☞ ALEX BYRNE
Interpretivism

1. Interpretivism introduced

In the writings of Daniel Dennett and Donald Davidson we find something like the following bold conjecture: it is an a priori truth that there is no gap between our *best judgements* of a subject's beliefs and desires and the *truth* about the subject's beliefs and desires. Under ideal conditions a subject's belief-box and desire-box become transparent.

To make this picture more vivid, let us introduce the familiar device of the *Ideal Interpreter*: an idealisation of a human being. The Ideal Interpreter, according to the bold conjecture we are considering, is capable of discovering exactly what you believe and desire. For now, let's be silent on the details of the Interpreter's knowledge and powers—with the following exception. The Interpreter's stipulated initial stock of knowledge had better not include knowledge of your propositional attitudes, else the conjecture will be trivialised.

Is this a kind of behaviourism? It does not have to be. For it may be that we wish to make our Interpreter take into account various non-behavioural facts when determining what someone believes or desires. For example, perhaps the discovery that someone is a hollow shell controlled from Mars would lead us to reject the hypothesis that he is a believer and desirer, no matter how convincing his behaviour. Very well, if there are such defeaters of the hypothesis that someone is a believer and desirer (or any other relevant inner facts), we can let our Interpreter know if any of them obtain in a par-

ticular case.[1] But both Dennett and Davidson think that the Interpreter does not need to open up the subject's head, and I shall not dispute this.

This intriguing picture of the mind—that, as an a priori matter, the facts about mental content are precisely captured by the judgements of some Ideal Interpreter—I shall call *interpretivism*.[2]

It is important to note that the Ideal Interpreter's presence (equivalently, the reference to our best judgements) is an indispensable component of interpretivism. Suppose it is (implausibly!) claimed that believing that snow is white is a priori equivalent to being disposed to produce the string 'snow

---

1.   The Mars example is taken from Peacocke 1983, p. 205. Peacocke imagines "The Body"—a debrained human body controlled by a computer from Mars which "has been given the vast but finite number of conditionals specifying what a typical human would do with given past history and current stimulation; so it can cause The Body to behave in any circumstances exactly as a human being would" (p. 205). Peacocke's Body, according to him, "does not have propositional attitudes at all: it is just a Martian marionette" (p. 205). I think Peacocke's verdict on the case debatable, but the important point is that even if Peacocke is correct, this does not spell disaster for the view which is the topic of this paper. It *would* spell disaster if the Martian puppeteers, with their evil intentions, were an essential component of Peacocke's example. For then any Ideal Interpreter would need to know whether her subject was controlled by *thinking* puppeteers, and this means that the Interpreter's database includes intentional facts. This problem cannot be avoided by taking the Interpreter to deliver the facts, with no intentional initial data, about whether the subject is controlled by thinking beings. The Interpreter might then be off on an infinite regress: the Martians themselves might appear to be Venusian marionettes, who in turn might appear to be Plutonian marionettes, and so on. In such a case, the Interpreter can have no final opinion whether the Martians are genuine thinkers, unless she knows whether the Venusians are genuine thinkers, and she cannot know *that* unless she knows whether the Plutonians are genuine thinkers...

Fortunately, the Martians with their evil intentions, and even the fact that the computer is distant from The Body, are surely dispensable (and misleading) parts of Peacocke's story. If the computer is shrunk to the size of a human brain, and placed in The Body, this *ought* to make no difference to our considered opinion about The Body's mental life. Or so I think.

2.   I borrow the term from Johnston 1993b. Interpretivism should be distinguished from what Dennett calls *interpretationism*, which "likens the question of whether a person has a particular belief to the question of whether a person is immoral, or has style, or talent, or would make a good wife" (1979, p. 15). Dennett takes interpretationism (avowedly not his own position) to involve some sort of relativism.

is white' (and likewise for other beliefs). This theory could be expressed to sound superficially like interpretivism, using an Ideal Interpreter who has as initial data the subject's dispositions to produce strings of symbols, and the relevant alleged a priori equivalencies. But here the Interpreter is otiose: she could be eliminated if desired, and so this theory is not a version of interpretivism.

It is worth noting that interpretivism holds the promise of elegantly vindicating the supervenience of intentional states on matters physical, with no reductionist assumptions. No commitment to reduction, because the Interpreter (a believer and desirer herself) is an essential part of the story. But a vindication of supervenience, because surely the Interpreter would not interpret differently without a physical difference in the subject of interpretation, or a physical difference in the subject's environment.[3] And as the Interpreter delivers the *facts*, this means there could be no difference in beliefs or desires without a physical difference. (This argument for supervenience should not be attributed to either Dennett or Davidson.)

Fodor and Lepore have, unfortunately, appropriated 'interpretivism' to label the thesis that "[t]here is an 'element of interpretation' in content ascription" (1992, p. 259). (The phrase 'element of interpretation' is taken from Dennett 1987b, p. 342.) But this is so imperspicuous as to be hardly worth naming.

After this paper was completed, I discovered that William Child (1994) discusses at length something similar to interpretivism as I formulate it (he calls it 'interpretationism'). But Child's interpretationist—who comes in a bewildering variety of guises—never quite becomes my interpretivist. The reader should consult Child's book for a different and more sympathetic way of approaching these issues.

3. Both Dennett (e.g., 1987a, p. 41) and Davidson (e.g., 1979) accept versions of Quine's indeterminacy thesis. They hold (at least) that it is sometimes indeterminate exactly what a subject believes or desires. Consider two physical doppelgangers S_1 and S_2, in physically identical environments. Given indeterminacy, might not an Interpreter permissibly interpret S_1 as believing that p, and S_2 as not believing that p? If so, it's wrong to say that the Interpreter would not interpret differently without a physical difference. However, the answer to the question is "no". What we should say instead is that the Interpreter interprets *both* S_1 and S_2 as indeterminately believing that p. Hence, even given indeterminacy, S_1 and S_2 are interpreted alike, and the argument for supervenience goes through as before. (NB: for simplicity, in what follows I shall assume that there are no cases of indeterminacy.)

2. Textual evidence

I suggested that Daniel Dennett and Donald Davidson are both interpretivists. More cautiously, there is a strong dark thread of interpretivism running through their views. Let us briefly survey the textual evidence.

As is well known, Dennett distinguishes three "stances", or predictive strategies, one might adopt towards some system—a chess-playing computer, to take Dennett's example.[4] First, we might adopt the *design stance*, and predict the computer's behaviour on the basis of what we know about its design. Second, we might adopt the *physical stance*, and base our predictions on the physical makeup of the machine. However, as Dennett observes, when it comes to predicting the moves a chess-playing computer will make, neither one of these stances will be of much use in practice—such machines are far too complicated.

There is a third predictive strategy: one may view the machine "rather like an intelligent opponent" (1971, p. 5), treating the system *as if* it had propositional attitudes. And certainly, whether or not it is a mere *façon de parler*, the language of beliefs and desires becomes almost irresistible in such cases. For instance, it might be said of the computer that it thinks it ought to get its queen out early (to use an example of Dennett's that is now something of a cliché in the philosophy of mind[5]) and no one would find such mentalistic talk at all odd. This third predictive strategy is the famous *intentional stance*.

It is somewhat obscure, to me at least, just what Dennett's claim was in the 1971 paper, "Intentional Systems". There were hints that it was merely convenient to assume that certain complex systems (including ourselves) had beliefs and desires, even though, strictly speaking, this assumption is mistaken. But in any event, Dennett's position has clarified and developed, and by 1978 we find him saying:

x believes that $p \equiv x$ can be predictively attributed the belief that p[6]

This schema is elaborated in the 1979 paper "True Believers", which, Dennett says, has replaced "Intentional Systems" "as the flagship expression of my position" (1987, p. 3). There he makes the following self-styled "perverse claim":

---

4.   In the classic paper "Intentional Systems" (1971).

5.   The example occurs in Dennett 1977.

6.   Dennett 1978, p. 67. See also Dennett 1981, p. xvii.

[a]ll there is to being a true believer is being a system whose behaviour is reliably predictable via the intentional strategy, and hence *all there is* to really and truly believing that *p* (for any proposition *p*) is being an intentional system for which *p* occurs as a belief in the best (most predictive) interpretation (p. 29).

Both this "perverse claim" and the earlier one appear to be clear statements of interpretivism. Dennett shows absolutely no sign of wanting to show how talk of interpretations, interpreters, the intentional stance, and so on, is a mere heuristic device, in principle eliminable from his theory. The failure of reductions in general has in any case been one of his major themes.

As to Davidson, one quotation will suffice. In "A Coherence Theory of Truth and Knowledge", he says:

What a fully informed interpreter could know about what a speaker means is all there is to learn; the same goes for what the speaker believes (1986, p. 315).

It is plain that this is intended as a statement of interpretivism, for two reasons.

First, Davidson is not making an empty claim. "[F]ully informed" is not an abbreviation for "fully informed about all a speaker's meanings and beliefs". He here has in mind his "Radical Interpreter" (Davidson 1973), of whom more anon.

Second, although Davidson's remark is consistent with the idea that some of a speaker's beliefs and meanings are not knowable by an Interpreter, because they are completely hidden, and so unlearnable, and so no part of "all there is to learn", this is not Davidson's intent. For one thing, it is implicit throughout the paper from which the quotation is taken that "[w]hat a fully informed interpreter could know" includes all a speaker's beliefs.

But now, in order to make the thesis of interpretivism more precise, we need to take a short detour.

3. Judgement-dependent concepts

In recent work, Crispin Wright (1988, 1989, 1992, 1993) and Mark Johnston (1993a, 1993b) have separately investigated *judgement-dependent concepts* (Johnston's terminology).[7] The idea is this. Take some

---

7.   Definitional epicycles aside, a concept *C* is judgement-dependent (Johnston) just in case judgements that something falls under *C* are "extension-determining" (Wright). Judgement-dependence is a species of what Johnston calls "response-dependence".

concept C. There is its extension, on the one hand, and our *best opinion* as to its extension, on the other. The two may coincide—either because we are lucky or skillful, or because their coincidence is an a priori matter. In the latter case—where it is an a priori truth that C's extension coincides with our best opinion of its extension—we say that C is judgement-dependent. Of course, we must avoid, to use Crispin Wright's phrase, trivializing "whatever it takes" specifications of "our best opinion": if our best opinion is characterised as the *true* opinion, then trivially all concepts will turn out to be judgement-dependent.

Let us have an example. There is something fishy, many would say, in the idea that the moral facts could outrun our idealised dispositions to make moral judgements. There couldn't be moral facts that we are forever unable to discover. So perhaps something like the following biconditional is a priori:

A state of affairs A is morally good iff the Impartial Observer would be disposed to judge that A is morally good.

This—understanding the Impartial Observer to be an idealised impartial human being—is a judgement-dependent account of moral goodness. If the Impartial Observer would not judge two states of affairs differently unless they differed descriptively—which seems plausible—then it follows that the moral supervenes on the descriptive.

Note three features of judgement-dependent accounts. First, they are not reductive: the concept in question—moral goodness in this case—appears on both sides of the biconditional (although a judgement-dependent account might be used as a stepping stone to a genuinely reductive account[8]). Second, a judgement-dependent account of 'x is F' is not happily formulated as an account of *what makes it true* that x is F, or of what *constitutes* the fact that x is F. We have a biconditional, for one thing, and no obvious motivation for taking one side of it to be any more basic than the other.[9] Third, judgement-dependent accounts are not necessarily relativistic. To

8. That is, a reductive account *of the concept C*. Take it as a stipulation on judgement-dependent accounts that the idealisations of ourselves mentioned in the biconditionals cannot themselves be reduced away.

9. I here draw on Appendix 3 to Johnston 1993a. Dennett himself mistakenly takes a judgement-dependent account of qualia to imply that they are "*logical constructs* out of subjects' qualia judgements" and so "akin to theorists' fictions" (1988, p. 55).

take the above example: if some other moral observer—perhaps from a society quite different from ours—disagrees with the Impartial Observer, then he is, according to the account, simply mistaken. And by understanding the Impartial Observer to be an idealisation of human beings as they *actually* are, we can say that even if our moral judgements had been different, the moral facts would have remained the same.

Interpretivism in general is the thesis that mental content is judgement-dependent: the facts about what propositional attitudes someone has are exactly captured by the (potential) judgements of some Ideal Interpreter (or Interpreters).[10]

Dennett's and Davidson's versions of interpretivism are *third-personal*: the Interpreter is someone who interprets someone *else*. But there are also *first-personal* accounts, where the Interpreter is taken to be the subject of interpretation, perhaps made ideally introspective.[11] But such accounts will not be our concern here.[12]

We now need an initial formulation of interpretivism that at least captures the spirit of Dennett's and Davidson's claims. Although the Interpreter is supposed to have knowledge of the contents of all your attitudes, for definiteness let us concentrate on belief. Let our initial formulation be the thesis that all biconditionals of the following form are a priori:

x believes that p iff if there were an appropriately informed Ideal Interpreter, she would be disposed to attribute to x the belief that p.

where Wright's strictures against a trivializing "whatever it takes" reading of the right hand side apply (so we do not, for instance, take the Ideal Interpreter by stipulation to be omniscient and infallible with respect to x's beliefs). We will fill in some gaps a little later. For now, there is one unique feature of interpretivism to note—not only does the *concept* expressed by 'believes that p' occur on the right hand side, but *facts* about belief are expressed by both sides. For in appealing to the Interpreter's *judgements*, we

---

10. As Johnston puts it: "the truth about meaning or content cannot outrun our idealized dispositions to grasp that truth" (1993a, p. 126).

11. For an example of a first-personal account, applied to intention, see Wright 1989.

12. For criticism of Wright's first-personal account (see above note), see Holton 1993; for a response to Holton, see Divers and Miller 1994.

are assuming that the Interpreter has beliefs. Interpretivism *presupposes* facts about mental content: namely facts about what an Interpreter would believe or, equivalently, facts about what *we* would believe under ideal conditions. So we have a reason specific to interpretivism (as opposed to general considerations about judgement-dependent accounts) to deny that it is a metaphysical thesis concerning what *constitutes* facts about content.

4. Interpretivism evaluated

We are now in a position critically to examine interpretivism. Our initial formulation was this:

x believes that *p* iff if there were an appropriately informed Ideal Interpreter, she would be disposed to attribute to *x* the belief that *p*.

We now cannot put off the question of the Interpreter's knowledge and powers any longer. Who is Dennett's Ideal Interpreter, and is Davidson's Interpreter significantly different?

Well, despite Dennett's attempt (1987b) to reconcile his views with Davidson's, it seems to me that they have substantial disagreements.[13] Dennett's Interpreter employs the intentional stance—and so do we, although not as well. Dennett's Interpreter employs whatever tacit methodology we

---

13. Here is as good a place as any to mention two other philosophers who have made use of the device of an Interpreter (or something like it)—Quine and David Lewis. Lewis is certainly not an interpretivist. As he puts it, "my problem of radical interpretation is not any real-life task of finding out about Karl's beliefs, desires, and meanings. I am not really asking how we could determine these facts. Rather, how do *the facts* determine these facts?" (1974, p. 110). Quine's relation to interpretivism is more complex. Quine's (1960) device of the radical translator is supposed to show that facts about meanings are not determined by facts about behavioural dispositions, and since Quine thinks that meaning-facts must supervene on behavioural dispositions, he concludes that there are no meaning-facts (see Van Cleve 1992). However, Quine need not appeal to some a priori connection between the deliverances of an idealised human translator and facts about synonymy to reach his conclusion. The radical translator could be God, and the argument would be unaffected. (And if the radical translator is God, Quine is simply asking how "*the facts* determine these facts".) But there is a taint of interpretivism in Quine's reason—namely, verificationism—for thinking that behavioural dispositions are the only candidates for the meaning-determining facts.

employ when determining what someone believes. We can think of her as an invisible onlooker, carefully observing your movements, your interactions with your environment, and so on, and following the interpretive practices—whatever they may be—that we human beings follow when deciding what someone believes. Let's call her the *Homely* Interpreter.[14] Unlike most of us, however, the Homely Interpreter never misapplies the tacit rules of interpretation, has unlimited concentration, is not perceptually defective, is never mistaken about the non-intentional facts, and so on. Anything which ordinary folk would take to be (non-intentional) evidence relevant to discovering what you believe, is part of the Homely Interpreter's database. A reminder: we can't allow the Interpreter *by stipulation* to have knowledge of all your thoughts, or beliefs, or actions (intentionally described), and the like. The Homely Interpreter is supposed to *find out* exactly what your propositional attitudes are, not be handed them on a definitional plate. But if sweating, reddening of the ears, blinking, whether you are drinking H_2O or XYZ, or whatever, are deemed relevant by the folk, then we will let the Homely Interpreter know these facts. We could think of the Homely Interpreter aided in her task by hidden high resolution cameras, shooting you from all angles; the Interpreter then views the resulting films in the comfort of her own home, aided by her compendious memory of what she saw herself.

Davidson's Ideal Interpreter is quite different, however: she is a *Radical* Interpreter, employing a methodology alien to the one we use in everyday life when determining what someone believes (Davidson 1973).[15] Such an Interpreter is one who begins the task of interpretation as might a scientifi-

---

14. Not to be confused with Blackburn's "homely interpreter", whom he contrasts with the "bleak interpreter". The former uses the "principle of humanity"—roughly, maximise intelligibility when interpreting; the latter, the "principle of charity"—roughly, maximise truth when interpreting (1984, pp. 277-81). This distinction is independent of the one being made in the text.

15. What about children learning their first language? Are they—according to Davidson—approximations to Radical Interpreters? Even if they are, it remains true that we who have learnt language do not employ the methodology of radical interpretation in our ordinary interactions with others. However, Davidson's position appears to be, not that children are Radical Interpreters, but rather that any Radical Interpreter would learn just what a child learns, given the same initial data (see Davidson 1994, p. 124). For more on language learning, see the Third Objection in section 4 below.

cally-minded visitor to an utterly strange land. The Radical Interpreter's data will be of this austere kind: system K produced the inscription 'es regnet' in circumstances of precipitation.

(Here, for three paragraphs, I interpolate a tricky question of Davidson exegesis, that I shall for the most part set aside in what follows.

The Radical Interpreter's official data includes, according to Davidson's 1973 and elsewhere, *intentional* facts—concerning what uninterpreted sentences the subject "holds true". But since Davidson's interpretivism seems to be at least in part motivated by Wittgensteinian "publicity" considerations—see the Third Objection below—he ought to hold that the Radical Interpreter's official data, at least in the final analysis, is purely *non*-intentional. For Davidson evidently thinks that it is not good enough, to answer worries about the public accessibility of meaning, to reply simply that we *directly observe* that someone believes that p, or utters a sentence that means that q. And he also evidently thinks that an *adequate* answer would be that by directly observing bodily movements, the production of certain sounds, and suchlike, we can recover just what a subject means and believes. But then Davidson ought not to rest content with "holdings true" as part of the Ideal Interpreter's initial database. Why should we find the intentional "holds true" epistemically unproblematic if belief contents and sentence meanings are disallowed as initial data? Surely these three categories should be in the same leaky boat.

But although I see no compelling reason to be this charitable, I will assume that "holdings true" are part of the Radical Interpreter's database.)

Bearing this exegetical digression in mind, let us return to the main thread. For our purposes, the important difference between the Homely and the Radical Interpreter is this: the Homely Interpreter needs to be in *perceptual* contact with the subject of interpretation (understanding this broadly enough to allow the Interpreter to use hidden cameras and similar devices). For there is little prospect of characterising, in non-intentional terms, information sufficient to enable us—*without the benefit of perception*, and using *our ordinary methods*—to work out what someone believes. That would be rather like giving a person a description of the distribution of color patches on a canvas, and expecting him to tell us whether the description is of a beautiful painting. Using ordinary methods of art appreciation, she will be at a loss: she needs to be *confronted* with the painting itself, not a description

of it.[16] And similarly in the case of interpreting others: we can only say that, *confronted* with some subject, we manage (here might follow a detailed account) to come to an opinion about what he believes.

But the Radical Interpreter, although Davidson normally situates her in the field, observing the subject with notebook in hand, need not be there. For there *is* supposed to be some complete characterisation of the Radical Interpreter's initial data, and hence we could simply take the Radical Interpreter simply to be supplied with it, never leaving her laboratory at all. Indeed, we can even take the Radical Interpreter to be entirely alone, off in some other possible world. (At any rate, I will assume that the preceding part of this paragraph is correct. As I shall go on to raise a problem for any version of interpretivism that requires the Interpreter to be in perceptual contact with the subject of interpretation, this assumption is a concession to Davidson.[17])

Before we get on to some objections, two minor and two major preliminary points. The minor points: let us take the necessary reference to the time period during which the subject has a particular belief to be implicit, and let us not fuss about the alleged influence of the subject's social milieu on his mental content (see Burge 1979).

The major points. First, although placing *some* knowledge of the subject's beliefs and desires, or knowledge of a theory of meaning for his language, in the Interpreter's initial endowment would not thereby trivialise interpretivism, I shall not consider these possibilities. Insofar as Davidson and Dennett are interpretivists, their interpretivism is of the stronger, more

---

16. Well, no doubt being confronted with a physical duplicate of the painting would be sufficient, which she might construct from a description. Similarly, the Homely Interpreter would not be impaired if she were confronted with a physical duplicate of the subject of interpretation (in a duplicate environment). So, strictly speaking, the Homely Interpreter need not be in perceptual contact with the *subject* of interpretation, just in perceptual contact with a *duplicate* of the subject of interpretation. But this qualification does not affect any of my objections, and so I will ignore it henceforth.

17. And in fact it seems that even if Davidson would accept the concession that the Radical Interpreter does not need to be in perceptual contact with the subject of interpretation, he thinks that *someone* must be. In his 1992 he gives what amounts to a transcendental deduction of the existence of next-door neighbours. He argues that someone could not "satisfy all the conditions for being interpretable—without actually being interpreted" (p. 260).

exciting kind. Second, a point that I have already touched on, and which cannot be stressed enough. *None* of my objections will assume the falsity of, or be objections to, supervenient behaviourism: the metaphysical thesis that facts about a subject's mental states supervene on facts about his behavioural dispositions. Interpretivism is the target, not behaviourism.

Here, then, are three objections to interpretivism. The first is only an objection to taking the Interpreter to be *Homely*. The other two are objections to interpretivism *simpliciter*. When we are through with these, I shall briefly examine whether a weakened version of interpretivism might fare any better.

First Objection: the curious case of the strange believers

The following examples are related to some that are familiar from secondary quality accounts of color (they are more readily dealt with in the color case, however).[18]

His argument for this truly extraordinary conclusion runs as follows. Suppose we interpret a child's use of 'table' as referring to tables. How do we do that? We need to find the salient cause of the child's utterance of 'table'. And "[i]t is the common cause of our response and the child's response" (p. 263).

> [This] kind of triangulation…while not *sufficient* to establish that a creature has a concept of a particular object or kind of object, is *necessary* if there is to be any answer at all to the question of what its concepts are concepts of. If we consider a single creature by itself, its responses, no matter how complex, cannot show that it is reacting to, or thinking about, events a certain distance away, rather than, say, on its skin.
>
> The problem is not, I should stress, one of verifying what objects or events a creature is responding to; the problem is that without a second creature responding to the first, there can be no answer to this question…So we can say…that if anyone is to speak a language, there must be another creature interacting with the speaker (p. 263).

If I understand this argument, a corollary of it is that if no one ever overhears me saying 'chair', then there can be "no answer" to the question of what I use this word to refer to. To reply that this is mistaken, because if an interpreter *had* been interacting with me, then he *would* have believed that 'chair' in my mouth refers to chairs, is obviously to undercut the entire argument.

18. See Johnston 1992, to which I am indebted. The present objection convicts our initial formulation of interpretivism of the "conditional fallacy" (Shope 1978).

First, there might be a believer who emitted rays capable of confusing any Homely Interpreter in the vicinity. The Homely Interpreter will then be ignorant of (many of) the subject's beliefs. Could we not say that the Interpreter interprets the believer as he would have been had he not produced the rays? Not obviously: the believer might live in a world where, as a matter of law, the rays are produced by certain *beliefs*. So if the believer had not produced the rays, he would have had different beliefs, if he had beliefs at all. Or perhaps one such believer believes on good evidence that he does produce these rays. And if the believer had not produced the rays, he would not have had this belief.[19]

Second, there might be an unobservable believer: perhaps invisible, perhaps of sub-microscopic size, or perhaps so large that only small parts of him can be observed. Again, the Homely Interpreter will be ignorant of (many of) the subject's beliefs. Could we not idealise the Homely Interpreter so that she has astonishing perceptual powers, sufficient to observe such believers? Perhaps. But the Homely Interpreter is a believer too, and once we allow believers to have astonishing perceptual powers, the following third problem arises.

There might be a believer, with astonishing perceptual powers, who believes on the basis of overwhelming evidence that he is not being perceived, and hence not interpreted by any Homely Interpreter. If a Homely Interpreter were to interpret him, no matter how subtly, he would no longer believe he was not being interpreted. So this believer has a belief that no Homely Interpreter can discover.

These problems are of no force against Davidson's vision of the Ideal Interpreter.[20] But they appear to raise serious difficulties for Dennett's account.

It might be suggested that these problems can be overcome by reformulating interpretivism in the style of Wright's "provisional equations", yielding:

If there were an Ideal Interpreter appropriately informed vis-à-vis x, then: x would believe that p iff the Interpreter would be disposed to attribute to x the belief that p.[21]

19. The counterpart of this character in the case of color is "killer-yellow", due to Saul Kripke in lectures.

20. Under my concessive assumptions, anyway. But cf. note 17 above.

21. See Wright 1992, p. 119, 1993. Here I am indebted to an anonymous referee.

Here the idea is that the Interpreter now only delivers a verdict on the subject's beliefs as they would be when he is being interpreted. Hence if the process of interpretation itself affects the subject's beliefs—as is the case in the last problem mentioned above—this does not matter. There are two obvious points to be made in response. First, as Wright notes[22], this kind of thesis lacks generality: nothing is said about the beliefs of a subject who is not being interpreted. Second, the reformulation does not help with the problem of the believer who induces confusion in the Interpreter. If there were an Ideal Interpreter appropriately informed vis-à-vis one of these believers then she would be disposed to judge mistakenly.

Second Objection: the problem of actuality

Suppose you are out walking alone. On a path through a forest, you glimpse a rabbit ahead, motionless in some tall grass. You think to yourself: Lo, a rabbit. You believe, at that moment, that there is a rabbit ahead. Preoccupied with your ruminations on *Word and Object*, before the end of your walk you completely forget about the rabbit. You definitely believed at the time (*t*, say) that there was a rabbit ahead. There is nothing unusual about this sort of case—it is utterly commonplace. But how do we get the Ideal Interpreter to deliver the right result?

We can suppose that the rabbit is not clearly visible to you, and although you look straight at it, you give no visible sign that you see the animal (you have seen enough rabbits to last a lifetime). To the Interpreter, you behave in just the way you would have had the rabbit slipped your gaze.

So if the Interpreter just knows of your *actual* movements, at best she will be of two minds about whether you have this belief (or perhaps two acceptable choices of Interpreter will disagree about whether you have it). This will then be a case of indeterminacy. At worst, she might be of one mind, and just not affirm that you have this belief—for perhaps you sincerely say later that your walk was entirely rabbit-free. In which case, according to interpretivism, you definitely do not have this belief.

Will going inside the head help? Not necessarily. For suppose the Interpreter knew that there was a rabbit-image on your retina at *t*, and even that a light-bulb marked 'rabbit' illuminated in your brain at that time, and such-

---

22. Wright 1992, p. 120.

like. Whether or not in fact this guarantees that you have the rabbit-belief, it is surely not an a priori matter that an Interpreter with access to such actual inner goings-on will be able to find out just what you believe (remember that an a priori reduction of beliefs to, say, inner light-bulbs, is not an option in the present context: reduction renders the Interpreter redundant).

This result—the missing rabbit-belief—is not acceptable. You definitely had that belief, if only for the reason that it might have influenced your action (although in fact it did not). Suppose you have a helpful disposition, and dislike rabbits. Had you met a rabbit hunter on the path a little after time *t*, you would have reported the rabbit sighting. The point here could apply equally to your ruminations on *Word and Object*—perhaps you come to believe during your walk that Quine is seriously mistaken, and out of embarrassment never give any indication then or later that this is what you believe.

Our initial formulation of interpretivism cannot be correct. The Interpreter must know, or be able to know, what the subject *would* have done, not just what he *actually* did.

Clearly, what *would* have happened had you met a rabbit hunter is relevant here. Could we perhaps equip our Interpreter with knowledge of a vast set of counterfactuals of the form: if so-and-so had happened, then you would have done such-and-such?

There are two immediate problems with this suggestion, apart from the obvious difficulty of specifying just what counterfactuals are supposed to count.

First, we cannot allow the counterfactuals to be couched in terms of belief and desire. It is sufficient to note that this suggestion is quite unmotivated— why should we allow the Interpreter to freely help herself to beliefs you *would* have had, but not to beliefs you *do* have? It is often *harder* to divine a subject's *counterfactual* beliefs than to find out what he *actually* believes.

Second, irrespective of the vocabulary in which the counterfactuals are supposed to be couched, giving the Interpreter such knowledge is to firmly distance her from all examples of the process of interpretation that are given in the literature. Note that problem of actuality arises because knowledge of how the subject *actually* behaves—both inside and outside the head—will not deliver the right results. Therefore we have to give the Interpreter knowledge of counterfactuals that *cannot* be known on the basis of how the subject actually behaves. But the ordinary person adopting the intentional

stance, or the anthropologist approaching a totally alien tribe, are simply not in possession of enormous quantities of counterfactual knowledge that cannot be gleaned from actual observation of the subject.

Now I have been unable to find a clear argument for interpretivism in the writings of either Dennett or Davidson. Their many insights concerning the mind can all be, and should all be, sharply divorced from this epistemological thesis. Be this as it may, *if* the literature contains any arguments for interpretivism, it seems fair to suppose that they depart either from the premise that the Interpreter is an idealisation of an ordinary person adopting the intentional stance, or from the premise that the Interpreter is an idealised anthropologist. But the problem of actuality shows that interpretivism, if the Interpreter wears either of these hats, is false. Therefore *if* there are any arguments for interpretivism in the literature, they are mistaken.

Third Objection: the problem of language

Language-using believers do not all use the same language, of course, and this poses an immediate problem for Dennett. For any Interpreter has got to understand the language of those she interprets, and directly building this into the Interpreter's powers amounts to smuggling in facts about the subject's beliefs. But it is quite unclear how else we are to give the Homely Interpreter (who, recall, interprets his subjects as you might interpret me) the resources she requires. There would appear to be little hope of arguing that some idealised version of myself, transplanted to China, *must* be capable, using *everyday* interpretive methods, of learning Chinese. Here Davidson is evidently on firmer ground, so let us examine how the Radical Interpreter might fare.

Davidson apparently has some general line of thought leading to the conclusion that the Radical Interpreter must be able to learn exactly what a subject's language means. But what this line of thought is, I find very hard to discover. The places in Davidson's writings where the methodology of radical interpretation receives most discussion (1973, 1974b) do not, as far as I can see, contain it. However, in "The Structure and Content of Truth" (1990), Davidson writes:

As Ludwig Wittgenstein, not to mention Dewey, G.H. Mead, Quine, and many others have insisted, language is intrinsically social. This does not entail that truth and meaning can be defined in terms of observable behaviour, or that it is "nothing but" observable behaviour; but it does imply that meaning is entirely

determined by observable behaviour, even readily observable behaviour. That meanings are decipherable is not a matter of luck; public availability is a constitutive aspect of language (p. 314).

The argument is a little terse.[23] It does, however, hint at an explicit argument along these lines:

Premise

(1)   As Wittgenstein showed, there can be no "private languages". That is, if S speaks some language L, it must be possible for others to come to know (a) that S speaks L, and (b) all truths of the form 'sentence s means (in L) that p'.

Premise

(2)   Our reasons for thinking we have understood a speaker are ultimately based on observation (of him, or those others who speak his language).

Premise

(3)   When we observe a speaker, what we observe is behaviour.

Therefore:

(4)   Any language can be learnt by observing some speaker of that language.

Premise

(5)   If observable behaviour did not determine meaning, then we could not learn a language by observing some speaker of it.

Therefore:

(6)   The meaning of any language is determined by the observable behaviour of some speaker of that language.

(4) gives the desired conclusion about radical interpretation, which is just as well, because the move from (4) to (6) is unwarranted. (5) is clearly false. If the condition which it states were imposed generally, most learning

---

23. It bears comparison with Dummett's arguments for anti-realism (see e.g. Dummett 1973). For some discussion, see Craig 1982 and Devitt 1991, chapter 14.

would be impossible. We have learnt much about the world that is not "determined" by what we have learnt it from.

Turn now to the argument for (4). Let us grant the first premise, for the sake of the argument (for what I take to be the correct rebuttal to—what is usually regarded as—the "private language argument", see Craig 1982). Now suppose you and I are both speakers of *L*. It would seem that I can have good reason to think that this is so, through our complex social interaction (how does it work this well if we speak different languages?). It is, perhaps, harmless to take this reason, as (2) declares, to be "ultimately based on observation". Turn now to the meaning of 'behaviour' in (3). Is this "behaviour described in folk psychological and semantic terms"? It can't be, on pain of giving the Interpreter knowledge of the subject's beliefs and desires from the start. But then (3) is highly contentious. For (3), in conjunction with (2), requires that my reasons for thinking that you speak *L* are "ultimately based" on observation of mere bodily movements plus "holdings true". That is hardly obvious. When we observe others, we do not only observe "raw behaviour" or that they hold certain symbol strings true—we also observe them *acting*, or *stating something*, and so forth. And it is these *folk psychologically/semantically described* deliverances of observation on which my reasons for thinking that you speak *L* are "ultimately based". At any rate, no considerations have been offered to make us doubt this natural way of describing matters.

Even if the first three premises are granted, (4) does not follow. (4) says that any language can be learnt by observing some speaker of that language. But the first three premises do not even imply *that language learning is possible*, let alone that all languages are learnable. The first premise, for instance, says that if *A* speaks *L*, then it must be possible for some other speaker *B* to know that *A* speaks *L*, and to know what is stated by the translation manual for *L*. Of course, to do the latter, *B* must understand *L* himself. However, there is no obvious reason why *B*, or anyone else, should have *learnt L*, or be capable of doing so. And it is clear at a glance that the second and third premises are equally unconnected with language learning. To be sure, we humans learn the languages we speak, but this seems to be an empirical fact, not one discoverable a priori.

In any case, (4) is too weak. (4) could at most support: for any language *L*, *L* must be learnable by the methodology of radical interpretation. But it does not follow from this that, given any speaker, *S*, of *L*, *L* must be learnable by

radically interpreting *S*. Yet that strong conclusion seems to be what David-son needs. For take some speaker of *L*, *S*, whose language cannot be learnt by radically interpreting *S* (perhaps because conversation is not *S*'s strong suit). The Radical Interpreter must be able to discover that *S* speaks *L*, if she is to find out what *S* believes. The reconstructed argument, even if sound, won't allow us to conclude that the Radical Interpreter can find out that *S* speaks *L*.

The line of argument suggested by Davidson's remarks, therefore, does not appear promising; the "problem of language" remains.[24,25]

5. Interpretivism weakened

I took interpretivism to be the a priori thesis that:

x believes that *p* iff if there were an appropriately informed Ideal Interpreter, she would be disposed to attribute to *x* the belief that *p*.

The Ideal Interpreter, according to Dennett, is a *Homely* Interpreter; according to Davidson, she is a *Radical* Interpreter. I raised a number of diffi-culties for both versions of interpretivism. Jointly, I think them conclusive. But it might be that there is a weaker kind of third-personal interpretivism about mental content, substantial enough to be of philosophical interest, yet immune to the previous objections. That is what I shall now examine.

---

24. A fourth problem might be called the "problem of anthropocentrism". The Ideal Interpreter, recall, is an idealisation of a *human being*. And if the reference to a human being is taken seriously, then it would seem interpretivism implies that there could not be believers who spoke a language uninterpretable by human beings. And Davidson, of course, has explicitly argued that, indeed, there could not be such believers ("On the Very Idea of a Conceptual Scheme", 1974a). How-ever, I believe that Davidson's paper contains materials which point to the very opposite conclusion (a rare vindication of deconstruction!). (Mark Johnston has expressed similar views in a seminar—see Hurley 1992.) So here we have a potential *reductio* of Davidson-style interpretivism.

25. In his 1994, a reply to Fodor and Lepore 1994, Davidson for the most part says things quite compatible with interpretivism. But there is one glaringly recalci-trant passage. Davidson denies he has "ever argued for the claim that radical interpretability is a condition of interpretability" and says that he has "never argued that every language is radically interpretable". "I do not think", he contin-ues, "that radical interpretation of natural languages *must* be possible; I have only argued that it *is* possible" (p. 122). If I understand these remarks correctly, Davidson is claiming simply that, as a matter of contingent empirical fact, human natural languages are radically interpretable.

With one exception to be consigned to a footnote, my objections to interpretivism were directed at the claim that the interpretivist's biconditionals provide *necessary* conditions for a subject to believe that p.[26] For example, I complained (the Third Objection) that there is no a priori guarantee that the Interpreter could learn the subject's language, and so no a priori guarantee that the Interpreter could find out all the subject's beliefs (in particular, those that are linguistically expressed). But that only shows that the Interpreter might fail to recognise a subject's belief, not that she might be wrong in making a positive belief attribution. So let us weaken interpretivism to provide merely *sufficient* conditions, as follows:

> x believes that p if the following conditional holds: if there were an appropriately informed Ideal Interpreter, she would be disposed to attribute to x the belief that p.

Here we are allowing that the Interpreter might miss some beliefs, in her final report of what a subject believes. But, if a belief attribution appears in the report, then (according to the weakened thesis), the subject definitely has it.

However, there is a fundamental difficulty, afflicting both Dennett and Davidson alike, which can be extracted from the problem of actuality (the Second Objection). Our attributions of belief are, for the most part, defeasible. If we observe what someone actually does, and ascribe beliefs to him on

---

But this is inconsistent, it seems to me, with some of Davidson's other writings. For example, in his 1986, he argues against scepticism roughly as follows. First, he claims that a Radical Interpreter would find the subject of interpretation as holding mostly true beliefs, by the Interpreter's lights. Then he claims that "an interpreter who is omniscient about the world" (p. 317) is one permissible choice of Radical Interpreter. Hence we (possible subjects of an Omniscient Interpreter) have mostly true beliefs by the Omniscient Interpreter's lights, which is to say mostly true beliefs *simpliciter.* I do not understand how this argument against the sceptic could get off the ground if it requires a contingent empirical premise about natural languages.

26. The exception is the last problem considered under the First Objection. That example concerned a subject who believes that he is *not* being interpreted, but who would change his mind if he were being interpreted. Supposing that the Interpreter would interpret the subject as believing that he *was* being interpreted, the initial formulation of interpretivism gives the result that the Interpreter's judgement is mistaken. Hence the Interpreter's judging that the subject believes that p is not sufficient for the subject to believe that p. But this problem can be overcome, as I discuss, by appeal to Wright's "provisional equations". I ignore this complication in the text below.

this basis, we must be prepared to revise our opinion if certain new evidence comes to light. For example, suppose that someone is looking in good light at a cup on a table, affirms that there is a cup on the table, reaches out for it when the tea is ready to pour, and so on. Clearly he believes that there is a cup on the table! But not necessarily. He may take himself to be the victim of some crafty cup-illusion, and not wish to give the impression that he is wise to the trick. Admittedly, this hypothesis is a little wild. But as it is a conceivable explanation of the subject's actual behaviour, the Interpreter must be able to rule it out, if her opinion that the subject believes that there is a cup on the table is guaranteed to be correct. If the hypothesis can be conclusively ruled out at all, the Interpreter will have to know what the subject *would* have done in various circumstances, not merely what he *actually* does. But this is the problem of actuality all over again. Therefore, either we restrict the Interpreter's knowledge to the subject's actual behaviour, in which case the Interpreter's firm opinions—if she has any—will be scant indeed, or else we squarely face the problem of actuality. But I do not see how that problem can be solved.[27]

In "A Study in Scarlet", Dr. Watson reported his reaction to an article written, as he later discovered, by Sherlock Holmes, as follows:

It struck me as a remarkable mixture of shrewdness and absurdity. The reasoning was close and intense, but the deductions appeared to me to be far fetched and exaggerated. The writer claimed by a momentary expression, a twitch of a muscle or a glance of an eye, to fathom a man's innermost thoughts. Deceit, according to

---

27. There is a final corner into which an interpretivist might retreat, following the failure of the merely sufficient conditions proposal. Take the Interpreter now simply to have an unimpeachable positive opinion that the subject of interpretation has a mind—not whether he believes that p, or desires that q, but simply whether he has beliefs and desires at all. That is:

x is a thinking being if the following conditional holds: if there were an appropriately informed Ideal Interpreter, she would be disposed to judge that x is a thinking being.

Arguably, the Interpreter would give the same judgements in the same physical circumstances, so this proposal would—if true—vindicate a supervenience claim: that the physical facts are at least metaphysically *sufficient* for the existence of minds. The proposal also has the merit of solving the problem of other minds (a virtue shared by the stronger versions of interpretivism). For I think it reasonable that an idealised version of myself would continue to *judge* that you have a mind, and therefore I can reasonably conclude, if the present version of interpretivism is correct, that you *have* a mind.

him, was an impossibility in the case of one trained to observation and analysis. His conclusions were as infallible as so many propositions of Euclid.

Interpretivism—at least in the versions we have considered here—is indeed, to borrow Watson's phrase, "far fetched and exaggerated". There is a certain irony in the fact that Dennett, and also Davidson, are prone to see Cartesianism as the root of all error in the philosophy of mind. They rightly repudiate the Cartesian picture of the mental as an inner stage to which the subject has complete and infallible access. But it seems to me that the picture still lingers on: the audience of one has been expanded, but the view of the stage remains as clear as ever. In denying that the Cartesian theatre is essentially private, interpretivism simply opens it up to the public.[28]

The present version is closely connected with the well-known Turing test. In "Computing Machinery and Intelligence" (1950) Alan Turing proposed, roughly, that if some entity could *convince* a human interrogator via a teleprinter that it was thinking, then it *was* thinking. In other words, passing the "Turing test" (not Turing's phrase) is sufficient for having a mind. This characterisation *is* a little rough. In the first place, Turing proposed to "replace" the question "Can a machine think?"—which he found "too meaningless to deserve discussion" (p. 49)—with the question "Can a machine pass the Turing test?" And in the second place, the machine had to do more than convince the interrogator that it was thinking. For Turing's actual proposal was that the machine must fool the interrogator into taking it to be a *human being.* But there is no obvious reason why it should have to do *that*—why is simply convincing the interrogator that he is talking to a thinking thing not sufficient? This unnecessary restriction on what counts as passing the test led Turing to consider objections of the form: a machine can't fall in love or be telepathic, etc., but human beings can. These objections are irrelevant to the question of whether a machine could pass the Turing test conceived of as a test for *thinking*, rather than as a test for humanity. All this is a little ironic, since the Turing test has been widely criticised for being too *easy*, not too *hard.*

Now in fact I think that this etiolated form of interpretivism can be made very hard to refute, in the sense that there are no convincing counterexamples. But I suspect that its appeal is precisely the appeal of the notorious "paradigm case argument", much discussed in the heyday of linguistic philosophy. "Surely *you* are a *paradigm case* of another mind. Therefore there *are* other minds." The difficulties with this are well-known.

Department of Linguistics & Philosophy
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02139-4307
USA
abyrne@mit.edu

References

Blackburn, S. 1984 *Spreading the Word*, Oxford: Oxford University Press.

Burge, T. 1979 "Individualism and the Mental," in P. French, T. Uehling, and H. Wettstein, eds., *Midwest Studies in Philosophy* 4, Minneapolis: University of Minnesota Press, 74–121.

Child, W. 1994 *Causality, Interpretation and the Mind*, Oxford: Oxford University Press.

Craig, E. 1982 "Meaning, Use, and Privacy," *Mind* 41, 541–64.

Davidson, D. 1973 "Radical Interpretation," Reprinted in Davidson 1985, 125–39.

Davidson, D. 1974a "On the Very Idea of a Conceptual Scheme," Reprinted in Davidson 1985, 183–98.

Davidson, D. 1974b "Belief and the Basis of Meaning," Reprinted in Davidson 1985, 141–54.

Davidson, D. 1979 "The Inscrutability of Reference," Reprinted in Davidson 1985, 227–41.

Davidson, D. 1985 *Inquiries into Truth and Interpretation*, Oxford: Oxford University Press.

Davidson, D. 1986 "A Coherence Theory of Truth and Knowledge," in E. Lepore, ed., *Truth and Interpretation: perspectives on the philosophy of Donald Davidson*, Oxford: Basil Blackwell, 307–19.

Davidson, D. 1990 "The Structure and Content of Truth," *Journal of Philosophy* 68, 279–328.

Davidson, D. 1992 "The Second Person," in P. French, T. Uehling, and H. Wettstein, eds., *Midwest Studies in Philosophy* 17, Minneapolis: University of Minnesota Press, 255–67.

Davidson, D. 1994 "Radical Interpretation Interpreted," in J. Tomberlin, ed., *Philosophical Perspectives* 8, Atascadero, Ca.: Ridgeview Press, 121–28.

Dennett, D. C. 1971 "Intentional Systems," Reprinted in Dennett 1981, 3–22.

Dennett, D. C. 1977 "A Cure for the Common Code?" Reprinted in Dennett 1981, 90–108.

Dennett, D. C. 1978 "Three Kinds of Intentional Psychology," Reprinted in Dennett 1987, 43–68.

Dennett, D. C. 1979 "True Believers," Reprinted in Dennett 1987, 13–35.

Dennett, D. C. 1981 *Brainstorms*, Brighton, England: The Harvester Press.

Dennett, D. C. 1987 *The Intentional Stance*, Cambridge, Mass.: MIT Press.

Dennett, D. C. 1987a "Evolution, Error, and Intentionality," in Dennett 1987. 287–321.

Dennett, D. C. 1987b "Mid–Term Examination: Compare and Contrast," in Dennett 1987.

Dennett, D. C. 1988 "Quining Qualia," in A. J. Marcel and E. Bisiach, eds. *Consciousness in Contemporary Science*, Oxford: Oxford University Press, 42–77.

Devitt, M. 1991 *Realism and Truth*, second edition, Oxford: Basil Blackwell.

Divers, J., and A. Miller 1994 "Best Opinion, Intention-Detecting and Analytic Functionalism," *Philosophical Quarterly* 44, 239–45.

Dummett, M. A. E. 1973 "The Philosophical Basis of Intuitionistic Logic," Reprinted in Dummett 1978, 215–47.

Dummett, M. A. E. 1978 *Truth and Other Enigmas*, London: Duckworth.

Fodor, J. A., and E. Lepore 1992 *Holism: a shopper's guide*, Oxford: Basil Blackwell.

Fodor, J. A., and E. Lepore 1994 "Is Radical Interpretation Possible?" in J. Tomberlin, ed., *Philosophical Perspectives* 8, Atascadero, Ca.: Ridgeview Press, 101–19.

Holton, R. 1993 "Intention Detecting," *Philosophical Quarterly* 43, 298–318.

Hurley, S. L. 1992 "Intelligibility, Imperialism, and Conceptual Scheme," in P. French, T. Uehling, and H. Wettstein, eds., *Midwest Studies in Philosophy* 17, Minneapolis: University of Minnesota Press, 89–108.

Johnston, M. 1992 "How To Speak of the Colors," *Philosophical Studies* 68, 221–63.

Johnston, M. 1993a "Objectivity Refigured: Pragmatism Without Verificationism," in J. Haldane and C. Wright, eds., *Reality, Representation, and Projection*, Oxford: Oxford University Press, 85–130.

Johnston, M. 1993b "The Missing Explanation Argument and Its Impact On Subjectivism," Ms.

Lewis, D. K. 1974 "Radical Interpretation," Reprinted in his *Philosophical Papers*, vol. 1, Oxford: Oxford University Press, 1983, 108–18.

Peacocke, C. 1983 *Sense and Content*, Oxford: Oxford University Press.

Quine, W. V. O. 1960 *Word and Object*, Cambridge, Mass.: MIT Press.

Shope, R. K. 1978 "The Conditional Fallacy in Contemporary Philosophy," *Journal of Philosophy* 75, 397–413.

Turing, A. M. 1950 "Computing Machinery and Intelligence," Reprinted in M. A. Boden, ed., *The Philosophy of Artificial Intelligence*, Oxford: Oxford University Press, 1990, 40–66.

Van Cleve, J. 1992 "Semantic Supervenience and Referential Indeterminacy," *Journal of Philosophy* 89, 344–61.

Wright, C. 1988 "Moral Values, Projection and Secondary Qualities," *Aristotelian Society Supplementary Volume* 62, 1–26.

Wright, C. 1989 "Wittgenstein's Rule-Following Considerations and the Central Project of Theoretical Linguistics," in A. George, ed., *Reflections on Chomsky*, Oxford: Basil Blackwell, 233–64.

Wright, C. 1992 *Truth and Objectivity*, Cambridge, Mass.: Harvard University Press.

Wright, C. 1993 "Realism: The Contemporary Debate—W(h)ither Now?" in J. Haldane and C. Wright, eds., *Reality, Representation, and Projection*, Oxford: Oxford University Press, 63–84.