, 2. <http://www.philosophersimprint.org/002002/>.
- (2004). 'What is compositionality?', *Manuscript*.
- (2005). 'Reason and language', in C. MacDonald and G. MacDonald (eds.), *McDowell and His Critics*. Oxford, Blackwells, 22–45.
- Higginbotham, J. (1989). 'Knowledge of reference', in A. George (ed.), *Reflections on Chomsky*. Oxford, Basil Blackwell, 153–74.
- Lewis, D. (1985). 'Languages and language', in *Philosophical Papers, vol. 1*. Oxford, Oxford University Press.
- McDowell, J. (1998a). 'Anti-realism and the epistemology of understanding', in McDowell 1998d, 314–343.
- (1998b). 'In defense of modesty', in McDowell 1998d, 87–107.
- (1998c). 'Meaning, communication, and knowledge', in McDowell 1998d, 29–50.
- (1998d). *Meaning, Knowledge, and Reality*. Cambridge MA, Harvard University Press.
- Quine, W. V. O. (1970). 'Methodological reflections on current linguistic theory', *Synthese* 21: 386–98.
- Rumfitt, I. (1995). 'Truth-conditions and communication', *Mind* 104: 827–62.
- Stanley, J. (2000). 'Context and logical form', *Linguistics and Philosophy* 23: 391–434.
- Strawson, P. F. (1971). 'Meaning and truth', in *Logico-Linguistic Papers*. London, Methuen, 170–89.
- Wellman, H. M. (1990). *The Child's Theory of Mind*. Cambridge MA, MIT Press.
- Wiggins, D. (1997). 'Languages as social objects', *Philosophy* 72: 499–524.