

RADICAL PAROCHIALISM ABOUT REFERENCE

Will Gamester and J. Robert G. Williams¹

University of Leeds

Penultimate version. Final version forthcoming in *Noûs*.

We can use radically different reference-schemes to generate the same truth-conditions for the sentences of a language. In this paper, we do three things. (1) Distinguish two arguments that deploy this observation to derive different conclusions. The first argues that reference is *radically indeterminate*: there is no fact of the matter what ordinary terms refer to. This threat is taken seriously and most contemporary metasemantic theories come with resources intended to rebut it. The second argues for *radical parochialism about reference*: it's a reflection of our parochial interests, rather than the nature of the subject matter, that our theorizing about language appeals to reference rather than another relation that generates the same truth-conditions. Rebuttals of the first argument cut no ice against the second, because radical parochialism is compatible with reference being determinate. (2) Argue that radical parochialism, like radical indeterminacy, would be shocking if true. (3) Argue that the case for radical parochialism turns on the explanatory purposes of "reference"-talk: on relatively "thin" conceptions, the argument goes through, and radical parochialism is (shockingly!) true; on richer conceptions, the argument can be blocked. We conclude that non-revisionists must endorse, and justify, a relatively rich conception of the explanatory purposes of "reference"-talk.

1. A Tale of Two Arguments

1.1 *Indeterminacy vs. parochialism*

A *reference-scheme* for a language L is an account that tells us which of the expressions of L are names and what each denotes, which expressions are predicates and what each is true of, and so on. A familiar formal result is that we can use radically different reference-schemes for the subsentential expressions of L to generate the same truth-conditions for its sentences.²

¹ Authors listed alphabetically by surname.

² Classic discussions of this result are in Quine (1960), Wallace (1977), Putnam (1981). For a very general formulation, see Williams (2008).

Here is a simple example. Suppose that L is English, and let S_0 be a sensible reference-scheme that includes:

- (1) "Boris Johnson" is a name that denotes Boris Johnson (and nothing else).
- (2) "is a Tory" is a predicate that is true of Tories (and nothing else).

Taking for granted the "projective rule" (3), we can derive truth-conditions for the sentence "Boris Johnson is a Tory":

- (3) If N is a name and F is a predicate, then " NF " is true iff $\exists x$ (N denotes x and F is true of x).
- (4) "Boris Johnson is a Tory" is true iff Boris Johnson is a Tory.

Now, let φ be a one-one mapping of individuals onto individuals. For example, suppose that φ maps Boris Johnson onto my toothbrush, my toothbrush onto Boris Johnson, and everything else onto itself. We can then "apply" φ to S_0 to generate an alternative reference-scheme S_φ , which includes:

- (1*) "Boris Johnson" is a name that denotes whatever bears φ to Boris Johnson – i.e. my toothbrush – (and nothing else).
- (2*) "is a Tory" is a predicate that is true of things that bear φ to Tories (and nothing else).

Given (3), it follows:

- (4*) "Boris Johnson is a Tory" is true iff whatever bears φ to Boris Johnson bears φ to something that is a Tory.

Whatever bears φ to Boris Johnson bears φ to something that is a Tory iff Boris Johnson is a Tory, so S_0 and S_φ agree on the truth-conditions for "Boris Johnson is a Tory" despite attributing radically different referents to its constituent expressions: one says "Boris Johnson" denotes Boris Johnson; the other says "Boris Johnson" denotes my toothbrush. (One is a piece of bamboo with a mishmash of bristles on top that is incapable of rational thought,

and the other is...) We'll say that S_0 and S_ϕ are *truth-conditionally equivalent*.³ Since we can use *any* arbitrary one-one mapping to this end, it follows that, for any reference-scheme, there are radically many truth-conditionally equivalent reference-schemes.

There are two crucially different ways in which this observation can be employed. The most familiar way uses it to argue that it is radically indeterminate what "Boris Johnson" refers to. To formulate the argument, let a reference-scheme count as *co-correct* iff it is not determinately incorrect. The radical indeterminacy argument will aim to show that if a sensible interpretation on which "Boris Johnson" denotes a human being is co-correct, so is the twisted interpretation on which that name denotes a toothbrush. The schematic argument has the following premises:

- (a) A reference-scheme S is co-correct for a language L iff S satisfies constraints C_C .
- (b) If a reference-scheme S satisfies C_C , then so does any truth-conditionally equivalent reference-scheme.

For example, C_C might be the constraint that a reference-scheme maximize the truth of ideal theory, under which reading (a) is an interesting conjecture about what grounds reference, and (b) is obviously true.

Given the earlier formal result, the argument is then straightforward. The assumption is that the sensible reference-scheme is not determinately incorrect, i.e. is co-correct, so by (a) it satisfies C_C . By (b) the twisted reference-scheme also satisfies C_C . By (a) again, the twisted reference scheme is co-correct. Conclusion: it is indeterminate whether, in English, "Boris Johnson" refers to a human being, a toothbrush... or anything else you can think of.

The case for radical indeterminacy was pressed by, e.g., Quine (1960), Davidson (1977, 1979), and Putnam (1980). Consequently, the threat is taken seriously and most metasemantic theories these days come with resources intended to rebut it. The natural response to such an argument—canonically expressed in Lewis's (1984) response to Putnam—is to throw extra constraints into C_C until (b) is false. For instance, one might argue that the different extensions we can assign to a predicate are more or less "natural", and that a co-correct reference-scheme is one that does the best job of maximising the truth of ideal theory while also maximising the

³ Strictly, reference-schemes only generate truth-conditions in conjunction with "projective rules" like (3), so it is pairs of reference-schemes and projective rules – what Field (1975: 377) calls *Tarskian descriptions* – that are truth-conditionally equivalent. For simplicity we'll focus on reference-schemes, taking the projective rules as read – but truth-conditionally equivalent Tarskian descriptions can be generated by varying projective rules too. Projective rules are not safe from the arguments below.

naturalness of the extensions it assigns to our predicates. Another familiar metasemantic strategy is to argue for a causal constraint on denotation. In general, enemies of radical indeterminacy will insist that *something* more needs to go into C_C , even while disagreeing about what exactly that extra something is.⁴

A rather different argument takes the following form:

- (c) A reference-scheme S is adequate for a language L iff S satisfies constraints C_A .
- (d) If a reference-scheme S satisfies C_A , then so does any truth-conditionally equivalent reference-scheme.

where a reference-scheme is *adequate* iff it can be used to achieve the explanatory purposes that the relevant terminology – “refers to”, “denotes”, “true of”, etc. – is employed to serve.⁵ What follows from this argument is that there are radically many adequate reference-schemes – perhaps we happen to have alighted on a particular one of these schemes in our own explanatory practices, but this is a parochial fact about us. Call this *radical parochialism* about reference. In our view, the paradigm parochialist argument is given in Field (1975).⁶

The difference between (a) and (c) is that one gives conditions of co-correctness (lack of determinate incorrectness) and the other gives conditions of adequacy. A candidate reference-scheme that pairs an expression with something that it does not refer to will be incorrect, whatever its other virtues or vices.⁷ But there’s no general reason why a reference-scheme that pairs an expression with something that expression does not refer to will be inadequate, since adequacy-conditions aim to pin down whether a given reference-scheme *suberves the explanatory goals of reference*, not whether the reference-scheme *picks out the reference relation*. That is, there’s no *prima facie* reason why we should accept the premise:

- (e) An adequate reference-scheme for a language L is a co-correct reference-scheme.

⁴ Resources intended to rebut radical indeterminacy include causal constraints (Devitt, 1983), eligibility/naturalness (Lewis, 1983, 1984), structuralness (Sider, 2011), substantive rationality (Weatherson, 2013; Williams, 2020).

⁵ Cf. Field (1975: 384-5).

⁶ Note however that Azzouni (2000: 144) interprets Field as pushing an indeterminist argument. We think that is clearly the wrong reading, but the most important thing for us right now is the distinction between the two arguments, not who historically has advocated them.

⁷ Note that this is not sufficient to rebut the threat of radical indeterminacy for the language in which the terms “refers”, “denotes”, “true of” appear – it could be that these terms are themselves radically indeterminate in reference.

Premise (e) says that the relations we in fact pick out with our terms “refers”, “denotes”, “true of”, etc. are the only relations that are sufficient to satisfy the explanatory purposes for which those terms are employed. Why believe that?

The possibility that (e) is false means that radical parochialism can come apart from radical indeterminacy. In general, the parochialist can concede whatever further constraints the enemy of radical indeterminacy wants to throw into C_C : causal constraints, reference magnetism, or what-have-you. She can happily agree that there is a determinate fact of the matter about what the expressions of the language refer to, *in our parochial sense of the term “refers”*.⁸ But she will maintain that the *adequacy* conditions on a reference-scheme are significantly weaker. While *we* use the term “refers” to pick out one relation, which holds determinately between “Boris Johnson” and Boris Johnson, *we could* have used it to pick out another relation, which holds determinately between “Boris Johnson” and my toothbrush – and there would have been nothing to choose between these two different ways of talking, as far as the explanatory purposes of “reference”-talk are concerned.

For instance, let’s call the relation that is picked out by our sensible scheme of reference S_0 , reference_0 , and the relation that is picked out by our permuted scheme of reference S_ϕ , reference_ϕ . “Boris Johnson” refers_0 to Boris Johnson, while it refers_ϕ to my toothbrush. The radical parochialist accepts that the relation that we are *in fact* thinking and talking about when we think and talk about reference is reference_0 (i.e. $\text{reference} = \text{reference}_0$), and happily endorses a particular C_C (appealing to causal constraints, reference-magnetism, etc.) to back up this position. However, she insists that we could employ reference_ϕ instead of reference in our semantic theories, and nothing of explanatory interest would be lost. Any interest we have in *reference* itself is parochial, she therefore insists. It would be an interest in the relation that *we use* to do a job, not an interest in how the job *must be done*.

1.2 Feature or bug?

The paradigm radical parochialist, Hartry Field (1975: 393-394), argues that not only is radical parochialism true, but also mundane.⁹ In this subsection our target is to say why reasonable people should think, pace Field, that it would be a shocking conclusion. The following section shows how to avoid the shock.

⁸ This will include the concession that there is a determinate fact of the matter about what “refers” refers to. The parochialist’s position is not to be confused with a kind of semantically-ascended or meta-language indeterminacy.

⁹ Field calls radical parochialism “conventionalism”, but that’s a loaded term that we prefer to avoid.

Field offers an analogy with physics. He notes that we could have used the word “force” to pick out (what-we-currently-call) force multiplied by kinetic energy. According to Field:

“the novel “force-concept” could be used in developing physics just as well as the ordinary force-concept can: for in place of the law “Force = mass \times acceleration”, we could use the law “Force = mass \times acceleration \times kinetic energy”; in place of “Force = GMm/r^2 ”, we could use “Force = $GMm \times$ kinetic energy / r^2 ”; and so forth.”

Field’s idea is that it’s a parochial fact about us that we use “force” to pick out what it does. He concludes that parochialism is an uninteresting “feature of *any* theoretical concept” (1975: 393).

Field’s physics example (assuming that it could be fleshed out in full detail) illustrates a *moderate* kind of parochialism. *Force* is a real explanatory physical quantity, as is *kinetic energy*. It’s not at all a stretch to admit that the defined quantity *force times kinetic energy* is also a real, explanatory, physical quantity. Often, in developing good theories (in physics or elsewhere), there are different reasonable choices of sets of primitives, such that starting with one set of primitives one can “define up” the other, and vice versa. A clean case is: mereology axiomatized with “part” vs. mereology axiomatized with “overlap” (Simons, 1987).

The moderate parochialist insists that there isn’t anything interesting or deep at stake when we frame our theory using one set of primitives over another. In this, the moderate parochialist will disagree with, for example, Ted Sider (2011), who has been advancing an approach to metaphysics that insists that the world itself somehow “chooses between” interdefinable sets of properties such as parthood and overlap, or those involving force vs. force \times kinetic energy. Sider claims that theories are objectively better if framed with primitives that correspond to the true structural elements of fundamental reality. But Sider acknowledges that this commitment of his approach is surprising; moderate parochialism is the commonsensical approach. It wouldn’t be a shock if it were true.

It doesn’t follow from this sort of moderate parochialism that any way whatsoever of reframing physics in terms of other properties is similarly parochial. To get a more radical example, closer to the radical parochialism about reference being contemplated, consider “shadow-physics”, framed in terms of shadows of real physical properties like mass (Shoemaker, 1988: 211ff). For example, a space-time region has shadow-mass of k iff it is a region which, exactly 5 minutes before (in some reference frame), was occupied by something

of mass k . Similarly we may define “shadow-force”, “shadow-kinetic-energy”, etc. These shadows of real properties of regions of space-time are not themselves real, explanatory, physical properties (Goodman’s grue, after all, is not a real, explanatory, colour property, despite being definable out of such). Nevertheless, as Shoemaker observed, shadow-mass and shadow-force satisfy regularities just like the force-mass laws of standard physics. Modulo the definitions, one could use shadow-physics to predict the trajectories of real things. Nonetheless, we submit that it is not a parochial fact about us that we think and theorize in terms of force and mass rather than shadow-force and shadow-mass.

This distinction between moderate parochialism about physics (very plausibly true, though not uncontroversial) and radical parochialism about physics (surprising-if-true, but very plausibly false) parallels a similar pair of positions over indeterminacy about reference. Plausibly, most words in our language are *somewhat* indeterminate in denotation. For many the “problem of the many” will be a convincing example (Unger, 1980; Lewis, 1993; McGee, 2005). There are many Prime-Minister-shaped fusions of matter in the vicinity of Boris Johnson, give or take the odd semi-detached hair, and between the various local Boris-candidates, there’s no fact of the matter which is picked out by the term “Boris Johnson”. But such moderate, local indeterminacy does nothing to make more plausible the wildly radical indeterminacy of reference according to which it’s not a factual matter that my *toothbrush* isn’t the thing that “Boris Johnson” denotes.

Field’s moderately-twisted-physics example shouldn’t convince you that *radical* parochialism is a mundane feature of theories *in general*. Radical parochialism about physics would be a shocking and revisionary thesis. That said, we can certainly imagine situations in which a revisionary thesis of radical parochialism is correct. For example, suppose you had two communities of people, neither of which explicitly relativize their “length” talk to units. One community, the imperialists, say “the length of the table is N ” where we would say “the length of the table is N foot”. The other community, the metricians, say “the length of the table is M ” where we would say “the length of the table is M metres”. The conventional aspect of our talk about lengths is clearly marked in our language; not so for the imperialists or metricians. Now, the imperialists’ use of “the length of” picks out a certain function from objects to real numbers; the metricians’ use of the same-sounding phrase picks out a different function. Despite this obvious difference, there’s a clear sense in which they, and we, are *talking about a common subject matter*: namely, certain physical quantities. In and of itself, this subject matter does not privilege either way of talking. Any dispute between an imperialist and metrician about the “real” length of the table – “6.6!” avows the imperialist, “2!” insists

the metrician, table-thumps and teeth-gnashing abound—is a paradigmatically verbal dispute. These are simply different ways of representing exactly the same state of affairs. Each disputant may have pragmatic reasons for preferring their own way of talking—they find it easier, they’ve always done it this way—but neither way of talking latches onto an objective feature of the subject matter that the other misses out on. When it comes to the imperialists’ and the metricians’ “length”-talk, radical parochialism is exactly the right view. Each is adequate to the underlying explanatory purposes. The choice to use one rather than the other can only be made on pragmatic grounds or else arbitrarily.

We could tell the same story about reference. Suppose there is an isolated community, the Permutes, who speak a language otherwise very much like English, but whose “reference”-talk accords with our permuted reference-scheme, S_ϕ . Their use of “reference” picks out reference_ϕ . When a Permute says that “Boris Johnson” “refers” to my toothbrush, they mean it refers_ϕ to my toothbrush—i.e., what they say is true, in their own language. Our disagreement with the Permutes about what the name “refers” to is a merely verbal dispute.

Does the reader nonetheless feel the urge at this point to thump the table and insist that the Permutes’ way of theorizing about language fails to latch onto an objectively important feature of the subject matter, just as talk about shadow-mass fails to latch on to the objectively important features of physics? We do. Our talk about Boris Johnson bears an objectively important relation to *Boris Johnson*—*that* guy!, the one in Downing Street!—that it just doesn’t bear to my toothbrush, or anything else in the universe. If the radical parochialist is right, then our insistence on this point is like an imperialist’s insistence that the length of a table bears an objectively important relation to the number 6.6: our way of talking and the Permutes’ way of talking are simply different ways of representing exactly the same state of affairs. Surely they must be wrong.

But incredulous stares are cheap and ineffective in philosophy. After all, we have an *argument* for radical parochialism that we need to grapple with. Either that argument is somehow unsound, or its conclusion is true. If it’s unsound, we owe an account of where it goes wrong. In this section, we’ve shown why reasonable people might regard its conclusion as shocking. The question for the rest of the paper is: what do we have to commit to, to avoid the shock?

2. Parochialist Premises

In section 1, we argued that, even granting that the revisionary challenge posed by radical indeterminacy about reference can be met, radical parochialism about reference poses a further, distinct challenge. This section evaluates what it would take to resist this revisionary conclusion. Recall that the argument for radical parochialism takes the following schematic form:

- (c) A reference-scheme S is adequate for a language L iff S satisfies constraints C_A .
- (d) If a reference-scheme S satisfies C_A , then so does any truth-conditionally equivalent reference-scheme.

A reference-scheme is *adequate* iff it can be used to achieve the explanatory purposes that the relevant terminology — “refers to”, “denotes”, “true of”, etc. — is employed to serve. We thus suggest that the parochialist’s premise (c) is to be backed by premises of the following schematic form:

- (i) The explanatory purposes for which the terms “refers” (“denotes”, “true of”, etc.) are employed are P .
- (ii) A reference-scheme S satisfies P iff S satisfies constraints C_A .

This, in turn, gives us two ways of responding to the parochialist argument. A *purpose-internal* response concedes that P are the explanatory purposes for which we use “reference”-talk but denies that satisfying C_A is sufficient for satisfying those purposes. That’s denying (ii). A *purpose-external* response can concede that satisfying C_A *would be* sufficient to satisfy P , but denies that P exhaust the relevant explanatory purposes. That’s denying (i). To be fully satisfying, the purpose-external response will need to put forward a rival conception of the explanatory purposes of “reference”-talk, and argue that it is not the case that radically many reference-schemes satisfy *those* purposes. Where responding to radical indeterminacy requires getting a better fix on *what* we talk about when we talk about reference, responding to radical parochialism requires getting a better fix on *why* we talk about what we talk about when we talk about reference.

We’ve presented the radical parochialist challenge so far by talking about the reference relation for a *language*, L . But from now on we’ll be focusing on it as it applies to reference-in-*thought*. If you endorse a language-of-thought thesis, it should be no problem at all to see what’s going on here: simply take L to be mentalese (cf. Field, 1978). For presentational

purposes, we'll adopt this reading, but in fact, the parochialist challenge also arises for mental representation under the assumption that mental *content* has language-like structure, whether this is by way of structured Fregean senses or Russellian propositions. We show this in Appendix A. We consider the argument for radical parochialism about linguistic reference in Appendix B.

The parochialist argument runs on an assumption about the explanatory purposes for which (mental) reference is employed. Below, we examine three hypotheses in turn, namely that these purposes are:

1. Predicting the behaviour of others;
2. Predicting and explaining the behaviour of others;
3. Articulating the reasons for which others act.

We will show that if (1) is right, then the argument for radical parochialism succeeds; and if (3) is right, then the argument for radical parochialism fails. If (2) is right, then whether or not the parochialist argument goes through turns on what explanation is.

However, there is another possibility we want to flag from the start. (1)-(3) agree on at least one thing: our interest in reference-in-thought is ultimately instrumental to its relation to something *beyond* mental representation. An important rival hypothesis is that our purpose in theorizing about mental reference is simply an intrinsic, non-instrumental interest in mental representation in general, and mental reference in particular. Premise (i) in the parochialist's argument would then have to read: "The purpose for which "reference"-talk is employed is to express facts about what refers to what." This is a show-stopping move. Once explaining *reference* is included in the specification of the purposes, then *ipso facto* a reference-scheme will only satisfy the purposes if it gets the facts about reference right.

There are cases where an analogue of this move is appropriate. We earlier invited the reader to agree that physics (framed in terms of fundamental properties like mass and force) can't be replaced without loss by shadow-physics (framed in terms of shadow-mass and shadow-force). But why not? On a common view, we ask physics to do two things for us. First, we look to physics to tell us about the most fundamental structure and workings of the world we live in. Second, we look to physics to put us in a position to build better bridges and bombs, enabling us to achieve our practical ends. Suppress the first while retaining the second, and shadow-physics plays the residual physics-role very nicely. Underlying the commonsense judgement that radical parochialism about physics is wild, we submit, is intrinsic interest in

describing the most fundamental structure of the world. A sceptic might ask: why care about fundamental structure, rather than shadow-fundamental structure (where shadow-mass is shadow-fundamental iff mass is fundamental)? Here we are perfectly comfortable making the show-stopping reply: the fundamental structure of the world is of intrinsic interest.¹⁰ (By the same token, someone who really wants to *defend* radical parochialism about physics would be well-advised to mount an argument that it's improper to care intrinsically about fundamental structure, and that physical theorizing is purely instrumental to purposes like building better bridges – ultimately perhaps to securing food, shelter and physical security. To us, this seems just as implausible as radical parochialism itself, but these theses do hang together.)

Every theorist will have purposes for theorizing that are intrinsic rather than instrumental on pain of infinite regress. Show-stopping moves like this are just what happens when you hit this bedrock. The authors of this paper declare an intrinsic interest in the following general questions: What is the fundamental structure of the world?, What predicts what?, What explains what?, and: What is a reason for doing what? Ask us: *in constructing theories, why are you targeting explanation rather than φ -explanation?* and we answer, flat-footedly: *figuring out what explains what is part of why we're in the theory-business in the first place.*

But equally there are cases where the analogous show-stopping move is inappropriate. An imperialist might try to make the show-stopping move by declaring an interest in imperialist-length *as such*: the purpose of imperialist "length"-talk is to express that specific function from objects to real numbers. Likewise, a Nazgûl might try to make the show-stopping move by declaring an intrinsic interest in shadow-mass and shadow-force. But *those* purposes are perverse, silly, and fetishistic. The *proper* purpose of imperialist "length"-talk is the same as ours and the metricians': to capture facts about comparative length. The *proper* purpose of "mass" and "force"-talk is to describe the fundamental structure of the world.

Any theorist can declare what purposes they like, of course, but they cannot escape evaluation of purposes. Lurking behind any way of filling in the parochialist's schema is a normative question: what are the good, proper, appropriate purposes for this piece of discourse? And what putative purposes are silly or fetishistic? Our goal here is not to enter into this normative dispute, but to see what we need to take the (proper) purpose of "reference"-talk to be, if we are to avoid radical parochialism. One option, we've noted, is to

¹⁰ Intrinsic interest in the fundamental structure of the world in general gives a swift answer to parochialist challenges against theories framed via properties that are fundamental. So if one thought that not only mass and force (and consciousness? and persons?) were fundamental but also thought that (mental) reference was fundamental, the same swift response to reference-parochialism would be available. Add therefore to the list of presuppositions of the second half of this paper the non-fundamentality of reference.

add reference to our list of things of intrinsic interest. We do not rule this out, but we are unsure whether such intrinsic interest is ultimately defensible, or merely fetishistic. Our goal in what follows is thus to determine whether those who do *not* take the reference relation to be of intrinsic interest can nonetheless avoid the threat of radical parochialism.

In doing so, we will be relying on our assumptions about what *is* of intrinsic interest: that prediction, explanation, and reasons are bedrock means that we won't be launched on a regress if we can explain by appeal to them an instrumental need to theorize using reference rather than reference_φ. The sections below explore these hypotheses in turn. The reader might wonder: what settles whether putative purposes are legitimately taken to be matters of intrinsic interest? What arguments can we offer for *why* we carve issues into bedrock and non-bedrock as we do? These are heady meta-questions, perfectly proper matters for philosophical inquiry, but not on our agenda in this paper. We have to start somewhere. As the reader will shortly see, we have enough to do exploring the terrain defined by our (reasonable!) starting point.

2.1 *Predicting behaviour*

Perhaps *predicting* the behaviour of others is what ascription of contentful mental states does for us. The person passing me on the street could walk quietly by, or they could stick out a leg and trip me up. Insofar as I need to predict what they'll do, my predictions about their behaviour will often need to be based on assumptions about the way they want the world to be, and the changes they expect to result from actions open to them, under a presumption they are rational enough to do the things most likely to bring about their desired ends. *Theorizing* about beliefs and desires, on this telling, captures those aspects of the minds of others that rationalize their behaviour, and so allows us to predict the behaviour of the (rational-enough) creatures around us.

If this is what we're up to (if it *exhausts* what we're ultimately up to) then we can make the case for radical parochialism. Consider the kind of thing we're trying to predict, for example, an agent tripping up a passer-by. This event can be redescribed in arbitrarily many necessarily equivalent ways. The event of *Sally tripping up Sue* is equally *Sally tripping* up Sue**, where *tripping* x* is *tripping someone mapped to x under a permutation*, and *Sue** is the individual

mapped to *Sue* under that permutation. These are not distinct events, we say.¹¹ Whether your theory of mind ascribes Sally the intention to *trip up Sue* or instead to *trip* up Sue**, one and the same event will be brought about if the intention is realized. So all that's necessary for predictive purposes is to get the coarse-grained (up-to-necessary-equivalence) content of the intention correct. In order to fill in the rest of the rationalizing story, you'll want the coarse-grained content of beliefs and desires to stand in the right relation to that content (e.g. for that content to be the one that maximizes expected utility, by the lights of what your interpretation says about the agent's beliefs and desires).

(Note that, in saying that truth-conditionally equivalent reference-schemes are equally adequate to the purpose of predicting behaviour, we are not saying that *we* have no reason to prefer using one reference-scheme to another for this purpose. Perhaps we find some schemes easier to learn or use. Ease of use certainly recommends that we *carry on* using a reference-scheme that conforms with our previous use of "refers".¹² But such practical matters reflect features of *us* as limited, flesh-and-blood theorists, and are orthogonal to questions of what's predicted.)

This supports (ii-P) in the following version of the earlier schematic argument:

- (i-P) The purpose for which "reference"-talk is employed is to ascribe contentful mental states to others in a way that predicts their behaviour.
- (ii-P) A reference-scheme *S* satisfies these purposes if *S* generates a reasonable assignment of truth-conditions to individual belief/desire/intention states.¹³

If (i-P) is also true, then radical parochialism is (shockingly!) true.

2.2 *Explaining behaviour*

If what mental reference is for is predicting an event *E* – a bit of behaviour, say – all that's required is that you have a theory that, on appropriate inputs, spits out *E* as output. Adjust focus and ask whether a given theory *explains* the events it predicts, however, and we

¹¹ For those who disagree, and insist the metaphysics of events is very fine-grained, the point is that the distinctions between necessarily equivalent events are "don't cares" relative to what we want to predict. Trip up Sue or trip* up Sue*, Sue still ends up on the ground.

¹² Cf. Field (1975: 377-380).

¹³ The reductively-inclined could add that *S* needs to be reducible to non-semantic facts. But if *S* is reducible, then so are its truth-conditional equivalents (cf. Field, 1975: 381-384), so we can suppress this constraint.

have resources that may allow the parochialist argument to be resisted – if we take on some contentious commitments about explanation.

Suppose that one thinks of the explanatoriness of a theory as an objective (not just pragmatic) matter. For example, one might think that some theories are objectively simpler than others, and that simpler theories are more explanatory of the phenomena that follow. To fix ideas: an infinitely long list of every microphysical fact in the universe at the present moment (supplemented with a clause that tells us that these are *all* the microphysical facts) *predicts* every present microphysical fact, given a bit of deductive logic, but surely it's not explanatory. But a neat formulation of compact physical laws, together with a specification of the initial microphysical conditions of the universe, also allows derivation of these microphysical facts, and this time it does seem explanatory. On the view being considered, this is traced to differences in objective simplicity, perhaps measured in this case by complexity of formulation in an appropriate microphysical language.

Analogously, an objectively *simpler* reference-scheme explains behaviour *better* than a more complex one, even if the rival reference-schemes ultimately predict the *same* behaviour. And so, by including explanation in clause (i) of the parochialist argument, one can make a case, based on the objectivist conception of explanation, that point (ii) can't just require generating the right truth-conditions, but needs also an additional clause saying that the truth-conditions are generated in an optimally simple way. But now, to justify the premises they need, the parochialist will have to argue that the deviant *refers_φ* reference-scheme is as simple as the original *refers* reference-scheme on which it is based. The anti-parochialist says: they will not be able to do this, and the argument is no longer a threat.¹⁴

In rebuttal one might say that simplicity in particular, and explanatoriness in general, is not an objective matter, but a reflection of pragmatic relations between the theorist and theory. Simplicity might, for example, be compactness in whatever language the theorist

¹⁴ To be clear: we are saying the parochialist has no *justification* for the premise they need. We have not provided positive independent reasons for saying that the *refers* reference-scheme is the simplest candidate available (though this is a familiar and fairly plausible claim). But this is enough! The burden is on those who put forward arguments for revisionary conclusions to justify their premises—once such justification is undercut, we non-revisionists will happily *tollens* the argument to extract justification for endorsing the negation of the contested premise. This dialectical situation leaves non-revisionists with a lingering vulnerability to future discovery of new ways to reinstate the argument by uncovering a new way to justification—but such is life. For relevant discussion, see e.g. Williams (2007) on simplicity and canonical languages, as a reconstruction of Lewis's idea of the eligibility of semantic theories. A related idea, advocated by Sider (2011), is that theories whose primitive terms pick out more natural/structural aspects of the world are *ipso facto* more explanatory. While Williams (and Lewis, we think) is exploring a first-order metaphysics of reference that would appeal to objective similarity, Sider proposes a desideratum of explanatoriness for semantic theory itself, which is the kind of second-order evaluation of the theory at issue in this paper.

happens to be working with, or easiness-to-use relative to the theorist's skills and background. Simplicity, and explanation, might still matter for theory-choice, but on this picture they would not feature among the adequacy conditions for discharging the purposes of the theory, but among the pragmatic factors for choosing which to work with among the several theories which discharge those purposes.

Here's the upshot: the purpose-external shift from prediction alone to prediction-and-explanation makes room for appeal to objective constraints on explanatoriness that would rule out the deviant reference-schemes. But whether or not there are objective constraints on explanation is a substantive matter (not one we'll try to settle here). The take-home message is that prediction-and-explanation may or may not lead to parochialism, depending on the outcome of a purpose-internal debate about the nature of explanation.

2.3 *Articulating reasons*

Consider now the following instantiation of (i):

- (i-A) The purpose for which "reference"-talk is employed is to ascribe contentful mental states to others in a way that articulates the reasons for which others behave.

This final strengthening introduces *normative* considerations into the aim of theorizing about mental reference. Generating the right predictions, even generating them in a simple and explanatory way, is no longer sufficient. The job of a theorist of representation is to understand *the reasons for which* agents do what they do.¹⁵ In the sense we intend this, the aim is to articulate the normative reasons upon which the agent *bases* their intention to act. Normative reasons for an act are considerations that speak in favour of the action; for such considerations to be the reasons for which an agent acts, they have to be somehow registered by the agent. This registration often happens by way of mental representation – by the agent *believing* the consideration obtains, for example. Part of the task of ascribing beliefs to Sally,

¹⁵ That appeals to normative reasons (within a suitable interpretationist theory) solve *indeterminacy* puzzles is the key element of Weatherson's (2013) reconstruction of Lewis, which forms the basis for Williams' (2020) book length development of a mental metasemantics. But (i-A) is independent of such first-order proposals. On the one hand, Weatherson's Lewis and Williams can perfectly consistently reject (i-A) in favour of the earlier candidates for the purpose of "reference"-talk. On the other hand, causal theorists of reference who would look askance at introducing normative facts into the reductive base for reference, such as Devitt and early Field, can consistently endorse (i-A).

therefore, is to enable us to say which of the many possible reasons for her eating a sandwich is *her reason* for eating that sandwich.

Assume that *the sandwich being delicious* is a reason for Sally to eat the sandwich. A sensible interpretation of Sally that has her eating the sandwich *because she believes that the sandwich is delicious* has her basing her action on that reason, in the sense we intend this. A truth-conditionally equivalent interpretation, by contrast, has Sally eating the sandwich because *she believes that the object that bears ϕ to the sandwich has the property of bearing ϕ to something that is delicious* (that the ϕ -sandwich is ϕ -delicious, for short). But that some arbitrary object has some such gruesome property is not a consideration that speaks in favour of eating a sandwich, or so commonsense tells us.

What backs this commonsense verdict? Lying behind it is the fact that normative reasons are hyperintensional (more precisely: ' $___$ is a normative reason to ψ ' is a hyperintensional context): p may be a good reason for taking an action, while q is not, even if p and q are necessarily equivalent. The Euthyphro contrast is a case in point. Consider the view that DePaul (1987: 436-437) attributes to William Paley: we are morally obliged to do what God commands; and what God commands us to do is whatever is most conducive to happiness (what is "felicific"). As DePaul emphasises, by Paley's lights both *being felicific* and *being commanded by God* are necessarily co-extensive with *being morally obligatory*, but only the latter is what *makes* an action morally obligatory. In our terms: only the latter gives (moral) *reason* for, say, donating to charity. Of course, a theist who is not a divine command theorist might take the other side of the Euthyphro choice to Paley: they might maintain that, while God's commands necessarily track what is felicific, it is the fact that donating to charity is felicific, rather than that it is commanded by God, that is the moral reason to do it. Both sides would agree we have normative reason to donate to charity, while disagreeing with each other as to *why*. Such disagreement is only possible in a hyperintensional context. Here are three more examples of this:

[A] Quantified reasons: fairness involves treating *everyone* equally. Treating equally only those who meet some artificial, unnatural condition F , is not fairness (call it fairness*). Fairness is a consideration that speaks in favour of some actions; fairness* is not. It turns out that there are choices of artificial, unnatural conditions F such that any

proposition involving unrestricted quantification is truth-conditionally equivalent to one where all quantifiers are restricted to F .¹⁶

[B] *De re* reasons: the proposition *that x is y 's child* may be a special weighty reason for y to save x over others from drowning. It is the relational fact that features the individuals x and y and their special relationship that matters normatively. But again there are truths necessarily equivalent to *that x is y 's child* that concern other individuals entirely.¹⁷

[C] Disjunctive non-reasons: *that the sandwich is delicious* is a reason to eat the sandwich; *that the sandwich is either delicious, or delicious and poisonous* is not a reason to eat the sandwich. But, necessarily, the sandwich is delicious iff it is either delicious or both delicious and poisonous.¹⁸

The hyperintensionality of the reasons-relation makes space for an attractive metaphysics of reasons: normative reasons are facts built only out of *relevant* bits of the world – the people, properties, or domain of individuals that normatively matter. It also makes space for an attractive general metaphysical analysis of the reason-relation (cf. Faroldi & Protopopescu, 2019: 415). Consider the thesis that for p to be the total normative reason for Sally to ψ is for p to fully metaphysically ground the deontic fact that Sally ought to ψ (compare Broome (2013: 50) on *pro toto* reasons). But metaphysical grounding relations are hyperintensional: *that Socrates drank hemlock* and *that Socrates' singleton is the singleton of a philosopher who drank hemlock* are two facts that are necessarily equivalent, but the first, and not the second, is the metaphysical ground of the fact that *some philosopher drank hemlock*. Since the grounding relation is hyperintensional in both its slots, the relevant normative reason relation will be too.¹⁹

¹⁶ Skolem (1920/1967), Putnam (1980). These arguments normally focus on truth-values, rather than truth-conditions, but by applying the arguments (at least to a non-modal language) relative to each world, one can extract the stronger version alluded to in the text. This construction is discussed in Williams (2020: 49-52).

¹⁷ We can construct these, indeed, from the permutation schemes above (McGee, 2005; Williams, 2008; Williamson, 2013).

¹⁸ Faroldi and Protopopescu (2019: 412-413) use a similar observation to argue for a closely-related, but distinct, conclusion; namely, that ' p is a normative reason to ____' is a hyperintensional context.

¹⁹ As stated, this is an argument that the *pro toto* reason relation is hyperintensional. But what we need in this paper is that the *pro tanto* reason relation is hyperintensional. However: (i) in the work cited, Broome also endorses a (more complex) analysis of *pro tanto* reasons as elements in noncausal weighing explanations, and parallel considerations for hyperintensionality will be extractable from that context, we anticipate; (ii) the *pro toto* reason for N to F will be composed out all the *pro tanto* reasons for and against N Fing, and so we cannot see how the former can be hyperintensional and the latter not. So we think this supports the thesis we need.

The authors endorse the commonsense verdicts [A-C] and the picture of reasons as built out of normatively relevant bits of the world. They are neutral on the analysis of the reasons-relation just floated, though they note parochialists seeking to overturn commonsense to evade this argument cannot be similarly neutral on such questions—and so should motivate the contentious stance they are taking.²⁰

We can now spell out why the argument for parochialism fails. Given (i-A), the parochialist argument requires that the twisted interpretation of Sally is just as good as the sensible interpretation for articulating the reasons for which Sally acts. But on the twisted interpretation, Sally bases her action on the fact that the ϕ -sandwich is ϕ -delicious. So, the parochialist must maintain that the fact that the ϕ -sandwich is ϕ -delicious is a genuine reason for Sally to eat the sandwich. More generally, the parochialist must say the same for every other reason Sally has for each act she performs, and for every permutation, in doing so committing themselves to a superabundance of necessarily-equivalent reasons. We see only one principled, non-ad-hoc motivation for opening the door to this superabundance of necessarily-equivalent reasons: the general principle that if p is a normative reason for ψ , and q is necessarily equivalent to p , then q is a normative reason to ψ . Because the reason-relation is hyperintensional, the parochialist argument will be unsound.²¹

3. Conclusion

The case for radical parochialism about reference, we've argued, is determined by the explanatory *purpose* of (mental) "reference"-talk. If the point of engaging in such talk is merely to predict the behaviour of others, then radical parochialism is (shockingly!) true. On a richer conception of the purpose—e.g., as involving explanation (where explanatoriness is understood as a non-parochial matter) or reason-articulation—the argument can be blocked. Non-revisionists must therefore endorse, and justify, a relatively rich conception of the explanatory purposes of (mental) "reference"-talk.

²⁰ If the parochialist wants to get off the boat at this step, Williamson (forthcoming) contains a spirited statement of intensionalism in a relevant context.

²¹ A quick bookkeeping note. The intensionalist can if they wish identify reasons with fine-grained entities, such as structured states of affairs. They will then be committed to a superabundance of distinct reasons standing in the same reason-relation to a given action. It would be more elegant for them to identify reasons with coarse-grained states of affairs, such as true possible-worlds propositions, so that rather than a superabundance of necessarily equivalent reasons there is just a single reason, which is denoted by each of the terms <the fact that the sandwich is delicious> and <the fact that the ϕ -sandwich is ϕ -delicious>.

Whether a theorist goes for prediction, explanation, or reason-articulation—or some other option—as the *purpose* of “reference”-talk is *prima facie* independent of what that theorist says about the nature of the reference relation itself. Suppose one gives a causal account of the metaphysics of (mental) reference (Field, 1972, 1978). That theorist could, consistently with this, add that the point of employing that particular causal relation in theorizing about other minds is just to predict behaviour—leading to parochialism (Field, 1975). But they could with perfect consistency say that the point of employing that causal relation is to articulate others’ reasons for behaviour—blocking parochialism. Or again: some explore a metaphysics of (mental) reference that makes essential appeal to the normative reasons for which a person acts (Lewis, 1974; Weatherson, 2013; Williams, 2020). But that theorist could, again, add that the point of employing that particular normatively-based relation is just to predict the agent’s behaviour, and they would be parochialists (they would regard reference_φ , grounded in maximizing normative reasons $_\varphi$, as being equally adequate, since it predicts the same pattern of behaviour). Or such theorists might maintain that the point of employing the normatively-based relation is to articulate others’ reasons—blocking parochialism. The question of the *metaphysics* of reference—whether the relation is constitutively causal, normatively-based, or whatever—is independent of the question of the *purpose of theorizing about the relation so-constituted*. Avoiding revisionism in the theory of reference requires getting a fix, not only on *what* we talk about when we talk about reference, but *why* we talk about what we talk about when we talk about reference.

Appendix A

The argument for parochialism in section 1 focused on *linguistic* representation, and so applies to *mental* representation on the assumption that the vehicles of thought have a language-like structure (the “language of thought” hypothesis). This appendix sketches how to formulate this challenge if we drop the language-of-thought assumption. For illustration, suppose that structured contents of mental states are *Russellian propositions*, structured complexes of objects and relations. For present purposes, we can take an (atomic) Russellian proposition to be an ordered n -tuple composed of an $(n-1)$ -place relation and $n-1$ objects. Thus the proposition that Boris Johnson is a Tory is a 2-tuple composed of the property (1-place relation) of *being a Tory* or *Toryness*, and Boris Johnson himself, which we denote thus: $\langle \text{Toryness}, \langle \text{Boris Johnson} \rangle \rangle$.

Given an atomic Russellian proposition p_0 , we can use an arbitrary one-one mapping φ to generate a different proposition p_φ that is *truth-conditionally equivalent* to p_0 . To do so, we replace each element of p_0 with a permuted variant given to us by φ . For an object x , this is

$\varphi(x)$; for a property, F_{ness} , this is the property of *bearing φ to something that is F* (call this, φ - F_{ness}).²² So, if p_0 is $\langle \text{Toryness}, \langle \text{Boris Johnson} \rangle \rangle$, then p_φ is $\langle \varphi\text{-Toryness}, \langle \varphi(\text{Boris Johnson}) \rangle \rangle$. So, if φ maps Boris Johnson onto my toothbrush, p_0 is the proposition that Boris Johnson is a Tory, while p_φ is the proposition that my toothbrush is a φ -Tory. Given that a Russellian proposition is true iff its component objects (in order) instantiate the component relation, it follows that p_0 is true iff p_φ is true.

It's thus straightforward, given an atomic Russellian proposition, to use any arbitrary one-one mapping to generate a distinct but truth-conditionally equivalent Russellian proposition. This strategy generalizes in familiar ways to logically complex Russellian propositions and also to other conceptions of mental contents as structured (e.g. as structured Fregean senses). If we let an *interpretation* I be a mapping from an agent to a set of propositional attitudes, then we can use an arbitrary φ to generate a distinct interpretation I_φ that maps the agent to a set of propositional attitudes with truth-conditionally equivalent structured contents. If we say that these interpretations are *truth-conditionally equivalent*, the schematic argument for parochialism looks like this:

- (f) An interpretation I of an agent is adequate iff I satisfies constraints C_A .
- (g) If an interpretation I satisfies C_A , then so does any truth-conditionally equivalent interpretation.

So any view on which mental content has language-like structure faces the radical parochialist challenge—we do not rely on the vehicles of content, specifically, having this structure.

There are limits. The unstructured, coarse-grained view of Lewis and Stalnaker sidesteps the parochialist argument entirely. The parochialist argument allows that the assignment of truth-conditions to a propositional attitude is *not* a merely parochial matter. Since the coarse-grained view does not individuate propositional attitudes any more finely than by their truth-conditions, this leaves nothing to be a parochialist about—on this view, “truth-conditionally equivalent” interpretations are identical. So one way to avoid parochialism about *mental* reference is to endorse a coarse-grained view of mental content. However, such theorists still face the parochialist challenge for *linguistic* reference (see Appendix B).

²² It is straightforward to generate φ -variants for relations of greater adicity: $\phi(x)$ bears φ - R to $\phi(y)$ iff what $\phi(x)$ bears ϕ to bears R to what $\phi(y)$ bears ϕ to.

Appendix B

Section 2 considered instantiations of the argument for radical parochialism about *mental* reference, arguing that the argument succeeds if the purpose of theorizing about mental representation is merely to predict behaviour, but fails if the purpose is to articulate reasons for behaviour (and may either succeed or fail if the purpose is the explanation of behaviour, depending on the nature of explanation). In this appendix, we review how these considerations play out when applied to linguistic, rather than mental, representation.

First, even if radical parochialism is false of mental reference, and even if you take the view that mental representation is explanatorily or metaphysically prior to linguistic representation,²³ radical parochialism may yet be true of linguistic reference. For instance, suppose your favoured metasemantic theory says that a name N denotes_{*l*} x iff there is a concept C such that N expresses C and C denotes_{*m*} x ; while a predicate F is true of_{*l*} x iff there is a concept C such that F expresses C and C is true of_{*m*} x ; and so on (we use subscripts here to keep track of the distinction between linguistic and mental representation).²⁴ We can use an arbitrary one-one mapping φ to produce an alternative metasemantic theory that spits out a truth-conditionally equivalent reference-scheme: N denotes_{*l*} $\varphi(x)$ iff there is a concept C such that N expresses C and C denotes_{*m*} x ; F is true of_{*l*} $\varphi(x)$ iff there is a concept C such that F expresses C and C is true of_{*m*} x . The parochialist about linguistic representation can maintain that there is nothing to choose between the former and the latter as far as the purposes of linguistic “reference”-talk are concerned.

Second, the dialectic need not play out the same way for linguistic as for mental reference. The candidate purposes considered in section 2 were: prediction, explanation, and reason-articulation for behaviour. For linguistic representation, the thought might be that we can pare this down to the prediction, explanation, or reason-articulation of *linguistic* behaviour. But a major contrast now emerges. It is plausible that mental representational properties are directly involved in the production of behaviour, but equally plausible that linguistic representational properties are *not* directly involved in the production of linguistic behaviour—instead, what is most directly involved are mental states directed at linguistic

²³ E.g. Grice (1957), Lewis (1975); for recent examples, Davis (2003), Williams (2020).

²⁴ Many “head-first” metasemantic theories in fact go by pairing *sentences* with *propositional attitudes* (typically beliefs) and identifying the representational content of the former with the representational content of the latter. This story needs supplementing if it’s to tell us anything about the semantic properties of subsentential expressions. The question will be whether this supplementation helps us home in on an objectively important feature of linguistic representation or something of merely parochial interest.

representational properties. For example, it may be the fact that a subject *believes* that “Boris Johnson” refers to Boris Johnson, or (subpersonally) *cognizes* that this is so, that most directly predicts or explains their linguistic behaviour (e.g. Larson & Segal, 1995). If appropriate non-factive mental states are all we need to predict and explain linguistic behaviour, then the adequacy conditions that emerge could be met even if linguistic reference were merely a useful fiction! The parochialist conclusion would be that the fact that our semantic theories talk about a particular (actually instantiated) reference relation *rather than no relation at all* is merely parochial.

Articulating-reasons-for-linguistic-behaviour, on the other hand, gives rise to a potentially richer set of adequacy conditions than prediction or explanation would do. This will be so if one can establish that among the *normative reasons* for which one utters a certain sentence are semantic facts. For instance, suppose that among my normative reasons for uttering “Boris Johnson is a Tory” is the fact that “Boris Johnson is a Tory” means that Boris Johnson is a Tory (Heck, 2006). If so, then articulating the reasons for which I utter this sentence will require citing: (i) that I believe (or subpersonally cognize) that “Boris Johnson is a Tory” means that Boris Johnson is a Tory (for this to be a reason for which I acted); *and* (ii) that “Boris Johnson is a Tory” does, in fact, mean that Boris Johnson is a Tory (for this to be a reason for the utterance). In (ii) we have a role for linguistic representation itself, and not just non-factive mental states about language.

The moral is that the dialectic concerning the argument for radical parochialism about linguistic representation is not simply a replay of that concerning the argument for radical parochialism about mental reference—though some of the motifs are similar, they arrange themselves in different patterns. Also relevant are, e.g., contentious assumptions about the nature of linguistic competence and reasons for linguistic behaviour. That opens up the intriguing possibility that radical parochialism is false of mental reference, but a genuine insight concerning linguistic reference.^{25, 26}

References

²⁵ Schwarz (2014) attributes something like this combination of views to Lewis.

²⁶ Thanks Thomas Brouwer, Daniel Elstein, Pekka Väyrynen, Jack Woods, and two anonymous reviewers for *Noûs* for useful feedback and discussion! Material from this paper was presented by JRGW as “Semantic bookkeeping and semantic teleology” to departmental seminars in Edinburgh, Syracuse, Leeds, and Cambridge Trinity UG society, and as “What we talk about when we talk about talking about” in Oxford. JRGW thanks all with whom he discussed it on those occasions! The research leading to these results has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement no. 818633); and from the British Academy (Grant No. PF2\180082).

- Azzouni, J. (2000). *Knowledge and Reference in Empirical Science*. London: Routledge.
- Davidson, D. (1977). Reality without reference. *Dialectica*, 31, 247-253. doi: 10.1111/j.1746-8361.1977.tb01287.x
- Davidson, D. (1979). The inscrutability of reference. *Southwestern Journal of Philosophy*, 10, 7-19. doi: 10.5840/swjphil197910229
- Davis, W. (2003). *Meaning, Expression, and Thought*. Cambridge: Cambridge University Press.
- DePaul, M. (1987). Supervenience and moral dependence. *Philosophical Studies*, 51, 425-439. doi: 10.1007/BF00354046
- Devitt, M. (1983). Realism and the renegade Putnam. *Noûs*, 17, 291-301. doi: 10.2307/2215148
- Faroldi, F. & Protopopescu, T. (2019). A hyperintensional logical framework for deontic reasons. *Logic Journal of the IGPL*, 27, 411-433. doi: 10.1093/jigpal/jzz012
- Field, H. (1972). Tarski's theory of truth. *The Journal of Philosophy*, 69, 347-375. doi: 10.2307/2024879
- Field, H. (1975). Conventionalism and instrumentalism in semantics. *Noûs*, 9, 375-405. doi: 10.2307/2214521
- Field, H. (1978). Mental representation. *Erkenntnis*, 13, 9-61. doi: 10.1007/BF00160888
- Grice, P. (1957). Meaning. *Philosophical Review*, 66, 377-388. doi: 10.2307/2182440
- Heck, R. (2006). Reason and language. In C. Macdonald & G. Macdonald (Eds.), *McDowell and His Critics* (pp.22-44). Oxford: Blackwell Publishing.
- Larson, R. & Segal, G. (1995). *Knowledge of Meaning: An Introduction to Semantic Theory*. Cambridge, MA: The MIT Press.
- Lewis, D. (1974). Radical interpretation. *Synthese*, 27, 331-344. doi: 10.1007/BF00484599
- Lewis, D. (1975). Languages and language. In K. Gunderson (Ed.), *Minnesota Studies in the Philosophy of Science* (pp.3-35). University of Minnesota Press.
- Lewis, D. (1983). New work for a theory of universals. *Australasian Journal of Philosophy*, 61, 343-377. doi: 10.1080/00048408312341131
- Lewis, D. (1984). Putnam's paradox. *Australasian Journal of Philosophy*, 62, 221-236. doi: 10.1080/00048408412340013
- Lewis, D. (1993). Many, but almost one. In J. Bacon, K. Campbell, & L. Reinhardt (Eds.), *Ontology, Causality and Mind: Essays in Honour of D. M. Armstrong* (pp.23-37). Cambridge: Cambridge University Press.
- McGee, V. (2005). Inscrutability and its discontents. *Noûs*, 39, 397-425. doi: 10.1111/j.0029-4624.2005.00507.x

- Putnam, H. (1980). Models and reality. *Journal of Symbolic Logic*, 45, 464-482. doi: 10.2307/2273415
- Putnam, H. (1981). *Reason, Truth and History*. Cambridge: Cambridge University Press.
- Quine, W. (1960). *Word and Object*. Cambridge, MA: Harvard University Press.
- Schwarz, W. (2014). Against magnetism. *Australasian Journal of Philosophy*, 92, 17-36. doi: 10.1080/00048402.2013.765900
- Shoemaker, S. (1988). On what there are. *Philosophical Topics*, 16, 201-223. doi: 10.5840/philtopics19881617
- Sider, T. (2011). *Writing the Book of the World*. Oxford: Oxford University Press.
- Simons, P. (1987). *Parts: A Study in Ontology*. Oxford: Clarendon Press.
- Skolem, T. (1920/1967). Logico-combinatorial investigations on the satisfiability or provability of mathematical propositions: a simplified proof of a theorem by Löwenheim and generalizations of the theorem. In J. van Heijenoort (Ed.), *From Frege to Gödel: A Source Book in Mathematical Logic, 1879-1931* (pp.252-263). Cambridge, MA: Harvard University Press.
- Unger, P. (1980). The problem of the many. *Midwest Studies in Philosophy*, 5, 411-468. doi: 10.1111/j.1475-4975.1980.tb00416.x
- Wallace, J. (1977). Only in the context of a sentence do words have any meaning. In P. French & T. Uehling (Eds.), *Midwest Studies in Philosophy 2: Studies in the Philosophy of Language* (pp.144-164). Morris: University of Minnesota Press.
- Weatherson, B. (2013). The role of naturalness in Lewis's theory of meaning. *Journal for the History of Analytical Philosophy*, 1(10), 1-19. doi: 10.4148/jhap.v1i10.1620
- Williams, J. R. G. (2007). Eligibility and inscrutability. *Philosophical Review*, 116, 361-399. doi: 10.1215/00318108-2007-002
- Williams, J. R. G. (2008). The price of inscrutability. *Noûs*, 42, 600-641. doi: 10.1111/j.1468-0068.2008.00694.x
- Williams, J. R. G. (2020). *The Metaphysics of Representation*. Oxford: Oxford University Press.
- Williamson, T. (2013). *Modal Logic as Metaphysics*. Oxford: Oxford University Press.
- Williamson, T. (Forthcoming). Moral anti-exceptionalism. In P. Bloomfield & D. Copp (Eds.), *The Oxford Handbook of Moral Realism*.