

Undoing the Truth Fetish: The Normative Path to Pragmatism

Shamik Dasgupta
Draft of January 2019

What is pragmatism? There's no succinct definition—the term encompasses a cluster of ideas associated with the American philosophical tradition running from Charles Peirce and William James through thinkers such as John Dewey, C.I. Lewis, Nelson Goodman, and Richard Rorty.¹ But a central theme is an attempt to diminish the importance of *truth* and shift the focus onto *us*; onto our interests, practices, and needs. Thus, while others see scientific inquiry as primarily aiming to uncover truth, pragmatists tend to see it more as serving our interests. Insofar as they talk of truth, some pragmatists define it as whatever inquiry would converge on in the long run—witness Peirce's claim that truth is 'settled belief'. This is to regard truth as "immanent" rather than transcendent, insofar as it cannot outrun our (long run) practices of inquiry and the interests they serve. Pragmatists therefore tend to reject the correspondence theory of truth and indeed any theory of truth based on inflationary word-world (or mind-world) relations of "representation" or "reference". Likewise, they reject accounts of meaning in terms of truth-conditions and focus more on practical consequences instead.

These then are four central theses of the pragmatist tradition:

Pragmatic theory of inquiry: The aim of inquiry is to serve our interests. A theory is a tool that serves our interests; our best theories are those that best serve those interests.

Immanent theory of truth: Truth is nothing other than that which inquiry would converge on (in the long run). It's therefore impossible for an ideal theory—a theory as confirmed as can be—to be false, and (vice-versa) impossible for there to be truths that inquiry could never uncover.

Anti-representationalism: Inflationary relations of "representation", "reference", and "correspondence" between language (or mind) and world should be rejected as philosophically idle.

Pragmatic theory of meaning: The meaning of a hypothesis is given by its practical consequences.

This will all seem ridiculous to philosophers who continue to give truth a central place in philosophy. These "truthers" think it *obvious* that there are truths on which inquiry will never converge, such as how many blades of grass existed during the year 1749; hence truth transcends inquiry *contra* the immanent theory of truth. They'll insist that when my daughter asks whether cows are ungulates I should find out the *truth*, not just whatever gets her off my back, *contra* the pragmatic theory of inquiry. As for the pragmatic theory of meaning, truthers will likely regard it as no more promising than the related verificationist doctrine that meaning consists in observable consequences. Some truthers will go on to say that the spectacular successes of recent truth-conditional semantics in philosophy and linguistics pretty much *refutes* any approach to meaning that ignores truth, and indeed demonstrates the importance of notions

¹ Misak (2013) presents an illuminating history of this tradition.

like reference and representation. To the truthers, then, these four pragmatist theses come across as vestiges of an outdated naiveté.

Nonetheless, I believe these pragmatist theses are by and large correct—and indeed that they follow from a normative claim that’s widely accepted by contemporary truthers! My aim here is to chart out this normative path to pragmatism.

The key normative claim is that notions like truth and reference are, in themselves normatively inert. To be sure, deflationists have long argued that they’re *explanatorily* inert, but the claim is that even if there are explanatory notions of truth and reference, they nonetheless hold no intrinsic normative significance. Sections 1-5 articulate this claim; sections 6-8 then argue that it leads to the pragmatist theses above.

To be clear, I won’t argue for the normative claim in any seriousness. To my mind it’s one of those claims in philosophy that becomes almost undeniable once articulated properly, which is why half the paper tries to do just that. And like I said, many self-identified truthers already accept it anyway. But I’ll end by describing the picture of reality that follows if you reject it. I find that picture incredible, so I embrace the pragmatist theses. But committed truthers may instead read the paper as revealing the picture they’re (implicitly) wedded to.

I also won’t do much in terms of scholarship to compare the pragmatism I develop with the pragmatist tradition. Pragmatists come in many varieties and I suspect that some will regard my reading of the four theses as not *real* pragmatism, while others will see me as merely retracing well-trodden pragmatist ground. To the latter I suggest there may be value in retracing this ground in terms that truthers will likely accept, and to the former I say that I don’t much care what we *call* it; my aim is just to develop an interesting position in the spirit of these four theses.²

1. How to go on

Our starting point is the Kripkensteinian question of *how to go on*.³ Consider the sign ‘+’. You’ve used it a finite number of times to record a finite number of calculations, writing down things like

$$\begin{aligned}5 + 7 &= 12 \\17 + 8 &= 25\end{aligned}$$

Now you try calculating ‘68 + 57’, and imagine you’ve never calculated numbers greater than 50 before. What’s the answer? We think it’s obviously 125, but Kripke demurs. What’s obvious is that *if you mean addition by ‘+’* then the answer is 125. But what if you mean quaddition instead? Quaddition is the function that maps two numbers x and y to their addition when x and y are both smaller than 50, and to the number 5 otherwise. If you mean quaddition by ‘+’, the answer to your calculation is 5! And Kripke’s “skeptical argument” purports to show that there’s no fact of the matter as to whether you mean addition rather than quaddition. Your past usage of ‘+’ is consistent with both meanings, by hypothesis. And the argument is that no other facts about you—your dispositions, intentions, etc.—determine what you mean either; hence there’s

² Readers familiar with Stich (1990) will find that my path to pragmatism resembles his in many respects. But I won’t compare our approaches in detail.

³ See Kripke (1982); Wittgenstein (1953).

no fact of the matter what the right answer to your calculation is. Put otherwise, there's no fact of the matter whether the sentence

(1) $68 + 57 = 125$

or

(2) $68 + 57 = 5$

is correct. Or in more Wittgensteinian terms, there's no "right" way to go.

The issue is whether there's a *correct* answer—whether there's a *right* way to go on using '+'—but this needs clarification. By the "correct answer", I don't mean what you'd be *justified* in believing, given your evidence. If you mean addition by '+' then the correct answer is 125 even if you don't justifiably believe so—indeed, even if you don't *know* that you mean addition. Thus, the issue does not concern the epistemology of meaning or what "guides" a subject in applying a sign. While some discussions of this topic focus on issues like these, I do not.⁴

Moreover, what's not in question is the mathematical fact that 68 added to 57 equals 125; what's at issue is whether '+' means addition. More explicitly, then, the issue can be put like this. You're given the question ' $68 + 57 = ?$ ' and your aim is to find the *correct* answer. Schematically, say that a sentence S is correct iff two conditions are met: (i) S means that p, and (ii) p. The second condition isn't at issue: we're assuming that 68 added to 57 equals 125. The issue is what (1) means, whether it means that 68 added to 57 is 125, or that 68 *quadded* to 57 is 125. If the former, it's correct; if the latter it isn't. Kripke's skeptical argument, as I'm interpreting it, purports to show that there's no fact of the matter what (1) means.

If you find talk of 'meaning' obscure, we can put the issue directly in terms of correctness-conditions instead. Schematically, a sentence S is correct iff (i) the correctness-condition of S is that p, and (ii) p. Then the issue concerns (1)'s *correctness-condition*—whether it's correct iff 68 added to 57 equals 125, or correct iff 68 quadded to 57 equals 125. Here I'll slide freely between meaning and correctness-conditions—by talk of 'meaning' I mean only that expressions have conditions under which they are correctly and incorrectly applied. Framed like this, Kripke's skeptical argument purports to show that no facts about us fix these correctness-conditions.

But if the issue is *correctness* and not justification, isn't it really an issue about *truth*? To answer correctly is to answer truly, the thought is, so the issue must concern (1)'s *truth-condition*—whether it's true iff 68 added to 57 makes 125, or true iff 68 quadded to 57 makes 125. And that depends on whether '+' *refers* to addition or quaddition. And what Kripke's skeptical argument purports to show is that there's no fact of the matter about *that*.

I think this is a mistake. Fix reference and truth-conditions all you want, that still doesn't determine the correct way to go on!

2. The truth fetish

⁴ See Wright (2001) and Merino-Rajme (2015) for discussions that focus on the epistemology and phenomenology of meaning and rule-following.

To see why, consider the theory of reference associated with David Lewis (1984) on which reference is fixed by use plus naturalness. The idea is this. Some functions fit your usage of ‘+’ better than others: if in the past you’ve tended to produce their addition when computing ‘ $x + y$ ’, then addition fits your usage well while subtraction doesn’t. But as Kripke showed, that’s not enough to fix reference since quaddition fits equally well. More is needed, and the idea is that some functions are *more natural* than others. Here we appeal to the metaphysical thesis that some entities (properties, relations, functions, whatever) are *perfectly natural*, and that amongst the rest some are *more natural* than others. Green is more natural than gerrymandered properties like grue, for example. Then the proposal is that the referent of an expression is whatever best satisfies these two constraints of fit with use and naturalness. If addition is more natural than quaddition (and other gerrymandered functions that fit usage), ‘+’ refers to addition.

One can then use reference to define truth in the standard Tarskian manner. For example, a base clause might say

A sentence of the form ‘ a is F ’ is true if and only if the referent of ‘ a ’ has the property referred to by ‘ F ’.⁵

Appropriately developed, such a theory would imply that (1) is true iff the *addition* of 68 and 57 is 125. *Et voila*: we have a theory on which (1) has a determinate truth-condition after all, *contra* Kripke’s skeptical argument!

On this theory, Kripke is right that facts about *you* don’t fix whether you refer to addition by ‘+’. But reference isn’t fixed just by you. It’s fixed by you *and the world*, and the world helps determine that ‘+’ refers to addition and not quaddition. I said this theory is *associated* with Lewis, not that Lewis accepts it—his considered view was rather more complex.⁶ Still, it’s a simple example of a theory on which truth and reference are fixed by some mixture of usage and the world. I claim that *theories like this imply nothing about correctness*—they settle *nothing* about the right way to go on using ‘+’! This Lewisian theory illustrates the point well because of its simplicity.

To see this, assume that the theory works as advertised: addition is perfectly natural and quaddition isn’t, so ‘+’ refers to addition and not quaddition. Still, if there’s a property of perfect naturalness that addition has and quaddition lacks, there’s also a property of perfect *quaturalness* that quaddition has and addition lacks. After all, properties are cheap: there is a property for every set. At least, this was Lewis’ own view of properties, and likewise for relations and functions: there’s an n -place relation for every set of n -tuples, and a function for every relation that relates each thing in its domain to a unique thing in its range. He posited perfect naturalness precisely to distinguish a select few of these entities from the rest. This works *to an extent*: it divides those that are natural from the rest. But if properties are cheap, there’s also a property of perfect quaturalness that distinguishes a *different* select few entities from the rest. Addition may be perfectly natural, but quaddition is perfectly quatural!

⁵ This is obviously a simplification. A more realistic theory might associate predicates with satisfaction-conditions or extensions rather than referents, but these details are not relevant here.

⁶ See Weatherston (2012), Schwarz (2014), and Williams (2015) for further discussion.

So, if the relation of *reference* is fixed by use plus naturalness *per* the Lewisian theory, there's equally a relation of *queference* fixed by use plus *quataturalness* in a precisely analogous manner. The referent of '+' is that function which best satisfies the two constraints of fit with use and naturalness, namely addition. And the *queferent* of '+' is that function which best satisfies the two constraints of fit with use and *quataturalness*, namely quaddition! Your expression '+' stands in *both* relations at once: it *refers* to addition and *quefers* to addition at the very same time. Just because the Lewisian theory didn't mention quataturalness and queference doesn't mean they're not there.

And just as the Lewisian theory defined truth in terms of reference, one may analogously define a notion of *quuth* in terms of queference. The base clause would be

A sentence 'a is F' is *quue* if and only if the queferent of 'a' has the property queferred to by 'F'

If the Lewisian theory implies that (1) is *true* iff 68 *added* to 57 is 125, this theory will imply that (1) is *quue* iff 68 *quadded* to 57 is 125. These theories aren't in conflict, for truth and quuth are different properties: (1) has a truth-condition *and* a quuth-condition at the very same time.

We started with a normative question of how to go on using '+'. When computing '68 + 57' for the first time, what's the right thing to say? The hope was that the Lewisian theory settles this: if '+' refers to addition then the correct answer is 125, not 5. But don't be fooled: it depends on whether it's right to go on in accordance with the referent of '+' or its queferent. If the former the correct answer is 125, but if the latter it's 5! Remember, both relations are out there: '+' refers to addition *and* quefers to quaddition at the very same time. Absent some further claim to the effect that reference is normatively significant—that the evaluation of speech as correct or incorrect should organize around reference rather than queference—*nothing* in the Lewisian theory implies that the correct answer is 125.

Similarly for truth. Grant that on the Lewisian theory (1) is true and (2) is not. Still, (2) is *quue* and (1) is not! So what's the right answer to your computation, (1) or (2)? That depends on whether the right (correct) answer is the one that's *true* or the one that's *quue*. Absent some further claim to the effect that truth is normatively significant in this regard, nothing in the Lewisian theory implies that (1) is the right answer.

You might say it's a platitude that truth is normatively significant; that 'true' is a *label* for whatever makes speech correct. But this just shunts the issue upstream. For the question is then whether 'true' labels use plus naturalness *per* the Lewisian theory, or use plus quataturalness instead. And *that* depends on whether naturalness or quataturalness plays this normative role of determining what counts as correct speech. This is just our question of normative significance all over again.

The point is that for all the Lewisian theory says, truth is just one property out there fixed by one particular mixture of usage and the world. But quuth is another property fixed by a different mixture, and there are countless other mixtures besides. Which mixture plays this normative role of determining how to go on, of fixing what counts as the right thing to say? Without the further claim that one mixture is normatively significant, no progress has been made. To think otherwise is to *fetishize* one mixture over others. It is this fetish I want to undo. If you insist on

using 'true' to denote the Lewisian mixture of use plus naturalness, fine—the point is that we must not fetishize truth.⁷

This is not Putnam's "just more theory" objection.⁸ His point was that there is no *extra* constraint of naturalness that helps fix reference along with fit with usage; there's just fit with usage, which includes fit with our use of the term 'naturalness'. More generally, his objection was that there *couldn't* be extra constraints over and above fit, since any constraint would collapse into fit with our use of the terms used to express the constraint. But my point is the opposite. It's not that there *couldn't* be extra constraints, it's that there are *too many*. There's an extra constraint of naturalness that helps fix reference, granted, but there's also a constraint of quaturalness that helps fix queference. Without an extra claim that naturalness is normatively significant, nothing follows about whether to go on using '+' according to its referent or its queferent.

Nor is this Kripke's objection to dispositional theories of meaning. Kripke argued (i) that our dispositions run out too fast (beyond some limit we have no dispositions to use '+' one way or another) and (ii) we can be disposed to use an expression incorrectly. The Lewisian theory employs naturalness to solve these problems: once our dispositions run out naturalness takes up the slack to fix a determinate referent, and since naturalness competes with usage the referent may diverge from how we're disposed to apply the term. My point is that *if* naturalness does this work, quaturalness also does a parallel job: it can take up slack when our dispositions run out to fix a determinate *queferent*, and this queferent can diverge from our dispositions of use for the same reason. Again, absent a further claim that naturalness is normatively significant, nothing follows about how to go on using '+'.

I've discussed the Lewisian theory, but the same goes for theories on which reference and truth-conditions are fixed by causal relations between language and world. For relations are cheap, so along with causation and there's also quausation, a relation that overlaps with causation in familiar cases but diverges elsewhere. If causation fixes the relation of reference, quausation fixes *another* relation of queference in a precisely analogous manner. Your expression stands in *both* relations to the world at once: it has a reference and a queference at

⁷ I've heard it said that truth is special because it satisfies the equivalence scheme:

'S' is true if and only if S

For example, suppose the Lewisian theory implies that 'snow' refers to snow and 'white' refers to whiteness. Then the Tarskian definition of truth gestured at above would imply

'Snow is white' is true if and only if snow is white.

And the thought is that this doesn't go for quuth. Suppose the mix of use and quaturalness implies that 'snow' *quefers* to grass and 'white' quefers to greenness. Then the corresponding definition of quuth would imply

'Snow is white' is que if and only if grass is green.

But this is too quick. For the equivalence scheme ranges over *interpreted* sentences of the meta-language, and the meta-language can be interpreted according to reference *or* queference. If the latter, we shouldn't say that 'snow' quefers to grass; we should say that 'snow' quefers to snow! The definition of quuth then yields

'Snow is white' is que if and only if snow is white.

and we have the equivalence scheme after all. Indeed, so interpreted we can't anymore say that 'snow' refers to snow and so *truth* won't satisfy the equivalence scheme! The upshot is that truth and quuth both satisfy the equivalence scheme, and for the same reason. If we interpret the object language according to reference then truth satisfies it and quuth doesn't. But equally, if we interpret according to queference then quuth satisfies it and truth doesn't.

⁸ Putnam (1977) objected to a causal theory of reference along these lines; Lewis (1984) discussed how the objection would apply to the Lewisian theory under discussion. Thanks to Rohan Sud for a helpful discussion of this point.

the very same time. And the normative question remains as to which relation determines the right way to go on using the expression.

Indeed, the same goes for any theory of reference and truth on which they're fixed by some mixture of usage and the world. For such mixtures are cheap, other mixtures will fix other word-world relations, so the normative question remains as to which mixture determines the right way to go on. Likewise if you say that reference or truth-conditions are primitive features that don't reduce to *other* facts about usage and the world. There may be primitive whatnots out there, but the normative question remains as to whether they determine the right way to go on.

3. Realism and anti-realism about meaning

To be clear, I haven't said that *no* mixture is normatively significant. I've just said that absent a claim that one of them *is* significant, nothing follows about what the correct answer to '68 + 57' is. Thus, it's not enough to identify some mixture and call it 'reference'; one must also claim that it has the normative status of determining the *right* way to go on using '+'.

Suppose then that one mixture has this special normative status. What would make it special? Why does *it* determine the right way to go on? Here we must distinguish the view that it's special *because of us* from the view that it's special *independently of us*. I call the first view anti-realism about meaning, the second view realism. This distinction is central to what follows, for anti-realism is the normative claim that, I'll argue, leads to our four pragmatist theses.

But first let me clarify the distinction. Our words stand in *countless* different relations to the world, as we've seen. Communication requires that we organize our linguistic practices around one relation and use it to evaluate speech as correct or incorrect. According to anti-realism all relations are on a par "objectively speaking"—i.e. when considered independent of us, our needs and natures and projects. We use the one we do just because it suits us well; because it serves our "interests", broadly construed. If that turns out to be the Lewisian relation fixed by use plus naturalness, fine; the point is that it plays this role of determining correct speech only because *we* are well served by it. Other communities of organisms with different natures, goals, or needs might do better to coordinate around something else, perhaps queference, and in doing so they would be doing exactly the right thing given their interests. By contrast, the realist thinks that one relation is "objectively" special in the sense that it determines what *really* counts as correct speech independently of anyone's interests. On this view, at least one community is going "wrong" insofar as they aren't organizing their linguistic practices around *that* relation.

Note that the issue here is not whether reference depends on us. The realist need not deny the platitude that had we used 'snow' differently it would have referred to something else. What's independent of us, on his view, is that *reference* and not queference determines correctness-conditions.

This point is worth marking with terminology. Semantics, let's say, is the theory that assigns semantic values—referents, satisfaction-conditions, truth-conditions, whatever—to expressions of a language. Meta-semantics, then, is the theory of what fixes the semantic value of an expression. The Lewisian theory that reference is fixed by use plus naturalness is an example. At the risk of multiplying hyphens, "meta-meta-semantics" is then the theory of why it's *semantic* values, not quemantic values, that *matter*. It's the theory of why semantics is in the business of assigning expressions to their *referents* rather than their queferents in the first place. Realism

and anti-realism are theories of meta-meta-semantics. That's why realism is consistent with the platitude that the referent of 'snow' depends on how we use it, for that's a platitude of meta-semantics. Equally, anti-realism is consistent with theories of semantics and meta-semantics which, like the Lewisian theory, give reference and truth a central role. It can agree that the wordly facts that help fix reference and truth-conditions, such as facts of naturalness, hold independently of us. What it says is that reference and truth (and whatever fixes them) play this normative role of regulating our linguistic practice because they serve our interests.⁹

I'm using 'interests' loosely here. Perhaps organizing our linguistic practice around a particular word-world relation furthers our preferences or goals; perhaps it suits our history or culture; or perhaps it fits well with our human psychology or biological make-up. These would all count as "serving our interests" in the loose sense I intend. Anti-realism is therefore a broad church. While some anti-realists will say that the relevant interests can vary from culture to culture, others will think they're uniform across humanity.

But when the anti-realist says that reference is special because of our *interests* (whatever that means), she must not take this too seriously. For along with interests that make reference special, we also have quinterests that make queference special. Quinterests are strange properties to which we pay little attention, but since properties are cheap there must be *some* such properties that stand to queference just as interests stand to reference. Thus, the anti-realist may say "Reference is special in virtue of our interests", but that's just because she already speaks a language that includes those terms. In reality, our interests confer no more significance on reference than our quinterests confer on queference! So at bottom the anti-realist picture is just that *we do what we do*. Inside our language game (as it were) we can say that reference is special because of our *interests*, but outside of it there's nothing really special about interests.

This is a queasy result, no doubt, but one avoids it only if one says that interests have some prior normative significance over quinterests. And the anti-realist cannot say this. For interests and quinterests are structurally analogous properties of us, so if one has significance over the other that cannot be because of anything "about us" in the sense I have in mind. It would therefore be significant "independently of us"—its significance would be conferred from elsewhere, or else be primitive. Consequently, any significance that reference enjoys over queference would ultimately have its source independently of us too. And that is realism—at least, it violates what I mean by the anti-realist slogan that reference is special "because of us".

Relatedly, anti-realism includes views on which truth and reference are special because of the constitutive nature of our practices. Suppose one says that truth is a constitutive aim of assertion—that one counts as making an assertion only if one aims to speak truly—hence truth is special because only true assertions fulfill their aim. Still, all that follows is that *if* one is in the game of making *assertions* then truth is special. There is of course the alternative game of

⁹ Readers familiar with Sider (2011) may note a resemblance between anti-realism and what he calls a "projectivist" theory of beauty on p. 57. This theory holds that some physical property P causes an aesthetic reaction in us; we use 'beautiful' to denote that property P; but P is highly disjunctive, so that things with P have nothing in common except their causing this reaction in us. On this view, the fact that Turner's paintings are beautiful is just the fact that they have P, a fact that holds independently of us. Nonetheless, we've organized our aesthetic lives around P only because of its effect on us; there's nothing independently significant about it. Anti-realism holds that 'reference' and 'truth' work similarly.

quassertion, which is like assertion except its constitutive aim is quuth. If one is in the business of quasserting then the right thing to do is utter *quuths*, not truths! This is anti-realism, for truth is special on this view only insofar as it fits the game *we* play.¹⁰

I've focused on linguistic meaning, but the same issue of realism and anti-realism arises for thought. You're computing '68 + 57' mentally, and the question is what mental state counts as the right answer. You form a belief you'd verbalize as '68 + 57 = 125', but is it correct? Schematically, it's correct iff (i) its correctness-condition is that p, and (ii) p. As before, condition (ii) isn't at issue: we're assuming the mathematical fact that 68 added to 57 is 125. The issue concerns (i), whether your mental state is correct iff 68 added to 57 is 125, or correct iff 68 quadded to 57 is 125.¹¹

Suppose you now offer a theory of *mental content* on which the content of a mental state is fixed by some mixture of mind and world. Perhaps the mix involves causal relations, or perhaps it's a Lewisian mix of fit plus naturalness like before (where "fit" now involves fitting the functional role of the mental state, not the use of a linguistic expression). Suppose your theory implies that the content of your belief is that 68 added to 57 is 125. Does it follow that your belief is correct? No, for if your belief has a content fixed by some mix of causation or naturalness, it also has a *quontent* fixed by an analogous mix of quasation or quaturalness. Its content is that 68 added to 57 is 125, but its quontent is that 68 *quadded* to 57 is 125—it has a content *and* a quontent at the very same time! Is your belief *correct*? Is it the *right* answer? That depends on whether beliefs are to be evaluated with respect to their contents or their quontents! Absent some further claim that the particular mix described by your theory plays this normative role in fixing the right way to go on in thought, *nothing* follows as to whether your belief is correct. To think otherwise is to fetishize one mixture over others. If you use 'content' to label your particular mixture, fine—the point is that we must not fetishize content.

As before, I'm not saying that your mixture *isn't* normatively significant in this regard. Perhaps it is. But if so, we can ask what makes it special. Why should beliefs be evaluated relative to *it*? Here we find the issue of realism and anti-realism again, of whether the mix is special because of us (anti-realism) or independently of us (realism). Our mental states have myriad properties—contents, quontents, etc.—fixed by myriad mixtures of mind and world. For the anti-realist, we focus on the one we do because it suits our interests to evaluate mental states with respect to it rather than the others. Whereas for the realist, one of these properties is "objectively" special in the sense that it is the *real* standard of evaluation regardless of whether it's in anyone's interest to attend to it.

To be clear, this question of realism vs anti-realism isn't the question of why true beliefs are valuable—at least, not as the latter has been discussed in recent epistemology literature.¹² That literature takes a theory of content and correctness for granted: it assumes that beliefs have contents fixed somehow or other, and that the correctness-condition of a belief is given by its

¹⁰ Enoch (2006) made a similar point about ethical norms: if they're grounded in the constitutive nature of certain practices this does nothing to confer significance upon those practices themselves.

¹¹ Earlier the bearers of correctness were sentences, here they are mental states. It's less obvious how to individuate mental states than sentences, but the issues here would detract from the main thread. So long as we can individuate them without explicitly mentioning their correctness conditions—e.g. by their computational role—the same issues arise. See Boghossian (1990), section 6.

¹² For example, see Loewer (1993), Goldman (1999), Alston (2005), and Grimm (2009).

content not its content. The question is then what value there is in accumulating correct beliefs. This is not obvious, for it can pay to hold incorrect beliefs: if forced by gunpoint to believe that 68 added to 57 is 5, I'll do my best to oblige. But our question has to do with what fixes correct-conditions in the first place. If one uses 'content' as a placeholder for correctness-conditions, our question is what mixture of mind and world fixes content and in particular whether it plays this normative role of determining correctness because of us or independently of us.¹³

This question of realism vs anti-realism arises for any theory of content, but it can take work to expose it. Consider the theory proposed by David Lewis (1983) on which content is fixed by fit plus humanity. Humanity, he says, calls for "interpretations according to which the subject has attitudes that we would deem *reasonable* for one who has lived the life that he has lived" (p. 375; my emphasis). So on this view content is fixed by a mixture of fit and *reason*. One might then say that this mix is normatively significant precisely because of the latter ingredient: reason is an evaluative standard, so the mix of fit and reason is too. Is this realism or anti-realism about content? On the face of it, that depends on what makes *reason* normatively significant. If it has this status as an evaluative standard independently of us, we have realism; if not, we have anti-realism. Lewis doesn't explicitly answer this question, but let's see what it amounts to on his view. In saying what makes an attitude reasonable, i.e. one that humanity calls for, Lewis invokes naturalness: humanity will "impute a bias toward believing that things are green rather than grue" precisely because green is more natural than grue (p. 375). But if naturalness helps fix reason, *quaturalness* will fix a parallel virtue of *queason*. It may be reasonable to believe that things are green, but if grue is more quatural than green then it's more queasonable to believe things are grue! And just as humanity favors reason, a corresponding principle of quumanity favors queason. And so if mental states have contents fixed by fit plus humanity *per* Lewis' theory, they also have contentents fixed by fit plus quumanity. It follows that reason and content are normatively significant only insofar as *naturalness* is normatively significant. Thus, the focus on reason was, in a way, a smokescreen: for Lewis, the question of realism vs anti-realism really amounts to the question of whether *naturalness* has significance independently of us or because of us.

4. Deflationism

Anti-realism says that a word-world (or mind-world) relation is made special *by us*. But how can this happen? How could *we* confer significance on it? One answer is found in the deflationary theory of meaning developed in Field (1994) and (2001). It will prove useful to have this theory in hand as a concrete model of anti-realism.¹⁴

Field's theory is really a theory of meaning *attribution*; of what we do when we say what something means. According to this theory, meaning attributions report *good translations*. Thus, to say that 'neige' means snow is to say that 'neige' is well translated as 'snow' in one's own idiolect. And to say that 'neige et blanc' means that snow is white is to say that 'neige et blanc' is well translated as 'snow is white' in one's idiolect.

¹³ A similar distinction is found in ethics. Normative ethicists might take a notion of moral goodness for granted and ask what reason one has to promote good states of affairs. By contrast, meta-ethicists might ask what makes a state good in the first place, for example whether goodness consists in certain natural properties of the state.

¹⁴ Field develops a variety of deflationist views; here I focus on the Quinean variety developed towards the end of (2001).

There are two core ideas here. The first is that in stating what an expression E means I state a relation between E *and another linguistic expression*, not between E and some non-linguistic entity that is its “meaning”. Thus, in saying that ‘neige’ means snow I state a relation it bears to ‘snow’, not to the stuff found on mountain tops. Likewise, in saying that ‘S’ means that p I relate ‘S’ not to some intensional entity such as a proposition or truth-condition, but to the sentence ‘p’.

The second idea is that the relevant relation between expressions is good translation, where what makes for good translation is a pragmatic matter. Suppose I want my friend to give me a cookie. What should I say? If she’s British, saying ‘cookie’ may be ineffective. So I look for a word she typically associates with cookies and causes her to hand over cookies. If her word ‘biscuit’ does this, I’ll translate it as ‘cookie’. This is a good translation not because it reflects some prior fact about what the terms mean, but because it gets me a cookie.

So on this deflationary view it’s not that each expression has a *meaning*, such that the job of translation is to pair expressions with the same meanings. Nor is there a prior relation of *sameness of meaning* between expressions that the job of a translation is to reflect. Rather, there are just facts about which translations into one’s own idiolect work well for the purposes at hand. Meaning attributions simply report these facts.

What is an idiolect? For Field, my idiolect consists in my current understanding of language, which in turn consists in the computational role that each expression plays *for me, now*. Idiolects are therefore carved thinly: it is the idiolect of *a person at a time*. So understood, deflationism comes to this: to say that ‘neige’ means snow is to say that ‘neige’ is well translated as ‘snow’ *as I currently understand it*.

This is just a sketch, but we can already see how this theory bears on the Kripkensteinian question of how to go on.¹⁵ You’re trying to compute ‘68 + 57’ for the first time. Your aim is to give the correct answer, where a sentence S is correct iff (i) S means that p, and (ii) p. Suppose you utter

(1) 68 + 57 = 125.

Remember, the mathematical fact that 68 added to 57 equals 125 is not in question; what’s in question is whether your utterance means that 68 added to 57 makes 125. Is that what it means? Well, for *me* to say that it means this is to say that it’s well translated as

(1-add) 68 added to 57 equals 125.

in my idiolect. And that’s just to say that a good translation would match your ‘+’ with my ‘addition’. And of course this is what a good translation would do—it would be *utterly perverse* to translate your ‘+’ as my ‘quaddition’!¹⁶

¹⁵ Horwich (1995) discusses how deflationary theory of *truth* might approach the question of how to go on. But it’s unclear whether the account of meaning he appeals to is as deflationary in Field’s sense.

¹⁶ Likewise, for *you* to say ‘(1) means that 68 added to 57 equals 125’ is just for you to say that your ‘+’ is well matched with your ‘addition’, which of course it is. Sometimes the Kripkean issue is framed as whether your current usage is correct given what your *past self* meant by ‘+’. But this just takes us back to the third-personal case, only now with your past self as the third person. For you to say that your past self meant addition by ‘+’ is to say that a

To be clear, this perverse translation is *consistent* with your prior usage—it wouldn't misrepresent some prior fact about *the meaning of '+' in your idiolect*. But translation isn't about reflecting such facts, it's about matching expressions to suit our needs and matching your '+' with my 'addition' obviously does the best job. Thus, this is not so much a solution to Kripkensteinian as a dissolution. The claim is *not* that there are facts about you and the world that fix what you mean *independently of our pragmatic aims in translation*. The claim is that meaning attributions aren't in the business of tracking such facts; they just record *useful translations*, and perverse translations are obviously bad on pragmatic grounds. Perhaps there are bizarre circumstances in which translating your '+' as my 'quaddition' would suit our needs better, but they'd be very alien to ordinary life.

Field calls this a “deflationary” theory of meaning insofar as truth and reference have played no role. The fact that '+' means addition hangs on useful mappings between idiolects, not what worldly entity it refers to. He contrasts this with “inflationary” theories in which truth and reference a central role. The Lewisian theory from section 2 is a paradigm example, insofar as meaning (correctness-conditions) is fixed a relation of reference between word and world. So too was the causal theory discussed there.

But while truth and reference have so far made no appearance in Field's theory, they may yet enter picture in two ways—one familiar, the other less so. The familiar way is through *disquotational* notions of truth and reference. To say that a sentence 'S' is disquotationally true, or true_d for short, is just to say that S; the two statements are cognitively equivalent. Thus, all instances of the equivalence scheme

'S' is true_d if and only if S

hold of 'conceptual necessity', as Field (1994) puts it, for any sentence 'S' in one's idiolect. It's controversial whether the English word 'true' is disquotational in this sense, but put that aside; what's uncontroversial that the deflationist can introduce true_d for its well-known utility of allowing one to formulate various infinite conjunctions and disjunctions.¹⁷ Likewise, she may introduce a disquotational notion of reference—or reference_d for short—on which 'the referent_d of 'N'' is cognitively equivalent to 'N'. These disquotational notions play no role in her theory of meaning, they're just useful logical devices.

The second way in which truth and reference can enter the deflationist picture is more subtle. To see how, note that a deflationist shouldn't deny that there are interesting relations between language and world. When I translate my friend's word 'biscuit' as 'cookie' I do so because she associates 'biscuit' with cookies, she tends to reach for cookies when I say "Give me a biscuit", and so on. So the fact that it's a good translation presumably turns on some complex causal relation between her word 'biscuit' and cookies.

But this causal relation might be what an inflationist called “reference” all along! If so, have we collapsed into inflationism? Here we must take care. One might argue that we haven't

good translation would match '+' in your past idiolect with 'addition' in your current idiolect, which is (obviously) what a good translation manual would do.

¹⁷ Some deflationists argue that this is the *only* utility of truth and argue on that basis that true_d is the English 'true'. See Horwich (1995).

collapsed on the grounds that while this particular causal relation is relevant to translating ‘cookie’, other words work differently. The idea is that gastronomic vocabulary is used for certain purposes, but moral vocabulary is used for other purposes and mathematical vocabulary for yet others, so that what makes for good translation might differ across these domains. Thus, good translation of gastronomic talk might track one word-world relation, but for mathematical and moral talk it might track others. Yet the term ‘reference’ is used uniformly: ‘good’ refers to goodness, ‘+’ refers to addition, and ‘cookie’ refers to cookies. If so, there’s no *non-disjunctive* word-world relation that can be identified as *reference* across the board. For some, this is a central argument against inflationism.¹⁸

I won’t put much weight on this point. No doubt different vocabularies serve different purposes—this is an important insight to which we’ll return. But to conclude that there’s no word-world relation of *reference* relies on a distinction between disjunctive and non-disjunctive relations, and in particular the claim that the former are in some way second-rate. These are questionable metaphysical commitments. So for the sake of argument I’ll grant that there is a word-world relation that good translation always tracks across the board, so long as we remember that it may strike us as a wildly disjunctive miscellany. This relation fixes good translation and hence correctness, and we can *call* it “inflationary reference” if we want. Still, on this view it has no *intrinsic* normative significance; there’s no fact *independent of our interests in translation* in virtue of which *it* determines what counts as correct. There are myriad relations out there—reference, queference, etc.—and the former fixes correct speech only because it happens to serve our translational interests.

If we’ve collapsed into inflationism, then, we’ve done so in an interesting way: we’ve collapsed into a manifestly anti-realist variety! As I see it, this is the central insight of Field’s theory. No doubt there’s *some* relation of “inflationary reference” that fixes correct speech, even if it’s just a disjunctive miscellany. And grant that this relation determines *some* (equally disjunctive) property that all and only correct sentences have, which we can *call* “inflationary truth” if we want. Still, what Field’s project implies is that they play the normative role of determining how to go on only because they *serve our translational interests*. This is anti-realism through and through.

Seen like this, Field’s deflationism is really a theory of *meta-meta-semantics*, a theory of why one word-world relation has normative significance. So understood, it’s in no tension with so-called “inflationary” theories of semantics and meta-semantics that give reference and truth a central role! It is sometimes said that deflationism is refuted by the tremendous fruits of recent truth-conditional semantics in philosophy and linguistics, which make essential use of notions like truth and reference. My point is that this is wrong insofar as we see deflationism as primarily a theory of meta-meta-semantics; a theory about why a particular word-world relation matters.

Admittedly, Field didn’t see deflationism like this. He said that “if deflationism is to be at all interesting, it must claim not merely that what plays a role in meaning and content not include truth conditions *under that description*, but that it not include *anything that could plausibly constitute a reduction of truth conditions to more naturalistic terms*.” (1994, p. 253). As I read him, his worry is precisely that our translational practices might turn out to track something that could constitute a reduction of inflationary truth-conditions. But I think he lost his nerve here. His worry is understandable insofar as he regards deflationism as an *alternative* to truth-conditional

¹⁸ See Price (2011, pp. 14-16). Field (1994, section 2) also makes a related point.

semantics and meta-semantics, but to my mind its real value lies in seeing it as a theory of *meta-meta*-semantics. Of course, it's a verbal issue what we *call* 'deflationism' and I'm happy to give Field the term. The important point is that his theory of meaning attribution is a model of anti-realism: it clearly illustrates how we might come to confer normative significance on a particular word-world relation.

I said that 'cookie' is a good translation of 'biscuit' because it gets you a treat. This might give the impression that a translation is good just insofar as it fulfills contingent desires. That's one view, but my point here stands so long as translation is a matter of serving our "interests" in the broad sense mentioned in section 3. Thus, being a good translation might have more to do with suiting aspects of our history or culture, or perhaps fitting well with our biological or psychological make-up. So long as *something* about us makes a translation good, and thereby confers significance on a particular word-world relation, anti-realism follows.

Still, as mentioned in section 3 the anti-realist should not take her own view too seriously. If one translation serves our interests, another will serve our *quinterests*. The former tracks the relation of reference, while the latter will track other relations like *queference*. For the anti-realist, there's nothing objectively special about interests over *quinterests*, so in reality our interests confer no more significance on reference than our *quinterests* confer on *queference*. We say that reference is special because of our interests, but that's just because of the language we already speak. Again, the picture is ultimately that *we do what we do*.

I've discussed deflationism about linguistic meaning, but one can extend the idea to *mental content* too. Here is one simple proposal: for me to attribute an agent X the belief (desire) *that p* is to say that X accepts (wants) a sentence 'S' which means that p; that is, a sentence that's well translated as 'p' in my idiolect. On this view, the propositional attitudes of belief and desire are explained in terms of attitudes to sentences, of 'accepting' and 'wanting' respectively, so the deflationist now owes an account of the latter. Presumably they are functionally specifiable, but I won't try to say how. Everyone owes some (presumably functional) account of belief and desire, so the debt here is not unique to the deflationist. In any case, if good translation tracks some word-world relation then by extension mental content is ultimately fixed by something involving that relation too and we can call this inflationism about content if we want. The point is that it's an explicitly anti-realist theory on which that relation has the normative significance of determining correct thought only because of our interests.¹⁹

6. Pragmatism I: Inquiry

So far I've distinguished realism from anti-realism about meaning and content. I'll now argue that anti-realism leads to the four pragmatist theses we began with. Since anti-realism is a claim of meta-meta-semantics, my main contention is that the pragmatist theses follow *when read as claims of meta-meta-semantics also*. This may or may not be how theorists in the pragmatist

¹⁹ Is there a viscous circularity in this account? I said that attributions of mental content depend on good translation—but doesn't good translation depend on mental content, on what I believe and desire? Not necessarily. As emphasized in the text, good translation is fixed by *something* about us, but that needn't be our contingent desires or other mental states. And even if it does involve our desires, we need not run in a circle. First come the sentences I accept and want. If I accept 'p', then since it's trivial that 'p' means that p, it follows that I believe that p. We've now fixed the contents of *my* beliefs, and likewise for my desires. My beliefs and desires then fix which translation into your idiolect is best, and that together with your attitudes to your own sentences fixes what *you* believe and desire—or at least, it fixes my attributions of them.

tradition always intended them to be read, but put that aside: my aim is just to show that there's *an* interesting reading on which they do follow.

Start with the pragmatist conception of inquiry:

Pragmatic theory of inquiry: The aim of inquiry is to serve our interests. A theory is a tool that serves our interests; our best theories are those that best serve those interests.

This contrasts with the “truther” theory on which the aim of inquiry is not to serve interests but to uncover truth, hence our best theories are those that best reflect the truth.

The issue is the “aim of inquiry”, but this phrase bears clarification. When the truther says that the aim is truth this could mean two things. It could be a “global” claim about the enterprise of inquiry writ large, namely that its end goal is to uncover all and only true propositions (or perhaps just those that are relevant in some sense). Or it could be a “local” claim about a particular attitude involved in inquiry, such as belief, to the effect that truth is its primary standard of evaluation. For example, Velleman (2000) argues that while believing that *p*, assuming that *p*, and imagining that *p* all involve “regarding *p* as true”, only belief involves an aim to do so only if *p* is true. On that view, truth is the primary standard of evaluation of a belief, for it is only true beliefs that fulfil their aim.²⁰

The global and local issues may be linked, but I'll focus on the latter. For the sake of neutrality let's call the main belief-like attitude involved in inquiry “endorsement”. Then what the pragmatic theory of inquiry says is that endorsing a theory involves an aim to endorse something that serves our interests; hence the primary standard of evaluation is whether it serves them. If like Velleman you think that belief aims at truth it follows that endorsement isn't belief; if instead you think endorsement is belief it follows that belief doesn't aim at truth.

This theory of inquiry is found in the classical American pragmatists. Charles Peirce (1955a) said that inquiry is a process that removes doubt and replaces it with “settled belief”. And doubt is worth removing, he says, because it paralyses action. Thus, when we settle belief we do so with the aim of believing something that serves action. And William James said that a theory is ‘an instrument: it is designed to achieve a purpose’ (1975, p. 33); thus, endorsing a theory involves an aim to endorse something that works.

I'll argue that anti-realism about meaning leads to this pragmatist theory of inquiry. To this end, note that I've assumed all along that the aim of inquiry is *correctness*. For I said back in section 1 that in calculating ‘ $68 + 57 = ?$ ’ one aims to give the correct answer, and this (we now see) was a local claim about the aim of inquiry. But this assumption was just a framework, not a specific theory. For I said schematically that a sentence *S* is correct if and only if (i) *S* is correct if and only if *p*, and (ii) *p*; and this leaves open how to understand the scheme. It leaves open what fixes correctness-conditions, and in particular whether it plays this normative role independently of us or because of us. This is the issue of realism vs anti-realism. What I'll argue is that the aim of inquiry hangs on this issue: anti-realism, but not realism, leads to the pragmatic theory that it aims at serving interests.

²⁰ The distinction between the global and local claim corresponds to Lynch's (2009) distinction between the “goal of inquiry” and the “norm of belief”.

This might be surprising, for I emphasized above that realism and anti-realism can both agree that correctness is truth. That is, both agree that some (perhaps wildly disjunctive) mixture of usage and the world fixes correctness-conditions, and both can agree to *call* these ‘truth-conditions’ if they like. So it follows on both views that a sentence is correct if and only if it is true; hence, for inquiry to aim at correctness is for it to aim at *truth*. But isn’t this just the truther theory that pragmatism is supposed to *reject*?

Not necessarily. It depends on the meta-meta-semantic question of *why* this property we call ‘truth’ constitutes correctness and hence is the aim of inquiry. For the realist, it has this normative significance independently of us and our interests. The aim of inquiry is therefore set independently of us: there is this property out there and the aim is to uncover *it* regardless of whether it serves our interests. But for the anti-realist, the property has no normative significance independently of us. It constitutes correctness, and hence is the aim of inquiry, only because it suits our interests to organize our practices around it. Thus, it’s our interests that really pull the strings in setting the aim of inquiry, just as the pragmatic theory states!

The point here is subtle. It may be that the property at which inquiry aims is specifiable independently of us—it could in principle be the Lewisian property fixed by use plus naturalness outlined in section 2. So the claim is not that our interests enter into the essential nature of the property at which inquiry aims. And as I said, if we *call* this property ‘truth’ we can *say* that inquiry aims at truth if we want. Still, what makes this a “pragmatic theory of inquiry” is the *meta-meta-semantic* claim about *why* that property is the proper aim, that it’s only because it serves our interests to evaluate inquiry with respect to it.

You might object that this reading of the pragmatic theory is too much of a stretch: anyone who *says* that inquiry aims at truth must be far removed from the concerns of the classic American pragmatists, regardless of what meta-meta-semantic claim they dress it with. I disagree, for three reasons. First, William James said that inquiry aims at what works. He also said that truth is what works, so it follows on his view that inquiry aims at truth. But James was no truther! Insofar as he was a pragmatist, we must interpret him *not* as denying that truth is the aim but as saying something about *why* it’s the aim. This is precisely what my reading of the pragmatic theory does.

Second, my reading leads to a kind of relativism associated with the pragmatist tradition. Our interests are served a property we call ‘truth’, but another group of organisms may have different interests served by a different property, *quuth*. Both groups aim for correctness in the broad sense, but for us this means seeking theories that are *true* while for them it means seeking theories that are *quue*. Neither group has the “better” aim: both groups seek the property appropriate to their respective interests and there is nothing further to break the symmetry. Thus, what counts as successful inquiry—what counts as “the final correct theory of everything”—differs for each group precisely because of their differing interests.²¹

Third, recall that what the anti-realist calls ‘truth’ may be highly gerrymandered. Mathematical, moral, and gastronomic vocabulary may have very different functions, each used for a different purpose, so that what makes for correct speech differs across those domains. I

²¹ Though this isn’t relativism in the semantic sense. The property we call ‘truth’ may be a monadic property that itself contains no relativization to us or our interests—like I said, it could be the Lewisian property fixed by use plus naturalness. What’s relative to a group is which (possibly monadic) property is normatively significant.

granted that the anti-realist can disjoin these correctness-making features and call the disjunction *truth*. But remember that the disjuncts may be a scattered miscellany, unified only insofar as each serves one of *our linguistic purposes*. Clearly, on this picture it is the *purposes* that come first: the aim of mathematical and moral inquiry is primarily to serve *them*; only in a secondary sense is the aim to uncover this disjunctive miscellany we call truth.

Realism leads to none of this. On that view interests play no role in setting the aim of inquiry; one property is special regardless of anyone's interests, and the aim of inquiry is to uncover *it*. Thus, there's no relativism: one of the above groups have the "wrong" aim insofar as they don't seek the special property. This is how I read the truther view that the aim of inquiry is truth, though it's better expressed as the view that truth is the *ultimate* aim. For the pragmatist can agree that the aim is truth—her point is that it's not the ultimate aim. The ultimate aim is serving our interests; the property she calls 'truth' is just a proximate means to that aim.

Truthers sometimes complain that pragmatists miss the obvious fact that inquiry aims at "getting the world right". Rorty invites this criticism in his more colorful moments, when he suggests that inquiry aims at "social solidarity".²² But one cannot criticize the pragmatism I've outlined on these grounds. For I started with the assumption that inquiry aims at correctness, and (schematically) a sentence is correct iff it means that *p*, and *p*. Thus, I've assumed all along that inquiry aims to "get the world right" in at least this minimal respect! To be sure, this assumption can be denied. Van Fraassen's "constructive empiricism" claims that science aims at empirical adequacy, and a theory can be empirically adequate yet not correct in my sense. But my goal was not to show that constructive empiricism is wrong. Rather, my goal was to start with an assumption congenial to the truther, namely that inquiry aims at "getting the world right" in *at least* the minimal respect of being correct, and show that without realism about meaning even that starting point collapses into pragmatism.

We saw earlier that Field's deflationism is a model of anti-realism. It's worth returning to that model, for it shows that there are in fact *two* respects in which the aim of inquiry can depend on interests. Above I characterized correctness with a scheme, and for the deflationist its substitution instances are sentences '*p*' in one's idiolect. Thus, correctness is defined by a potentially infinite disjunction:

(D) 'S' is correct if and only if 'S' means that snow is white and snow is white, or 'S' means that grass is green and grass is green, or...

with a disjunct for every sentence of her idiolect. Then one source of interest-dependence is obviously the deflationary account of 'means that'. For on that view, (D) amounts to this:

(D*) 'S' is correct if and only if 'S' is well translated as 'snow is white' in my idiolect and snow is white, or 'S' is well translated as 'grass is green' in my idiolect and grass is green, or...

Since good translation depends on interests, correctness depends on interests also.

The second respect of interest-dependence concerns the scope of the correctness scheme. We naturally assumed that it ranges over *every* sentence in one's idiolect, and hence that (D*) contains a disjunct for every such sentence. But now imagine deleting from (D*) a disjunct

²² See Rorty (1979); see Misak (2013, chapter 13) for a discussion of this objection to Rorty.

corresponding to an astronomically large sentence I'll never entertain. This defines a distinct notion of correctness—call it correctness(-)—and hence a distinct possible aim of inquiry. One might aim for correctness or correctness(-), and according to anti-realism there is, independently of us, no fact about which is more worthy than the other.

More generally, there's a distinct aim for *any* set of sentences in one's idiolect, each defined *via* the same correctness scheme we started with and differing only in which instances we include in its scope. For the anti-realist, which aim we should take up depends on us; on our needs, projects, capacities, and the like. Not so for the realist. On her view, a sentence is correct when it has the special property with independent significance, and inquiry should aim at *that* property. The different aims one gets by varying the scope of the correctness scheme are (at least for the most part) simply the *wrong* aim, regardless of whether they suit us better.

7. Pragmatism II: Immanent truth

Turn now to the pragmatist idea that truth is immanent in the sense that it cannot outrun our practices of inquiry. As I put it,

Immanent theory of truth: Truth is nothing other than that which inquiry would converge on (in the long run). It's therefore impossible for an ideal theory—a theory as confirmed as can be—to be false, and (vice-versa) impossible for there to be truths that inquiry could never uncover.

This idea is found in classical pragmatists such as Peirce, who said that 'the opinion which is fated to be ultimately agreed by all who investigate, is what we mean by the truth' (Peirce 1955b, p. 39). This is perhaps the most striking of our pragmatist theses, for I suspect that many philosophers will think it *obvious* that truth transcends us and our practices; that we may fail to uncover it no matter how long or hard we try.

Does anti-realism lead to the immanent conception? It's hard to see why it would. Sure, the anti-realist thinks that inquiry aims to serve *our* interests, but it doesn't follow that inquiry will be a success. More fully, her view is that our interests are best served by some property P, so inquiry aims at theories with P. And I've said she can call this property 'truth' if she wants. But couldn't P be a property that inquiry will never converge on?

Nonetheless, I claim that anti-realism does lead to the immanent conception. In short, my argument is that for the anti-realist there's no point in using 'true' to denote anything transcendent. So truth will be immanent *by definition*: it's immanent because 'true' is stipulated to denote that which inquiry would converge on. Indeed, Peirce's quote above suggests that he simply *defined* 'true' to be immanent also. Of course, the mere fact that one *can* define 'true' like this is not interesting—one can use words how one likes. The interesting question is why it makes sense for the anti-realist to use the word like this.

As a warm up, think about what *realism* entails about inquiry in the long run. On this view the aim of inquiry is to uncover the special property with independent significance. Suppose for concreteness that this is the Lewisian property fixed by use plus naturalness. Presumably there's no guarantee that inquiry will converge on this property, even in the long run, in which case that would be a *failure* of sorts precisely because that property was the aim.

It's this last bit about failure that the anti-realist rejects. She'll agree that inquiry isn't guaranteed to converge on the Lewisian property, but on her view this isn't necessarily a *failure* because inquiry doesn't necessarily aim at this property in the first place. Inquiry just aims to serve our interests, so unless we are best served by the Lewisian property there's no sense in which inquiry fails by not converging on it.

Indeed, there are *plenty* of properties that inquiry will never converge on. Consider a property had by all and only those sentences printed in the Lord of the Rings (properties are cheap, so there is some such property). Inquiry won't converge on this "Tolkeinian" property either, but this is *obviously* no failure since it was no aim of inquiry to do so! Inquiry aims to serve our interests, and we aren't served by uncovering the Tolkeinian property; hence convergence on this property is entirely irrelevant to evaluating the success of inquiry. If we aren't served by the Lewisian property either, the anti-realist will regard it as equally irrelevant.

The point is that for the anti-realist, once inquiry has converged on a theory that serves our interests there is *no further standard* against which to evaluate it. If it didn't converge on the Tolkeinian property this is no failure because it isn't the appropriate standard. Thus while the anti-realist could in principle use 'true' to denote the Tolkeinian property, it would make little sense *insofar as 'true' is supposed to denote a standard by which to evaluate inquiry*.

That's just by way of warm up. The original worry was that inquiry might not converge on a theory that serves our interests, and I've said nothing yet that speaks to this. Our interests may be served by some property P, but it doesn't follow that inquiry will converge on P.

Still, if inquiry converges at all it'll converge on *some* property; call it Q. And for the anti-realist, Q has no more and no less normative significance than P. Sure, P may serve our interests and Q may not. But Q will serve something else, our *quinterests*. Remember, when the anti-realist says that P is significant because of our interests, we should not take this too seriously. Our interests may confer significance on P, but our quinterests confer significance on Q and on her view interests have no prior significance over quinterests. So inquiry *has* converged on a property that serves *something* about us, namely our quinterests!

Ultimately, then, wherever inquiry ends up is as good a place as any. We may say "P is special because it serves our interests", but that's only because we already speak the language of interests. We could just as well say "Q is special because it serves our quinterests". Both are special; neither is special. Whatever inquiry converges on serves *something* about us, and for the anti-realist there is no further standard against which to evaluate it. Thus, while the anti-realist could in principle use 'true' to denote something further—something transcendent—it makes little sense *insofar as 'true' is supposed to denote a standard against which to evaluate inquiry*. That is why it makes sense to define 'truth' as immanent, as that which inquiry converges on.

A comparison with sport might help. Cricket is hard to play well. It requires being able to hit a six; that is, hitting a ball so it clears a boundary about 70 yards away without bouncing. Some people can do this better than others. Now, suppose humanity sets about developing the best cricketers it can, converging in the long run on "ideal" athletes that are as good at this as humanly possible. Still, they wouldn't be able to hit a ball 70 *miles*! So even if they were perfect at cricket they'd be terrible at *quicket*, a game like cricket but with boundaries set miles out. Have we thereby "failed"? Not at all, for there's nothing special about quicket over cricket

independently of *us*, our history of sport, our body type, and so on. They are both games out there in the space of logically possible games. To think that our ideal athletes are “failures” in any serious sense is to presume that quicket is a “natural joint” in the space of games, which is obviously absurd.

And what if they’re not perfect at cricket? Well, they’ll be perfect at *some* game with slightly shorter boundaries, call it *clicket*. But cricket is no more a “natural joint” in the space of games than quicket; hence it makes no more sense to evaluate our athletes with respect to cricket than quicket. Thus, if they’re not perfect at cricket the moral would not be that we’ve *failed* in any interesting sense, but that cricket was a silly aim all along; that *clicket*, not cricket, is in fact the game that suits *us* best.

Sport is therefore an immanent activity: once we converge on the “ideal” athlete there’s no important sense in which we’ll have failed. For the anti-realist, inquiry is immanent in the same sense. Once humanity converges on an ‘ideal’ theory, a theory as well confirmed as can be, there is likewise no important sense in which we’ll have failed. Sure, there are many properties out there we won’t have converged on, such as the Tolkeinian property. But they have no significance, independently of *us*, over the property we *did* converge on, so it makes no sense to evaluate our theory against them. So, insofar as ‘true’ is to denote a standard of evaluation, it must denote that which inquiry converges on.

We thereby arrive at the immanent conception of truth. Commentators sometimes describe it as an “epistemic” theory of truth, on which truth is analyzed in epistemic terms like inquiry of confirmation. But as I read it, that’s misleading. The property we converge on may be specifiable non-epistemically—indeed it could turn out to be the Lewisian property fixed by use plus naturalness. If so, an analysis of truth will identify it as *that* property. The immanent theory is rather a theory of *meta-meta-semantics*, a theory of why that property matters. Or rather, of why other properties do not matter; of why there is no further standard of evaluation.

I’ve expressed the immanent theory in terms of convergence, but this was unnecessary. We needn’t assume that inquiry converges on *anything*, nor that there is any such thing as its “end”. We may instead note that inquiry isn’t a random walk. We use rules and procedures—“norms”—to help us. There are norms of logic, norms of induction, norms of finding p-values less than 0.05, and so on. The claim that truth is immanent then amounts to this: that in addition to evaluating whether inquiry satisfies these norms there is, ultimately, nothing *further* against which to evaluate it.²³

The realist thinks there is something further: the special property he calls ‘truth’. Even if we satisfy the norms perfectly our theories may lack this property, and since this property is the aim of inquiry this would be a *failure*. This is then a transcendent conception of truth. Indeed, since this property of truth is the ultimate aim of inquiry, our norms earn their keep only insofar as they reliably track it. Herein lies the road to skepticism, for we cannot argue that our most basic norms like modus ponens and induction are reliable without using those very norms. For the

²³ Misak (2009) argues that Peirce never put much weight on “convergence” or “the end of inquiry”, and instead held a view more in line the one just expressed.

realist, then, inquiry is like stumbling through the desert in the dark. If we're lucky we'll keep heading where we want to go, but there's no way to tell from the inside.²⁴

By contrast, for the anti-realist the norms earn their keep insofar as they serve our interests, that being the aim of inquiry. But as before, we shouldn't take this too seriously. It's not that there is this special thing, our *interests*, such that we may now ask whether our norms serve *them*. That would lead back into skepticism. Rather, there will inevitably be *something* about us that our norms help serve, be it our interests or quinterests. So again, the picture is ultimately that we do what we do. At bottom our norms are "justified" thanks to being *ours*. Beyond that there's nothing further that makes our norms special. And this is the immanent conception of truth and inquiry: that in addition to evaluating whether inquiry satisfies these norms there's nothing further against which to evaluate it.²⁵

Note that on this view there may be no cross-domain unity to the norms. If mathematics, ethics, and science all serve different purposes, then the norms found of one domain may look very different from the norms of another. It's not, as the realist thinks, that they're all unified by serving as guides to his special property of *truth*. Rather, they may have nothing in common between them other than that they serve *some* purpose or other. This heterogeneity of norms is emphasized by contemporary pragmatists such as Misak (2009) and Price (2011).

I've argued that it makes sense for an anti-realist to define 'true' as immanent *insofar as it's supposed to denote a standard of evaluation*. But that's not the only function of 'true'. As we saw earlier, deflationists have argued that disquotational truth, true_d , is useful in allowing one to express potentially infinite disjunctions. And disquotational truth is not immanent. Field himself emphasized this point when he said that the deflationist can recognize a long disjunction of the form

"It might be the case that either the number of brontosaurus that ever lived is precisely 75,278 but we will never have reason to believe that; or the amount that Michael Jackson spent on underwear in his lifetime is exactly \$1,078,085.72 but we will never have reason to believe that; or..." (1994, p. 264).

Given disquotational truth and the translational theory of belief, this is equivalent to

It might be the case that either 'The number of brontosaurus that ever lived is precisely 75,278' is true_d but we will never have reason to accept that sentence; or "The amount that Michael Jackson spent on underwear in his lifetime is exactly \$1,078,085.72" is true_d but we will never have reason to accept that sentence; or...

That is,

²⁴ Contemporary epistemology is replete with attempts to disguise this ugly truth. Some emphasize that *if we're lucky* our norms will lead to knowledge, though of course we cannot tell whether we're lucky from the inside. Others re-define "tell" and "inside" so that it follows that we can tell from the inside, only now with different meanings attached to these terms. To my mind this is all just paper over the cracks, but I can't defend this here.

²⁵ Misak (2009) interprets Peirce's immanent theory of truth along roughly these lines. In particular, she argues that Peirce never put much weight on "convergence" or "the end of inquiry", despite some of his remarks to that effect.

It might be the case that there's a sentence S such that S is true_d but we'll never have reason to accept S.

Hence truth_d can outrun rational inquiry in the long run. But this is consistent with what was said above. For an anti-realist will take the same attitude to truth_d as she did to the Tolkeinian property: if inquiry doesn't converge on it, that's no failure. True_d may have its uses, but not as a standard of evaluation.

This corresponds to the point at the end of the last section: that for the deflationist there's a distinct aim of inquiry for *any* set of sentences in their idiolect, each defined *via* the same correctness scheme but differing in which instances are included in its scope. For the deflationist, which aim to take up and use as a standard of evaluation depends on which suits us best. The point here is that it'll be a restricted aim: it'll omit Field's sentence concerning brontosaurus from its scope.

8. Pragmatism III: Representation and meaning

It is now straightforward to show that anti-realism entails the other two pragmatist theses. Consider

Anti-representationalism: Inflationary relations of “representation”, “reference”, and “correspondence” between language (or mind) and world should be rejected as philosophically idle.

If “philosophically idle” means that they have no normative significance independently of us, this just *is* anti-realism!

I said at the beginning that this thesis leads pragmatists to reject theories of truth based on inflationary relations like these, such as the correspondence theory. But the disagreement must be stated carefully, for we're now reading anti-representationalism as a theory of *meta-meta-semantics*. Read like this, the pragmatist need not deny that there are inflationary word-world relations of reference, representation, correspondence; nor that they play a central role in a theory of truth. Her thesis is about *why* they *matter*, why it's *them* that fix correctness and the aim of inquiry. Her view is that they're not special in themselves; they matter only because they suit our interests. What she regards as philosophically idle, then, are these relations *absent any claim of where their normative significance comes from*.

Seen like this, the pragmatist can *in a sense* accept the correspondence theory! She could agree that a sentence is true iff it corresponds to a fact. Her point is just that relations are cheap: there is correspondence and quorrespondence, so a sentence will correspond to one fact and quorrespond to another at the same time. Why then does *correspondence* fix correct speech and the aim of inquiry? For the pragmatist, only because it suits us to organize our practices around it. Thus if she rejects the correspondence theory, what she really rejects is that the significance of correspondence is there anyway, independently of us. This may have been William James' criticism of the correspondence theory when he said that “it is not self-evident that the sole business of our mind with realities should be to copy them” (1904, p. 467). As I read him, his point here is not that correspondence (or copying) is objectionably obscure, though that complaint is often levied against it. Rather, his complaint is that the relation of copying has, in itself, no normative significance.

I also take this to be the core idea behind more contemporary anti-representationalism. For example, Huw Price writes that

“representation... is a theoretical category that we should dispense with altogether. The right thing to do, as theorists, is not to say that it turns out that none of our statements is a genuine representation; it is to stop talking about representation altogether, to abandon the project of theorizing about word-world relations in these terms. It is a bit like the familiar case of simultaneity: the lesson of relativity is not merely that we live in a world in which absolute simultaneity does not make itself manifest, but that we should abandon the notion of absolute simultaneity altogether, for theoretical purposes” (2011, p. 10)

But as he goes on to emphasize, he is *not* saying that “there is nothing to be said about the relation of our words to the natural world” (p. 11)—indeed, he is eager to emphasize that there are interesting word-world relations. His point, as I read him, is that it’s only insofar as they suit our needs that they have any significance, and indeed that different areas of discourses may serve different needs and hence revolve around different relations. This reading makes his comparison to simultaneity particularly fitting. Just as relativity theory allows one to talk of simultaneity *relative to a frame of reference*, so pragmatism allows us to talk of an inflationary word-world relation being normatively significant *relative to an interest*.

Consider finally the thesis about meaning:

Pragmatic theory of meaning: The meaning of a hypothesis is given by its practical consequences.

This so-called “pragmatist maxim” was central to pragmatism from the beginning. According to Peirce, ‘there is no distinction of meaning so fine as to consist in anything but a possible difference in practice’ (1955b, p. 33). And James wrote that ‘if no practical difference whatever can be traced, then the alternatives mean practically the same thing, and all dispute is idle’ (James 1975, p. 28).

Anti-realism about meaning is very much in the same ball-park. This is most clear on Field’s deflationary view that meaning has to do with good translation, which itself is a *practical* matter. But we needn’t use deflationism to make the point. We agreed in section 5 that *some* mixture of usage and the world fixes meaning (i.e. correctness-conditions); what anti-realist says is that that mixture plays this role only because it suits our *interests*. Here the focus is on “interests” rather “practical consequences”, but any difference here is minor: the point remains that what fixes meaning does so only because of its consequences *for us*. This is close enough to the pragmatist maxim for me.

Once again, this thesis emerges as a thesis of *meta-meta-semantics*. The claim is *not* one of semantics—it’s not that the meaning of an expression is identical with some set of practical consequences. Nor is it a claim of meta-semantics—it’s not that practical consequences help fix reference. It’s rather a claim about why the properties and relations mentioned in the semantics and meta-semantics matter; that they matter only because they suit us well.

This is all to the good, for if read as a thesis of semantics or meta-semantics it invites the kinds of objections that felled verificationism. Exactly which practical consequences fix meaning? Aren't there meaningful statements with no practical consequences, such as those of speculative cosmology or the higher reaches of mathematics? Do practical consequences compose in the right way? But when read as a thesis of meta-meta-semantics, these worries are not to the point.

9. Primitive normativity

I've tried to show that anti-realism about meaning leads to these pragmatist theses. Hence those who reject pragmatism—the “truthers”—must endorse realism.

But I haven't argued for anti-realism. I find it immensely plausible, and I hope my articulation of the view in sections 1-5 made you feel similarly. But that's no argument, and I don't intend to give one here. Still, it's worth showing that realism is committed to a certain radical picture of reality—a picture so striking that many would-be truthers will be loath to accept it!

The picture is that some normative facts are *sui generis* in the sense that they are not identical or reducible to natural facts about the cosmos, or even to super-natural facts about God and the like. Normativity is its own, extra component of reality—as real and mind-independent as matter but entirely distinct from it. This is sometimes known as non-naturalism or robust realism about normativity, but I'll call it 'primitivism'.²⁶

Why does realism need primitivism? It's not immediately obvious. A realist makes normative claims about truth and reference: that they fix the *right* way to go on in thought and talk, that inquiry *ought* to aim at them, and so on. And by definition, realism says that these facts hold independently of us; that they don't hold in virtue of anything to do with us. But why must they be primitive? Why couldn't they hold in virtue of something else about the world that's independent of us?

Well, they could. But properties are cheap. If a property P confers significance on truth, another property Q will stand to quoth just as P stands to truth; hence Q would confer equal significance on quoth. So for truth to be special, P must somehow be special too. But if some further property P* confers significance on P, there'd be another property Q* that conferred equal significance on Q. And so on. This is a disaster for the realist. As I tell my kids in darker moments, if everyone's special no one's special. Thus, for truth to be special in the way that the realist needs, whatever ultimately confers significance on truth must have normative significance *all on its own*, not in virtue of anything else. This is primitivism.

At least, that's the short story. There is more to say here, some of which I say in Dasgupta (2018). For example, a realist might say that truth is uniquely special because it's fixed by the Lewisian mix of use plus *naturalness*. After all, naturalness delineates “objective structure” in the world, and surely *that's* special! But as we saw in section 2, there is naturalness and quaturalness. What this appeal to naturalness needs, then, is a normative fact to the effect that

²⁶ “Non-naturalism” is not ideal since, as is often observed, the view in question distinguishes normative facts not only from natural facts but also from non-natural facts about God and spirits. And “robust realism” is too loaded—who wouldn't like a view that's *robust*?

naturalness has significance with regards to thought and talk that quaturalness lacks. If this fact is brute, we have primitivism. And if it isn't, then something *else* makes naturalness significant and we're off to the race described above. The point of Dasgupta (2018) was to make explicit that the appeal to naturalness in much contemporary philosophy requires primitive normativity.²⁷

Don't confuse primitive normativity with primitivism about *content*, the view that properties such as *referring to addition* or *having the content that p* are primitive properties of linguistic signs or mental states, not reducible to some mixture of use and world. Boghossian (1989) argues that the moral of Kripkenstein's question of how to go on is that we must reject 'the reducibility of content properties to naturalistic properties' (p. 540), and it's easy to read him here as defending content primitivism. But this doesn't help, for all we have now is an extra, primitive property of '+'—it's "content"—and it doesn't follow that it has any normative significance. For all we've been told it's just another property out there; nothing follows about whether *it* determines the right way to go on using '+'. What Boghossian needs is primitive *normativity*, a fact to the effect that content is normatively significant. But then his primitive content properties are redundant: the primitive normativity could have attached to some mixture of use and the world that was there anyway. (To be fair, perhaps by primitive content Boghossian meant what I mean by primitive normativity. If so, fine—the important point is that these views are distinct and that content primitivism, on its own, does nothing to address the Kripkensteinian question of how to go on.)

The upshot is that realism about meaning needs primitive normativity. Thus, anyone who rejects primitive normativity must embrace anti-realism and the pragmatism it leads to! This includes normative naturalists, who reduce normativity to facts about the natural cosmos. It includes constructivists, who reduce normativity to human natures or preferences. It includes relativists, for whom normativity varies from culture to culture. And so on. That is why I said at the beginning that my path to pragmatism proceeds from a claim that many contemporary truthers would accept. For it is only a minority of truthers, the normative primitivists, who are in a position to deny it. The rest must renounce the truther movement.

10. Conclusion

Having clarified that realism needs primitive normativity, I won't try to refute it. But let me end by indicating why I find it incredible. In brief, even if there were facts of primitive normativity I see no reason to *care* about them—indeed it strikes me as pathological to pay it any regard. This objection to primitive normativity is well known; see Dreier (2015) for an overview. I developed the objection in a particular style (Dasgupta 2017) and applied it to norms of metaphysics in Dasgupta (2018). But it applies equally to norms of language and thought.

To see why, consider the word-world relation that our linguistic practices have organized around. We call it 'reference'. It holds between "birds" and birds, "pigs" and pigs, and so on. It goes along with a property we call 'truth': 'Birds fly' is true if and only if birds fly; 'Pigs grunt' is true if and only if pigs grunt; and so on. Now suppose millenia of scientific research converges on an ideal theory, a theory as well confirmed as can be by our norms: it's explanatory, simple,

²⁷ Alternatively, a realist might try securing significance by definition: perhaps it's part of what it *means* for an action to count as an assertion that it aims at truth; that's why truth fixes the right way to go on. But as I said in section 3 this is just anti-realism, for truth is significant only insofar as we assert rather than quassert.

illuminating, whatever. And suppose it's even *true*. We naturally celebrate our feat. But according to realism, our celebration is premature. For on that view it's possible that reference and truth aren't in fact primitively normative; that another property and relation, "Real Truth and Reference", have this property of primitive normativity instead. Thus our ideal theory could be true, but Really False. Indeed, it could be that 'pigs' refers to pigs but Really Refers to birds, in which case 'Pigs fly' is false but Really True!

What of it? One objection here is epistemic, that realism makes coherent a strange skeptical scenario in which everything we think is true *is indeed true* but Really False. Button (2013) discusses this, but my objection is different. It's that *this* kind of "error" shouldn't bother us in the slightest. Imagine an oracle informs us that our ideal theory is true but Really False. It tells us that "Pigs fly" is false (obviously) but Really True. Why should we care? Our theory is ideal by our standards and does everything we want of it—it's even *true*! All we've been told is that it lacks some extra primitive property, but so what? That's a minor curiosity at most. No one should seriously suggest that we scrap 500 years of post-enlightenment science and develop an entirely new theory on which pigs fly! Thus, if there is such a thing as Real Truth, it's just another property out there along with the Tolkenian property. To strike out in a radical effort to uncover it, insisting that *pigs really do fly*, would be, quite literally, madness.

References

- Alston, William (2005). *Beyond Justification*. Ithica, NY: Cornell University Press.
- Boghossian, Paul A. (1989). "The Rule-Following Considerations". *Mind* 98(392): 507-549.
- Button, Timothy (2013). *The Limits of Realism*. Oxford: OUP.
- Dasgupta, Shamik (2018). "Realism and the Absence of Value". *The Philosophical Review* 127(3): 279-322.
- Dreier, Jamie (2015). "Can Reasons Fundamentalism Answer the Normative Question?" In Gunnar Bjornsson, Caj Strandberg, Ragnar Francen Olinder, John Eriksson, and Fredrik Bjorklund (eds.), *Motivational Internalism*, pp. 167-81. Oxford: OUP.
- Enoch, David (2006). "Agency Shmagency: Why Normativity Won't Come From What is Constitutive in Action". *Philosophical Review* 155: 169-98.
- James, William (1975). *Pragmatism: A New Name for Some Old Ways of Thinking*. Cambridge: Harvard University Press. Originally published 1907.
- Lewis, David (1983). "New Work for a Theory of Universals". *Australasian Journal of Philosophy* 61(4): 343-377.
- Lewis, David (1984). "Putnam's Paradox". *Australasian Journal of Philosophy* 62(3): 221-236.
- Field, Hartry H. (1994). "Deflationist Views of Meaning and Content". *Mind* 103: 249-285.
- Field, Hartry H. (2001). "Attributions of Meaning and Content". In *Truth and the Absence of Fact*, pp. 157-175. Oxford: OUP.

- Horwich, Paul. (1995). "Meaning, Use, and Truth". *Mind* 104(414): 355-368.
- Goldman, Alvin (1999). *Knowledge in a Social World*. Oxford: OUP.
- Grimm, Stephen (2009). "Epistemic Normativity". In *Epistemic Value*, edited by A. Haddock, A. Millar, and D. Pritchard.
- Kripke, Saul A. (1982). *Wittgenstein on Rules and Private Language*. Cambridge MA: Harvard University Press.
- Loewer, Barry (1993). "The Value of Truth". *Philosophical Issues* 4: 265-280.
- Lynch, Michael P. (2009). "Values of Truth and Truth of Values". In *Epistemic Value*, edited by A. Haddock, A. Millar, and D. Pritchard.
- Merino-Rajme (2015). "Why Lewis' Appeal to Natural Properties Fails to Kripke's Rule-Following Paradox". *Philosophical Studies* 172: 163-175.
- Miller, Alexander (2001). "Introduction". In *Rule-Following and Meaning*, edited by A. Miller and C. Wright, pp. 1-16. Acumen.
- Misak, Cheryl (2009). "Pragmatism and Deflationism". In *New Pragmatists*, edited by C. Misak.
- Misak, Cheryl (2013). *The American Pragmatists*. Oxford: OUP.
- Peirce, Charles (1955a). "The Fixation of Belief". In *Philosophical Writings of Peirce*. New York: Dover. Originally published in 1877, *Popular Science Monthly* 12(1): 1-15.
- Peirce, Charles (1955b). "How To Make Our Ideas Clear". In *Philosophical Writings of Peirce*. New York: Dover. Originally published in 1878, *Popular Science Monthly* 12: 286-302.
- Price, Huw (2011). *Naturalism Without Mirrors*. Oxford: OUP.
- Putnam, Hilary (1977). "Realism and Reason". *Proceedings and Addresses of the American Philosophical Association* 50(6): 483-498.
- Rorty, Richard (1979). *Philosophy and the Mirror of Nature*. Princeton: Princeton University Press.
- Sider, Theodore (2011). *Writing the Book of the World*. Oxford: OUP.
- Stich, Stephen (1990). *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation*. Cambridge, MA: The MIT Press.
- Schwarz, Wolfgang (2014). "Against Magnetism". *Australasian Journal of Philosophy* 92: 17-36.
- Velleman, David (2000). "On The Aim of Belief". In *The Possibility of Practical Reason*, pp. 244-281.

Weatherson, Brian (2012). "The Role of Naturalness in Lewis' Theory of Meaning". *Journal for the History of Analytic Philosophy* 1(10): 1-19.

Williams, J. Robert G. (2015). "Lewis on Reference and Eligibility". In *Companion to David Lewis*, edited by B. Loewer and J. Schaffer.

Wittgenstein, Ludwig (1953). *Philosophical Investigations*. Oxford: Blackwell.

Wright, Crispin (2001). "Critical Notice of Colin McGinn's *Wittgenstein on Meaning*". In *Rule-Following and Meaning*, edited by A. Miller and C. Wright, pp. 108-128. Acumen.