

# The inferentialist guide to scientific realism

Tushar Menon

---

## ABSTRACT

Scientific realism is appealing in part because it captures the following naturalistic impulse: claims about what certain perceptually-inaccessible domains of the world are like should be determined our best physics. In this paper, I argue that this naturalistic impulse is in tension with the standard representationalist metasemantic foundations of scientific realism. To better capture the impulse, we should instead adopt an inferentialist metasemantics. Using this metasemantics, I develop a view that I call ‘inferential scientific vindicationism’. I argue that this view fares better than scientific realism, structural realism and constructive empiricism with respect to the naturalist impulse.

---

## 1 Introduction

Scientific realism was sold to me as an expression of an attractive sort of naturalistic impulse: if our inquiry into the world should take seriously the results of some practice, then that practice should be scientific practice. In particular, our inquiry into what certain perceptually-inaccessible domains of the world are like should take seriously the practice of our best physics that models those domains. This strikes me as a good impulse. I want to hold on to it.

In this paper, I argue that the impulse is hindered by the standard metasemantic foundations of scientific realism, according to which the meanings of linguistic expressions are grounded in some representation relation that they stand in to some part of the world. Call this view ‘representationalism’. My negative thesis is that we should drop representationalism. My positive thesis is that if we want to do justice to the naturalist impulse, then we should instead adopt an inferentialist metasemantics, according to which the meanings of linguistic expressions are grounded in the manner in which their use is guided by inferential norms. I call this view ‘inferential scientific vindicationism’,<sup>1</sup> of which standard scientific realism turns out to be a special case.

I am aware that there is a vast amount of philosophical real estate between the denial of representationalism, and the acceptance of inferentialism, and the latter might strike some readers as an extreme view. The argument presented in this paper, which takes us from anti-representationalism to inferentialism, is cumulative: each section introduces a subspecies of the

<sup>1</sup> The choice of terminology is a nod to (Feigl [1950]).

preceding view, characterised by decision points at which we choose between some competing additional commitments. I will highlight these decision points explicitly as we proceed. The payoff is that inferential scientific vindicationism satisfies three desiderata that follow from the naturalistic impulse introduced above, but which form an inconsistent triad for scientific realists, structural realists and constructive empiricists.

Arthur Fine captures the naturalistic appeal of scientific realism thus:

[I]t is possible to accept the evidence of one's senses and to accept *in the same way*, the confirmed results of science only for a realist; hence I should be one (and so should you!). ([1984], p. 35)

The phrase 'in the same way' is doing some heavy lifting here. I take it as attempting to capture the intuition that as a realist, I accept that if my senses are to be trusted, then there is a directly perceptually-accessible entity, a computer, which is described using ordinary empirical discourse (OED), and which I am currently manipulating. Similarly, if QFT is to be trusted, then as a scientific realist, I accept that there are directly perceptually inaccessible entities, electrons, that behave in the way that QFT claims they do.

Fine's phrase captures three sorts of uniformity. The first is metaphysical: to the extent that I accept a mind-independent existence of chairs, I accept a mind-independent existence of electrons. The second is semantic: 'chair' means to chair in the same way that 'electron' means electron. I use reference to imbue 'chair' with meaning, so I should do the same with 'electron'. The third is epistemic: I accept perceptual evidence about chairs by ascribing to chairs the properties that my senses say that I should, so I should accept scientific evidence about electrons, and ascribe to them the properties that QFT says that I should. This suggests that the naturalistic impulse can be usefully regimented as a 'uniformity thesis' with three conjuncts:

1. Metaphysical uniformity: There exists a mind-independent world, and both the observable and unobservable entities that populate this world have the same metaphysical status *viz.* their mind-independence.
2. Semantic uniformity: There should be no difference in the semantic treatment of observational and theoretical terms.
3. Epistemic uniformity: The epistemic status of (some) claims about unobservables is the same as (or approaches, in a sense to be made precise) the epistemic status of (some) claims about observables.

Two questions emerge from semantic uniformity, understood referentially: (i) what makes it the case that tokens of 'electron' successfully refer? and (ii) how should we characterise those referents? The options that have received the most air-time have traditionally understood the problem as one of 'gluing': we have a label, for example 'electron', now how do we figure

out which part(s) of the world it is appropriate to glue this label on to? Understood in terms of gluing, these two semantic questions pose a problem that is strikingly similar to what (Jackson [1998]) raises as a ‘location problem’ for physicalists: where, in the physical world do we locate the truthmakers—elements of the world whose existence grounds (or explains or necessitates) the truth—of moral, mental, modal or mathematical claims? Where, in other words, are the worldly entities onto which we stick the labels from these problematic areas of discourse?

When we recast the semantic questions as a location problem we see just how deeply entrenched the metasemantic assumption of representationalism is in the literature on scientific realism. Roughly speaking, representationalism is the (family of) view(s) according to which the meaning of a linguistic expression is grounded in the relation of representation that obtains between an expression and the proper part of the world that that expression picks out. ‘Referentialism’ is the special case of representationalism where the relation is reference, and the relata are, respectively, terms and parts of the world. It is the default semantic position for scientific realism, although scientific realists generally embrace a number of further semantic commitments. In §3, invoking arguments first articulated by Price ([2011b]), I argue that representationalism (*a fortiori* referentialism) should not be the default metasemantic underpinning of the semantic component of a naturalistically-motivated scientific realism.

You might worry that dropping representationalism undermines the subsequent position’s claim to being objective, and about (the perceptually inaccessible domain of) an external world. In §4, I argue that although taking science seriously, in the appropriate semantic sense, appears to require referentialism, this appearance is illusory. A Pricean subject-naturalism, which is a form of pragmatism, is a far better alternative. On this view, metasemantics is to be determined by, and answerable to, pragmatics. §5 offers an account of precisely how this determination works, by sketching a normative approach to grounding metasemantics in pragmatics. By dropping representationalism, the pragmatist appears to lose the representationalist’s ability to treat theoretical and observational terms on the same semantic footing. So subject-naturalists need an alternative account of semantic uniformity. This is provided by the final component of the view: inferentialist metasemantics, the topic of discussion in §6.

In §7, I put it all together, and present an account of inferential scientific vindicationism. The formal characterisation of a scientific theory makes explicit what is implicit in the discursive practices of the relevant scientific community. Attention to the inferential richness of these practices is what serves to characterise the contexts in which our best physics practices are to be understood as committed to the existence of unobservables. The focus on semantic uniformity generalises beyond propositionally-articulated theories to the highly mathematised explicitations of modern physics, whereas talk of theoretical and observational terms does not (for reasons of space, I will not say any more about this aspect of the view). This is why I claim that the traditional scientific realism debate is a special case of the broader problem of inferential scientific vindication. The discursive practices of scientists determine what their terms and other expressions mean, and consequently what our best science commits us to. As a result, the inferential scientific vindicationist can circumscribe the contexts in which our discursive prac-

tices support scientific realism, and those in which they do not. Finally, in §8, I argue that the naturalistic impulse that we began with is better served by inferential scientific vindicationism than it is by scientific realism, structural realism and constructive empiricism.

## 2 Scientific realism and anti-realism

Van Fraassen distils from Sellars, Boyd, and Putnam, a characterisation of scientific realism: ‘Science aims to give us, in its theories, a literally true story of what the world is like; and acceptance of a scientific theory involves the belief that it is true’ ([1980], p. 8). More contemporary discussions, following Psillos ([2005]), tacitly accept van Fraassen’s claim about the scientific realist’s understanding of the aim of science, and explicitly identify three constitutive aspects of scientific realism:

Metaphysical realism: There exists a mind-independent world, in the sense that entities posited by science exist and ‘their existence is independent of our knowledge and minds’ (Ladyman [2012], p. 158).

Semantic realism: ‘[i]Theoretical assertions are not reducible to claims about the behaviour of observables, nor are they merely instrumental devices for establishing connections between observables. [ii] The theoretical terms featuring in theories have putative factual reference. [iii] So, if scientific theories are true, the unobservable entities they posit populate the world’ (Psillos [2005], p. xvii).

Epistemic realism: ‘Theoretical claims [of science] give us knowledge of the world. (Chakravartty [2007], p. 9). (In addition, some of this is knowledge of a perceptually inaccessible domain of the world).

If satisfaction of the uniformity thesis is a desideratum of a naturalistic account of scientific theory interpretation, then it is easy to see why scientific realism is appealing for the naturalist: it is just a special case of the uniformity thesis. Metaphysical and epistemic uniformity follow from metaphysical and epistemic realism. There is a subtlety in the semantic component that is worth highlighting. Psillos’ statement of semantic realism is equivalent to a strong version of semantic uniformity introduced above, call it the ‘strong semantic uniformity thesis’: there is to be no difference in the semantic treatment of observational and theoretical terms, *and* that semantic treatment is referential. Rephrasing things in this way suggests a useful generalisation, one which embraces only the first conjunct. Call it the ‘weak semantic uniformity thesis’: there is to be no difference in the semantic treatment of observational and theoretical terms. This is the semantic component of what I called the uniformity thesis. Fine’s ‘in the same way’, should be understood as mandating semantic uniformity without also mandating reference.

In its most straightforward sense, where Fine’s ‘in the same way’ is understood to mean a referential semantics for both theoretical and observational terms, scientific realism comes under immense pressure from underdetermination and theory-change arguments, such as Laudan’s

([1981]) pessimistic meta-induction (PMI): successive successful scientific theories since at least the Scientific Revolution have made incompatible referential claims about what there is in the world. So (unless we are very epistemically lucky) we should not interpret these theories realistically, given their susceptibility to being replaced by a more successful theory.

The semantic realist thus has a dilemma: either accept that the appropriate semantics for all aspects of scientific theories is referential, but then they fall prey to arguments from theory change, or deny it, but then they undermine the scientific realist's (and naturalist's) central semantic uniformity claim. Scientific realists generally attempt to extricate themselves from this bind by arguing that one can accept deny an indiscriminate referentialism about all scientific terms without undermining semantic realism (see e.g. (Kitcher [1993]; Psillos [2005])). They believe that it is possible to modify a referential semantics just enough to avoid the theory-change argument but not so much as to slide into a semantic anti-realism. However, the literature on realism suggests, even if it does not conclusively establish, that there is simply not enough space between the two extremes to carve out a stable and attractive position (see e.g. (Ladyman et al. [2007])). So if these criticisms are correct, as I take them to be, then the scientific realist does not have the resources to satisfy the uniformity thesis; the theory change argument severely undermines both epistemic and semantic realism.

In what follows, I suggest a novel way out of the dilemma. My suggestion is to understand the 'in the same way' clause in the scientific realist impulse not as mandating the use of the same semantics for science as for everyday language, but instead as mandating the use of the same *metasemantics*—inferentialism— across both scientific and everyday language. What we end up with is exactly what we wanted: a version of semantic uniformity which does not rely on reference. This suggestion is not an ad-hoc patch for the semantic realist. Instead, as I detail in the rest of this paper, the view arises from broader considerations of the sort of naturalism that leads to positions like scientific realism, structural realism, and constructive empiricism in the first place.

### 3 Locating unobservables

Both the scientific realist and anti-realist ultimately need to vindicate the use of talk of unobservables (call this 'u-talk') in scientific discourse. Vindication, in this context, is just a demonstration that such talk is appropriate in some specified context; more on this in §5. We can rephrase the animating question as follows: is our scientific theoretical vocabulary malfunctioning when it purports to refer to unobservables like electrons? The realist says that it is not malfunctioning, and that such claims can be vindicated referentially. But they are then saddled with the burden of undermining arguments from theory-change. The anti-realist, on the other hand, says it is malfunctioning. But as a result, they are saddled with the burden of vindicating accepted scientific claims like 'electrons are negatively charged', without recourse to reference.

This impasse is the traditional starting-point for the debate. In this section, following ar-

guments from (Price [2004]; Price et al. [2013]), I argue this is the wrong starting point: it builds in the significant presupposition of referentialism which, although not obviously false, does need to be highlighted and subjected to philosophical analysis. Both sides take for granted that, if u-talk is to be vindicated, then one needs to demonstrate that the referents of troublesome terms ('electron') are, in fact, the referents of some tractable terms. The scientific realist does this by arguing that 'electron' is simply not a troublesome term, and has an unobservable referent that is in perfectly good philosophical standing. The strict empiricist does this by shifting the referent of 'electron' to something constructed out of the referents of 'voltmeter', 'ammeter', 'wire', and so on.

In both cases, the domestication of the troublesome vocabulary occurs at the level of the objects referred-to, or, more broadly, at the level of the worldly truthmakers. Now, this might well be the correct way to think about vindication. It is a mistake, however, to presuppose this without argument. Price highlights that there is no *a priori* reason why vindication of u-talk has to proceed via reference.

The mere cogency of the suggestion of alternative sources of vindication draws attention to two important questions that need to be answered. The first is: what can vindicate u-talk? The standard answer is (successful) reference. This leads the second question: is vindication by reference possible? Consequently, we can characterise four positions, based on how they answer these two questions:

	Vindication by reference is possible	Vindication by reference is impossible
Vindication requires reference	Object-naturalism	Pessimism
Vindication does not require reference	Weak subject-naturalism	Strong subject-naturalism

We can use these characterisations to define the four positions as follows:

Object-naturalism: Successful reference is the only vindicatory strategy for u-talk.

Pessimism: There is no vindicatory strategy for u-talk.


Weak subject-naturalism: Successful reference is not the only vindicatory strategy for u-talk.

Strong subject-naturalism: Successful reference is never a vindicatory strategy for u-talk.


The subject/object-naturalist distinction is due to Price ([2013]). A weak subject-naturalist can believe in reference, as Price himself does ([2010]). What they deny is that the meaning of a linguistic expression is grounded in reference. According to them whatever mediates the meaning of an expression, and thus vindicates u-talk, also sometimes establishes a reference

relation between that expression and some proper subset of the world. But this is epiphenomenal, and the presence of a reference relation in these cases should not be taken as evidence that all, or indeed any, vindication of u-talk is achieved by reference.

At this stage, then, we have a choice between subject and object-naturalism; we choose the former.



We then have the further choice between a weak and strong version of subject-naturalism; once again we choose the former.



#### 4 Vindicating assertions

The referentialism that underpins object-naturalism serves two distinct and independent purposes. The first, as we have seen, is to ground the semantic treatment of linguistic expressions in such a way as to attempt to vindicate u-talk. The second is to provide a uniform account of the semantic treatment of what I might call ‘declarative’ or ‘fact-stating’ discourse, of which u-talk is an example. Call the proponent of the view that all declarative sentences warrant the same sort of semantic treatment a ‘weak semantic uniformitarian’, recalling our use of the term ‘weak semantic uniformity’ in §2.

In denying referentialism, the subject-naturalist is not forced to deny weak semantic uniformitarianism. They can accept that we need a uniform account of how declarative sentences declare, whilst denying that this account is provided by reference. Since the denial of referentialism does not entail the denial of weak semantic uniformitarianism, weak semantic uniformitarianism does not entail referentialism. Of course, we already knew this, given how we set up weak semantic uniformity, but now we can see how Price’s insight about the priority of subject-naturalism is what allows us to sever the link between semantic uniformitarianism and referentialism.

So how do declarative sentences declare? The standard referentialist view is that declarative sentences have components which pick out referents in the world, and ascribe to those referents certain properties and relations. The referents, together with the properties and relations they instantiate then serve as the truthmakers for those sentences. Call the combination of weak semantic uniformitarianism and referentialism ‘strong semantic uniformitarianism’. On

the strong semantic uniformitarianism view, the only manner in which declarative sentences declare is by describing (putative) truthmakers. Truth then becomes a substantive property of a true sentence, bestowed upon it by the truthmakers being a particular way. For the strong semantic uniformitarian scientific realist, ‘electrons are negatively charged’ is made true by the existence of unobservable entities which, according to our best physics, have properties that determine, for example, particular scattering amplitudes.

The subject-naturalist rejects this order of explanation. To be sure, an account needs to be provided of how we vindicate locutions like ‘electrons are negatively charged’, but, crucially, that account need not appeal to worldly truthmakers, and the associated referentialist presuppositions. One important component of the subject-naturalist’s account is what is sometimes referred to as ‘semantic minimalism’. Of particular interest at this stage is the semantic minimalist’s view about truth.

Minimalism about truth is the claim that, in circumstances in which the behaviour of the truth predicate is captured by the Tarskian T-schema, there is nothing more to truth than satisfying this schema.<sup>2</sup> Object-naturalists cannot be semantic minimalists; for them the T-schema is satisfied *because* their ontology of worldly truthmakers bestows upon truth a metaphysically substantial heft. This metaphysical heft is precisely what the semantic minimalist about truth denies. For the purpose of this paper, we need not delve into the specific details of semantically minimalist accounts of truth; all we need to invoke is the fact that any kind of semantic minimalism will, by construction, disrupt the link between semantic uniformitarianism and strong semantic uniformitarianism: if there are no worldly truthmakers for declarative sentences, in virtue of which they come to be declarative, then any account which simply vindicates declarative sentences automatically makes them true.

Of course this move shifts the focus of the discussion from truth to vindication, and presents the semantic minimalist with the following challenge:

Vindicatory challenge: How can a semantic minimalist, with no recourse to substantial reference relations and worldly truthmakers, vindicate u-talk such as ‘electrons are negatively charged’?

What makes the vindicatory challenge particularly difficult for the subject-naturalist is the fact that, absent the machinery of reference, it is not immediately clear how we should even understand the conceptual content of declarative sentences in the first place. How, in other words, can we understand what ‘electrons are negatively charged’ even means, if we do not first establish what the labels ‘electron’ and ‘negative charge’ stick to in the world? This worry leads to the a more basic challenge for the subject-naturalist:

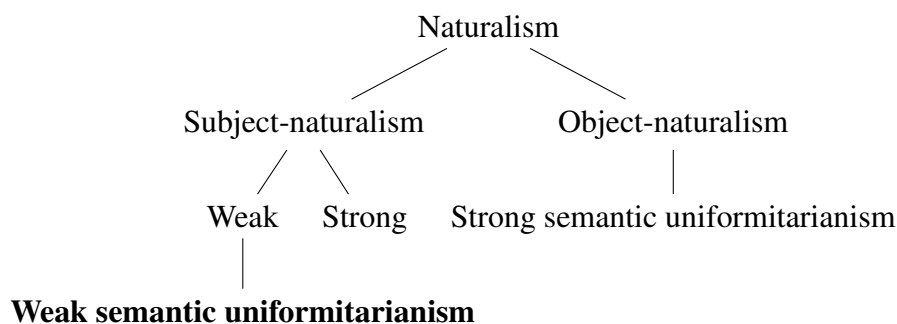
The conceptual challenge: How can a subject-naturalist, with no recourse to substantial reference relations and worldly truthmakers, articulate the conceptual content of a sentence like ‘electrons are negatively charged’?

<sup>2</sup> ‘S’ is true if and only if S, where ‘S’ names the sentence S.



In what follows, I will demonstrate how broader metasemantic arguments have been successfully wielded by the subject-naturalist in responding to these challenges. These arguments have, as a consequence, important upshots for the debate around scientific realism that, I think, have not been sufficiently appreciated.

Fine’s ‘in the same way’ edict is precisely the sort of injunction that Price endorses: to give u-talk the same semantic treatment as we give ordinary talk. However, under the hitherto hidden assumption of object-naturalism—that semantic uniformitarianism mandates referentialism—this is equivalent to understanding ‘in the same way’ as mandating the use of *reference* for both u-talk and ordinary talk. The resulting strong semantic uniformitarianism is what leads to the impasse between the scientific realists and anti-realists that we began this section with. But now we can see why the subject-naturalist argues that this is a bad starting point: their semantic minimalism shows that the move from weak semantic uniformitarianism to strong semantic uniformitarianism is a mistake. In terms of our decision tree, the weak subject-naturalist only has one option that does not collapse their position into object-naturalism: weak semantic uniformitarianism.



## 5 Normative pragmatic metasemantics

So far, I have been relying on an intuitive gloss of ‘metasemantics’, according to which metasemantic questions are questions about semantic properties and structures. When discussing minimalism about truth, for example, we were answering a metasemantic question of the form ‘what sort of thing is truth’? Following (Burgess and Sherman [2014]), I’d like to distinguish between two distinct metasemantic questions about meaning:

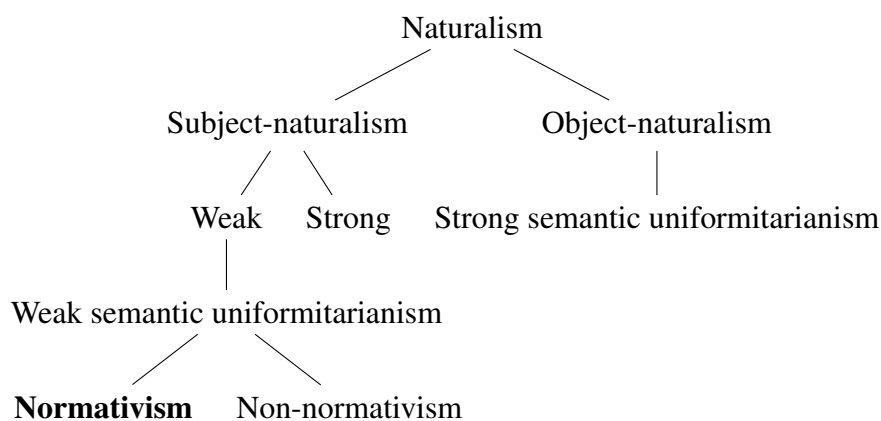
Basic metasemantic question: What facts determine or ground the basic semantic facts? (for example, facts of the form ‘*E* means *M*’)

Metaphysical question: What sort of thing is a meaning?

In what follows, following (Sellars [1974]), I offer a deflationary account of the metaphysical question: we should not be tempted by a slogan of the form ‘meaning determines referents (or representeds)’ into hypostasising meanings any more than we should take talk of doing something for someone’s sake to mean we should hypostasise ‘sakes’ (Quine [1960]). Consequently, the important question for us is the basic metasemantic one.

To answer the basic metasemantic question, we begin by looking at the ‘basic pragmatic question’: How should we account for how linguistic expressions are used? Prima facie, the basic metasemantic and basic pragmatic questions are distinct, and answerable, in principle, by appeal to very different fields of inquiry. One might, for example, understand the basic metasemantic question as an exercise in applied metaphysics, and the basic pragmatic question as an exercise in sociology or social anthropology. Call this the *independence thesis*. One version of strong semantic uniformitarianism might embrace the independence thesis, arguing that the semantic facts are grounded in the primitive relation of reference, and that pragmatic facts are beyond the remit of philosophers of language. An alternative version of strong semantic uniformitarianism might deny the independence thesis, and understand pragmatic facts as grounded in facts about successful reference: speakers’ uses of expressions are appropriate because the words they use latch on to the world correctly.

In what follows, we will explore a different independence-thesis-denying view called *normative pragmatism*, according to which questions about what meanings are and how expressions come to mean what they do are nothing over and above questions about how those expressions are correctly used (where correctness is relative to the rules of some discursive practice, hence ‘normative’). Having identified these two questions, we can ask what sorts of practical norms establish semantic facts like those about meaning. If normative pragmatism is to be in a position to satisfy the uniformity thesis, it must be able to establish weak semantic uniformitarianism. This requires a further component, which I will introduce in the next section. But first, let us add a new branch to our decision tree:



## 6 Inferentialist metasemantics

The subject-naturalist eschews referentialist metasemantics (even if they accept reference relations). So if they are to be taken seriously at all, they owe us, at the very least, a positive account of a replacement metasemantics. The minimal desideratum for such an account, I contend, is that it satisfactorily resolves the two challenges articulated in §4: the conceptual and the vindicatory challenges (of course, ultimately it also has to underpin an account of science that satisfies uniformity). In this section, I provide just such an account, highlighting the decisions

made with respect to a number of metasemantic and methodological choice points.

At the end of the last section, we noted that something needed to be added to normative pragmatism to ensure weak semantic uniformity. In this section, following (Brandom [1994], [2009]), I introduce that final element: inferentialism. Brandom is both a pragmatist and a weak semantic uniformitarian, arguing that all discursive practices, however varied, can bestow upon language a uniform semantic treatment. He establishes this by arguing that there is set of norms, adherence to which distinguishes a practice as discursive. In addition, these demarcating norms underwrite the assertoric nature of this practice. These are precisely the norms that the weak semantic uniformitarian was looking for.

Which norms? The norms of inference. In what follows, I highlight some of the important components of the normative-pragmatist-inferentialist view, before combining them all to articulate the core metasemantic thesis that underpins inferential scientific vindicationism.

**Inference is to be understood in a broad sense:** On an austere understanding of what an inference is—e.g. an ordered n-tuple of sentences—inferentialism is a non-starter as a metasemantic theory, since it plainly does not have the resources to anchor language onto the world. We therefore need to broaden the concept of inference to include what Sellars ([1954]) calls ‘language entry- and exit-transitions’: rules for how a worldly event or stimulus gets incorporated into a discursive performance (I see a snake and say ‘Yikes! A snake!’) and how discursive performances lead to intentional actions (you hear my claim, and move away). These moves are also treated as inferences.

**Sentences are the basic semantic units:** For the referentialist, the basic semantic unit is a term (e.g. a singular term, a name), out of which a sentence can be built using the appropriate compounding devices (connectives, quantifiers, predicates, etc.). By contrast, for the inferentialist, the basic unit is the sentence. Sub-sentential expressions like words are derivative. If substitution of one expression for another in a claim does not turn a good inference into a bad one, then those expressions are synonymous (Brandom [1994], Ch. 6).

**Meaning is contribution to good *material* inference:** Here is the standard account of logical inference: logical inferences are made on the basis of formally articulated rules, that invoke the (logical) form but not the content of claims. Grasp the rules, and you grasp the inferences. On this view, the correctness of an inference is purely a matter of logical form. The inference from ‘Anil Kumble has more wickets than Glenn McGrath’ to ‘Glenn McGrath has fewer wickets than Anil Kumble’ is valid only because it is an instance a schema with a suppressed premise, of the form ‘If A has more wickets than B, then B has fewer wickets than A’. (Sellars [1953]) distinguishes such logical inferences from a class of *material inferences*. That a material inference is valid for a community is just a primitive fact about that community’s discursive practice. The inference from ‘Anil Kumble has more wickets than Glenn McGrath’ to ‘Glenn McGrath has fewer wickets than Anil Kumble’ is materially valid, and does not rely on unsuppressing a hidden premise for its vindication. Material inferences, and not logical or formal inferences, serve as the basis for claims about meaning: the meaning of a claim is nothing more than the contribution that that claim makes to good material inferences.

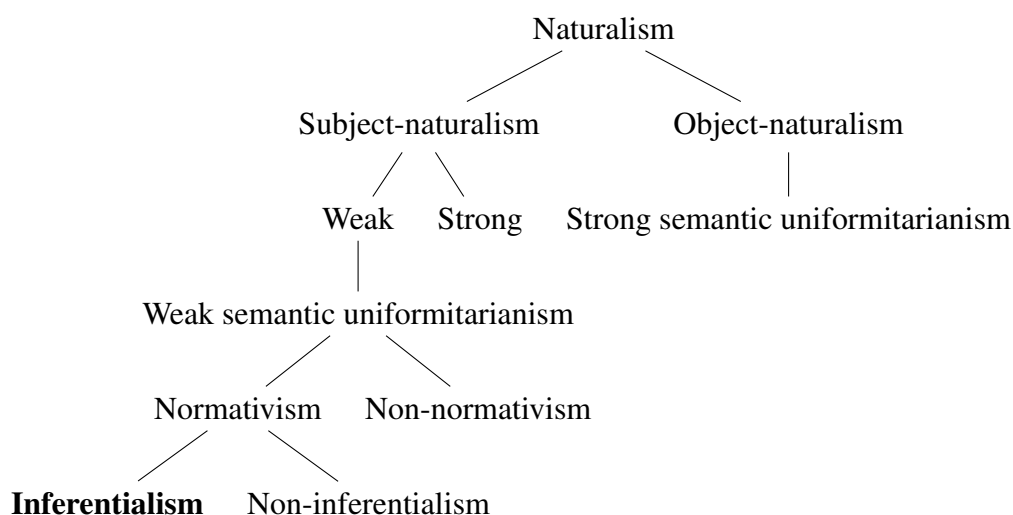
**Inferentialism is a form of weak semantic uniformitarianism:** Logical inference rules can be defined from material inferences in the following way: delineate (if even just by stipulation) certain vocabulary as ‘logical’. An inference can then be treated as valid in virtue of its logical form if and only if (i) it is a materially good inference and (ii) it cannot be turned into a materially bad inference by substituting non-logical vocabulary for non-logical vocabulary in its premises and/or conclusions. This is non-circular, because we take material goodness as primitive (what justifies this move is the fact that the world will eventually penalise agents for taking, as good, bad material inferences; more on this below). In the Kumble/McGrath example above, we can identify the indicative conditional as making explicit what was implicit in our endorsement of the material inference. The derivation of formal rules of inference from material inferences is an example of a process known as *explicitation*: it makes explicit, using some vocabulary (in this case logical), what is implicit in the practice of agents. In particular, it renders open to debate a particular agent’s epistemic commitments and entitlements. This generalises to all vocabulary, although different sorts of vocabulary will make explicit different aspects of an agent’s collection of commitments and entitlements. Here is another example: the use of ‘not’ makes explicit commitments to various sorts of incompatibilities.

As Brandom puts it: ‘paradigmatically rational processes... [depend] on the possibility of making implicit commitments explicit in the form of claims. *Expressing* them in this sense is bringing them into the game of giving and asking for reasons as playing the special sort of role in virtue of which something has a conceptual content at all, namely an inferential role, as premise and conclusion of inferences’ ([2009], p. 57). In this quote, we also see how Brandom understands inferentialism as a form of weak semantic uniformitarianism: the ‘game of giving and asking for reasons’, not a relation of representation, is what imbues certain acts with conceptual content and therefore marks them as assertoric, and thus uniform in their semantic treatment.

**Inferentialism systematises language use using a scorekeeping model:** The significance of a linguistic performance is measured by how it changes the commitments and entitlements that one attributes to a speaker, and that one undertakes oneself. This information, collectively, constitutes what Brandom calls a ‘deontic scorecard’. To identify which commitments and entitlements are to be kept track of (i.e. which commitments count towards the deontic score), we look to those which are articulated inferentially. The norms of inference are what determine the rules of scorekeeping. For example, if I say something like ‘Priya believes of the red cricket ball that it is hard’, I immediately attribute (i) commitments such as ‘Priya believes of the cricket that it is solid’, (ii) entitlements such as ‘Priya believes that Bangalore is the capital of India’ (it is not), and I myself undertake a commitment such as ‘the cricket ball is red’. Different inferential norms will lead to different ways of updating my deontic scorecard. And the explicitation of those norms is what articulates the reasons for my decisions to update in the way that I do.

Putting all this together, we can characterise the Brandomian inferentialist thesis as the following conjunction:

- (1) The meaning of a claim is exhausted by its contribution to the inferential valence of the material inferences in which it is caught up.
- (2) These material inferences (i) identify a set of practices as distinctively discursive and (ii) are primitive fact about that practice. Which material inferences are good is constantly being updated by the practice as it evolves through the constant interaction of its agents with their environment.
- (3) Metasemantic facts are reducible, without residue, to use facts, which are governed by implicit rules.
- (4) The expressive role of vocabulary is to make explicit, according to the norms governing the material inferential structure of an agent's linguistic practice, what is implicit in their claims.



## 7 Inferential scientific vindicationism

In this section, I present inferential scientific vindicationism. I begin, in §7.1 by discussing how the standard object-naturalist (referentialist) approach purports to realise the naturalist impulse. I then contrast that approach with my preferred alternative, the inferentialist approach. In §7.2, I demonstrate how, despite denying the metasemantic importance of reference, the inferentialist can nonetheless make sense of the appropriateness of reference-talk. I will use, as a representative example of the sort of locution to be vindicated by the scientific realist, ‘electrons are unobservable, and obey Pauli’s Exclusion Principle’.

### 7.1 How to vindicate claims inferentially

For the object-naturalist, the philosophical action is taking place at the level of (purported) objects in the world, and any interesting consequences for their descriptive language follow from the way the world is. So if we want to vindicate a speaker, *S*’s, utterance of ‘electrons are

unobservable, and obey Pauli's Exclusion Principle', then we need to look (i) to the referent of 'electron', for the identity of the object of the claim, and (ii) to the extensions of properties like 'unobservable' and 'obeys Pauli's Exclusion Principle'.

It will be helpful to introduce some vocabulary borrowed from modal metaphysics at this point. Following (Nelson [2023]), we define an attribution as 'metaphysically *de re* with respect to an object *o* just in case it directly attributes a property to *o*.' If we extend this terminology to cover propositional attitudes, then *S* expresses a *de re* attitude, if that attitude is about some part of the world. For a referentialist, the identity of that something is what is picked out by the reference relation, as the referent of the central term (in this case 'electron'). This is a semantically identical procedure to what goes on when *S* makes a claim about observables like chairs or tables.

We can extract from this a 'vindication schema':

Classification step: Identify the circumstances that signal that *S* is making a *de re* claim.

Content-extraction step : Identify the contents of the *de re* claim, in particular the part of the world that is the target of the claim. [This addresses the conceptual challenge.]

Vindication step: Vindicate the *de re* claim by assessing the contents of the claim. [This addresses the vindicatory challenge.]

The object-naturalist scientific realist instantiates the schema as follows:

Classification step: *S* is making a declarative claim. For an object-naturalist, a declarative claim is to be understood as a claim *about* some worldly object.<sup>3</sup>

Content-extraction step : A token of the type 'electron' refers to a token of the type electron. 'Electron' means electron because of the primitive referential relation between them. Therefore the content of *S*'s claim is (partially) grounded in this relation. Similar accounts can be given of the properties.

Vindication step: 'Electrons are unobservable, and obey Pauli's Exclusion Principle' is vindicated because it is true. It is true because the electrons, picked out by the reference relation, have the ascribed properties, and are thus truthmakers of this utterance.

The subject-naturalist inferential scientific vindicationist offers an alternative account. Let us examine each step more closely.

**Classification step** The inferentialist classifies a performance as linguistic, (therefore assertoric) by investigating whether that performance has inferential significance, in the broad sense described in §6. *S*'s utterance of 'electrons are unobservable, and obey Pauli's Exclusion Principle' counts as declarative, according to the inferentialist's scruples.

<sup>3</sup> This is the move that traditionally gets the object-naturalist into trouble over problematic discourse like mathematical or moral. But for the scientific realist, this is not at all a concern, since the spatiotemporal existence of electrons is never really disputed.

But it does not follow from this classification that *S*'s utterance is automatically a *de re* utterance, given the inferentialist's denial of referentialism. So the inferentialist needs some alternative classificatory criteria for a *de re* utterance. Brandom highlights the pragmatist underpinning of the entire enterprise by noting that the identification of an utterance as *de re* is tantamount to the ascription of an attitude to the speaker. Consequently, our target for analysis is not 'electrons are unobservable, and obey Pauli's Exclusion Principle', but rather '*S* says (or claims) of electrons that they are unobservable, and obey Pauli's Exclusion Principle'. And the target is an attitude ascription to *S*, rather than a property ascription to (some component of) *S*'s utterance.

Brandom ([1994], [2009]) offers the following demarcation criteria. The ascription of a *de re* attitude to *S* by me is signalled by the aspect(s) of the deontic scorecard that keep track of three things:

1. The commitments I attribute to *S*, given what I know of their background commitments: e.g. '*S* believes of electrons that they are invisible', '*S* believes of electrons that they have no colour', etc.
2. The entitlements I attribute to *S*, given what I know of their background commitments: e.g. '*S* believes of protons that they obey the Pauli Exclusion Principle', '*S* believes of Helsinki that it is the capital of Finland', etc.
3. The commitments I myself undertake, on accepting *S*'s claim, given my background commitments: e.g. 'Electrons are fermions', 'Electrons are constituents of atoms, not force fields', etc

The identification of these three expressive characteristics of an utterance from *S* allows me to ascribe to *S* a *de re* attitude. On its own this is not the whole story. What I need in addition, is a way to identify the *res*, the worldly entity, that I take to be the target *S*'s claim. Fortunately, Brandom's criteria provide the ingredients to make that identification, i.e. to achieve 'content extraction'.

**Content-extraction step** The demarcation criteria for *de re* attitude ascriptions presented above were simply stipulated. In explaining why this stipulation works (i.e. in explaining what it is about these specific expressive roles that makes such locutions *de re* attitude ascriptions, rather than something else), an account will emerge of how the inferentialist achieves the content-extraction step of the vindication schema.

The central role of a *de re* claim is to make a declaration about some part of the world. To what end? Communication. The subject-naturalist takes their declarative utterances to correspond to some way the world is and they make those utterances in order to communicate this fact to some listener(s). But so does the object-naturalist! The disagreement, of course, is over what grounds the vindication of these declarative utterances. What the object-naturalist's order of explanation occludes is that there are two aspects of coordination associated with *de*

*re* claims: what we might call *i-coordination* and *e-coordination*.<sup>4</sup> *i-coordination* is a relation between the conceptual contents of two (or more) speakers' locutions, and is achieved when speakers correctly take themselves to be making claims whose conceptual content (i.e. inferential profile) is substantially shared by them. A particularly vivid instance of this is when two speakers correctly take themselves to be *i-coordinated*: i.e. to be speaking to each other about the same part of the world (e.g. entity or object). *e-coordination* is a relation between a speaker's expression's conceptual content and some part of the world. It is achieved when a speaker's utterances successfully pick out some part of the world, such as an object, and the speaker's associated descriptions appropriately track the properties of that object. For the object-naturalist, the reference relation establishes an *e-coordinative* link between the contents of an utterance and some worldly truthmaker. So if all interlocutors refer successfully to some referent, then their claims about that referent are automatically *i-coordinated* as well.

The inferentialist provides a different account of these two sorts of coordination. *i-coordination* is achieved by locutions that update the deontic scorecard in the way described above, by updating both ascriptions of commitments and entitlements to their interlocutors, as well as their own consequent undertakings. Note that this is possible only because the conceptual content of those utterances is prior to any representational/referential significance. *S* and *I* can ensure we are talking about the same sort of thing, electrons for example, by coordinating on sufficiently many of the meaning-conferring inferences that imbue 'electron' with meaning. Indeed, part of the goal of communication is to make explicit what the salient meaning-conferring inferences are, in order to ensure we are on the same page.

The three aspects of the deontic scorecard that are relevant to the characterisation of an attitude ascription as *de re* serve to highlight the differences in perspective between *S* and me. As such, it provides a model for communication. As Brandom points out, '[e]xtracting information from the remarks of others requires grasping what is expressed when one offers *de re* characterizations of the contents of their beliefs—that is, to be able to tell what their beliefs

<sup>4</sup> These two forms of coordination are also achieved by Price's notions of *i-* and *e-representation* respectively (see e.g. ([2011a]; [2013])). There's always a debate to be had about whether two notions are sufficiently similar in a given context that they warrant the same label. In this paper, at the risk of further polluting the terminological space, I choose to use these slight variations, rather than Price's original terms, to highlight the subtle differences between the antecedent contexts of the discussions in which they were introduced. Price introduces his two notions of representation in order to highlight that they have unhelpfully been conflated and subsumed under the banner of 'representation' in canonical discussions of non-cognitivism. This does not quite align with my more specific purpose in this paper, which is to understand, from within the Brandomian inferentialist framework, (i) the distinctively social character of the coordinative aspect of the act of representing, understood in terms of identifying the appropriate aspect of the game of giving and asking for reasons (hence 'i-coordination') and (ii) the covariation of some parameter of discursive practice with some aspect of the world (hence 'e-coordination'). In other words, *i-coordination* is what is achieved by *i-representation* in an inferentialist framework, and *e-coordination* is *e-representation* in the same framework. I should also point out that *e-coordination* is, in the Brandomian framework, what Sellars called 'picturing' ([1962]). Another reason to opt for a new vocabulary is to signal that I reject Price's strongly naturalistic thesis that *e-representation* just is *i-representation* according to the appropriate scientific inferential practice; no part of the characterisation of scientific realism, at least as I have presented it in this paper, requires that science be the *only* guide to what there is. I hope that scope and motivations are sufficiently discernibly distinct that my terminological choices are acceptable.



would be true *of* if they were true. It is to grasp the *representational* content of their claims... [D]oing this is just mastering the *social* dimension of their inferential articulation.' ([2009], p. 182) And doing this is what extracts the content from a *de re* claim.

**Vindication step** But this only gets the inferentialist as far as i-coordination. The next step is to achieve e-coordination. This is achieved by the practice's inferential web being sufficiently rich. Richness is to be understood quantitatively, both in terms of the number of (i) intra-linguistic inferential links, as well as (ii) broadly-construed inferential links introduced in §6 which includes language entry- and exit-transitions. The richer the inferential web in the first sense, the more information we can extract from each impingement of the world onto our discursive practice. The richer the web in the second sense the more the impingements. The establishment of a sufficiently tight e-coordination will be contingent on the inferential richness of the practice in both senses. An example here might help.

Imagine a community of scientists who accept the inference from 'electrons are negatively charged' to 'there exist entities that are perceptually inaccessible, with properties that determine the behaviour of voltmeters and ammeters.' There are many ways in which the world can reward or penalise a community for accepting an inference like this. For example, a community that accepts this inference might impose the further commitment that electrons are literally point-sized, thus rendering their individual charge densities infinite. This view of electrons is incompatible with classical electrostatics, as experiments will demonstrate. Those experiments, therefore, are instances of the world impinging on, and forcing us to update, our discursive practices. In particular, an observation from an electrostatics experiment counts as the premise of a language-entry transition. But such a premise is only useful if we understand the material inferences it is caught up in. The more material inferences a particular observation, or agent-world interaction, is caught up in, the richer the inferential web, in the first sense. Now, suppose the community has no way of understanding, say, a piece of data from the LIGO telescope as having implications for the meaning of 'electron', then that piece of data does not serve as a language-entry premise for adjudicating a claim about electrons. The more observations, or agent-world interactions, we can bring to bear as language-entry rules, the richer the inferential web, in the second sense.

The observation of a conflict with electrostatics might be enough for the community to dial back on the commitment to the point-sizedness of electrons while still holding on to the commitment to electrons being perceptually inaccessible. The community reaches an equilibrium, with respect to this claim, until some other material inference that they accept, one of whose premises is that electrons are perceptually inaccessible, leads to a conflict with the world. Now, the implications of the commitment to the perceptual inaccessibility of electrons might not immediately filter out towards anything that can be vindicated. In this case, the community might continue to mean, by 'electron', something perceptually inaccessible (and by 'perceptually inaccessible', among other things, that it applies to electrons). And given the holistic nature of meaning-conferral, it is, in general, possible to hold onto a particular material inference (e.g. electrons are negatively charged so there exist perceptually inaccessible entities) at the expense

of dropping another.<sup>5</sup>

Vindication of a claim comes, therefore, from the discursive practice being sufficiently inferentially rich as to allow a collection of agent-world interactions to adjudicate on that claim. This is true whether or not the claim is a form of u-talk. With u-talk, different communities of agents will take, as good (and meaning-conferring), different material inferences, and depending on the contingencies of their interactions with the world, may or may not be forced to update the meaning of ‘electron’. After sufficiently many impingements on a community’s discursive practice, that community might be justified in claiming that electrons are, in fact, perceptually inaccessible objects, characterised in whatever way their best science says that it is. The community of users of that theory will have achieved e-coordination. Understood in these terms, the project of theory interpretation is to understand the extent to which i-coordination between a theory’s users’ expressions’ conceptual content is e-coordinated with some proper part of the world.

We can now write down the vindication schema for the inferentialist, completing our characterisation of inferential scientific vindicationism.:

Classification step: *S* is making a claim which updates my deontic scorecard in a particular way (i.e. it makes me update ascribed and undertaken commitments and entitlements). I attribute to *S* a *de re* attitude about electrons.

Content-extraction step : The content of *S*’s claim is inferentially articulated, and intelligible to me in virtue of the communication of information that is codified in my updated deontic scorecard.

Vindication step: ‘Electrons are unobservable, and obey Pauli’s Exclusion Principle’ is vindicated because it forms part of a dense web of inferences, with sufficiently many entry- and exit-points to allow the world to impinge on the practice enough that the bad inferences that it corrects do not render *S*’s claim false.

## 7.2 Reference for inferentialists

For the inferentialist, the representational capacity of a discursive practice does not need to be absolutely factorisable. Questions like ‘does the metric tensor represent spacetime?’ are akin to asking whether the letter ‘e’ represents some part of an electron, just because ‘e’ is part of ‘electron’. Representation relations, in general, exist between models and worlds, and theory factorisability, into relations between parts of models and parts of worlds is highly contingent. But even if we restrict ourselves to situations where relations exist between parts of models and parts of the world, it is far from clear that, from an inferentialist point-of-view, that reference should be prioritised. Nonetheless, inferentialists do owe us an account of why reference is, at

<sup>5</sup> As in the last section of Quine’s *Two Dogmas of Empiricism* ([1980]).

least in some circumstances, in good philosophical standing. For example, we now know that tokens of 'virus' refer to viruses, even though virus-talk is u-talk.

While the inferentialist is committed to all external impingements contributing to inferential valence, and hence meanings, of claims, they are not committed to an egalitarianism about how much these impingements contribute to the *referential* valence of a claim, understood in terms of how our deontic scorecard keeps track of *de re* ascriptions. In many circumstances, for many practical purposes, a local inferentialism is indistinguishable from a global one.

Understood in light of the derivation of representational relations from inferential relations, we can imagine a collection of word-world relations that look and behave a lot like reference relations locally, in the sense that any deviations from these relations' behaviour as reference relations do not manifest themselves under certain circumstances. Take, for example, the word 'mammal'. Almost all mammals do not lay eggs. But the class of referents of 'mammal' includes some egg-layers. But we can define a new class of animals which includes all and only mammals that do not lay eggs. This class is not the class of referents of 'mammal.' But we can define a new relation, one of 'quasi-reference', such that the class of referents of 'mammals who do not lay eggs' is the class of quasi-referents of 'mammal'.

Since the inferentialist does not rely on reference to imbue terms with meanings, reference is not philosophically privileged over quasi-reference. By 'mammal', the inferentialist means 'something whose referents include anteaters and duck-billed platypuses, but whose quasi-referents exclude them.' But under many circumstances (for example, if you want to discover whether this egg-laying creature you've just discovered in Sweden can fly), the reference-conferring link between egg-laying and 'mammal' might be safely ignored. The world might cooperate in allowing us, in many circumstances, to drop more and more meaning-conferring links, without making the class of referents of particular terms deviate significantly from the class of its quasi-referents. In many cases, selecting a 'for all practical purposes' quasi-referential relation by dropping, as reference-conferring, a number of inferential links, makes holistic e-coordination appear localised. So the inferentialist can account for the appropriateness of reference-talk: it applies in circumstances where the world is sufficiently cooperative with our discourse that it does not penalise us for using holistic quasi-reference relations as if they were atomistic reference relations.

Realism is standardly understood as a genus of views united by a commitment over what sort of thing has the authority to vindicate claims: worldly truthmakers. Scientific realism, in particular, is a referential view about what vindicates scientific discourse, in particular u-talk. The inferential scientific vindicationist argues that vindication of a scientific claim comes from a subject-naturalistic account of the appropriateness of the use of that claim in a particular set of scientific contexts. As I just demonstrated, sometimes, those contexts are sufficiently divorced from other circumstances that, for all practical purposes, the associated claims can be taken to be vindicated by objects being a particular way, where those objects are picked out by a quasi-local relation of quasi-reference to a word. In those contexts, the inferential scientific vindicationist is practically indistinguishable from the scientific realist. And in those contexts,

standard arguments for scientific realism apply; the mistake is to extend those claims to a global schema, and to treat the contingent cooperation of the world with our use of reference-talk in certain circumstances, at certain scales, as universal.

## **8 Vindicating inferential scientific vindicationism**

I have introduced inferential scientific vindicationism by contrasting its constituent commitments, at each point in the decision tree, to some alternatives. Although I do endorse many of those moves for independent reasons, ultimately, the justification of inferential scientific vindicationism comes from demonstrating that it beats the competition. In this section, I demonstrate that, with respect to the some plausible naturalistic desiderata (detailed below), inferential scientific vindicationism is preferable to both structural realism (8.1) and constructive empiricism (8.2), at least as they are canonically understood, i.e. as forms of representationalism.

The first desideratum of a naturalistic account of scientific theory interpretation is that it satisfy the uniformity thesis. But this should not be our only desideratum. After all, mere metaphysical uniformity on its own, is rather weak. It just enjoins us to treat the ontological categories at different levels of description on the same footing, without telling us what that footing ought to be. However, if we are moved by Finean considerations, as a number of naturalists are, then we have a proposal for what provides that footing: the objects of ordinary empirical discourse (OED). So a second naturalistic desideratum is that our naturalistic account of a scientific theory interpretation not be metaphysically revisionary with respect to the objects of OED.

The final desideratum is familiar from the vast literature on scientific realism: that it deal with the challenge of underdetermination. This challenge broadly speaking, takes two forms: synchronic underdetermination of which weak and strong underdetermination (see e.g. (Ladyman [2012])) are subspecies, and diachronic underdetermination, of which the pessimistic meta-induction (Laudan [1981]) is perhaps the best-known. What they all have in common is the following: they identify a problem with the plausible suggestion that science and metaphysics employ different levels of descriptive grain. If the level of descriptive grain of metaphysics is, as is suggested, finer than that of scientific theories, then scientific evidence underdetermines metaphysical descriptions.

In summary, we can identify the following three desiderata for a naturalistic account of theory interpretation:

1. Satisfaction of the uniformity thesis
2. Metaphysical conservatism about the objects of OED
3. Immunity to the underdetermination challenge

As we saw in §2, scientific realism satisfies the first two desiderata quite easily, but fails the underdetermination challenge. So unmodified scientific realism is not an option for us.

## 8.1 Structural realism

Structural realism is a family of views according to which the descriptive capacities of our best science are exhausted by claims about the structure of the world. Broadly speaking, there are two ways that this might be the case: either (i) there are other entities, but our science is limited to only being able to describe, at best, the structures that these entities instantiate or (ii) there are no further entities to instantiate structures, so the reason that our best science can only describe structure is that structure is all there is. Following (Ladyman [1998]), we can call former variant ‘epistemic structural realism’ and the latter, ‘ontic structural realism.’ Epistemic structural realism, by construction, denies epistemic uniformity, so it fails to satisfy our first naturalistic desideratum. So we will focus, in this section, on ontic structural realism.

Ontic structural realism satisfies the uniformity thesis relatively easily: structure is all there is, so there is metaphysical uniformity; scientific knowledge is structural knowledge at all levels of description, so there is epistemic uniformity; structure is all that is described by our science, so there is semantic uniformity. So our first naturalistic desideratum is satisfied.

The underdetermination challenge arises for the scientific realist because the fineness of their scientific descriptive grain is outstripped by the requisite fineness of metaphysical descriptive grain. The ontic structural realist simply denies that there is any difference in descriptive grain between science and metaphysics so no problem of underdetermination arises. So our third naturalistic desideratum is satisfied.

Things are a little more complicated with respect to the desideratum of a non-revisionary metaphysics of the objects of OED. The central claim of ontic structural realism is that (the worldly structure represented by) the mathematical structure of our best physics exhausts our fundamental ontological commitments. The elimination of objects from our fundamental ontology does not immediately preclude our being able to vindicate object-talk at a non-fundamental level, though. (Ladyman et al. [2007]) drop the tacit Quinean standard for ontological commitment, and adopt instead a Dennettian one, which (Wallace [2012], p. 50) summarises as:

Dennett’s criterion: A macro-object is a pattern, and the existence of a pattern as a real thing depends on the usefulness—in particular the explanatory power and predictive reliability—of theories which admit that pattern in their ontology.

Without delving too deeply into an assessment of Dennett’s criterion as a criterion for ontological commitments, we should note that this is deeply metaphysically revisionary, not only with respect to how inclusive it is (haircuts and symphonies become objects on this view), but also with respect to what it says about the objects of OED. It implies, as Ladyman and Ross emphasise, the scale-relativity of ontology, which says that ‘claims about what really... exists should be relativized to... scales at which nature is measurable’ (Ladyman et al. [2007], p. 200). So chairs and tables do not exist simpliciter; they exist-at-classical-physics-scale. Perhaps, for some, this is acceptable. It might even ground a Moorean shift against our second naturalistic desideratum. But for people of a more metaphysically conservative bent, inferential scientific

vindicationism offers an alternative, with all of the benefits of structural realism, but without the metaphysical revisionism; it is neutral with respect to claims about ontology, so you can pick your favourite one. Whereas OSR holds on to the uniformity thesis by being deeply metaphysically revisionary, with respect to the objects of OED, inferential scientific vindicationism holds on to the uniformity thesis by being semantically revisionary. As a result, the latter, but not the former, satisfies the second naturalistic desideratum.

## 8.2 Constructive empiricism

If read object-naturalistically, constructive empiricism is untenable; the object-naturalist, *qua* representationalist, does not have the resources to prise apart epistemic anti-realism from semantic anti-realism. This is because e-coordination is central to meaning ascriptions to u-talk. The most plausible way for an object-naturalist epistemic anti-realist to be able to speak meaningfully about electrons is to be a semantic anti-realist; to view u-talk as elliptical for OED. The less appealing alternative for the object-naturalist epistemic anti-realist is to be a semantic realist, and on top of that, to hold a negative or agnostic epistemic attitude towards the entities whose existence grounds the meanings of theoretical claims. Although not obviously logically inconsistent, this is Moore-paradoxical ([1942]); on this view the constructive empiricist who wants to vindicate claims like ‘electrons are negatively charged’ is committed to locutions of the form:

(E\*): ‘Electrons are negatively charged’ is vindicated by the fact that ‘electron’ refers to an unobservable electron (and *mutatis mutandis* ‘negatively charged’), but I don’t know that ‘electrons are negatively charged’ is vindicated by the fact that ‘electron’ refers to an unobservable electron (and *mutatis mutandis* ‘negatively charged’), rather than the fact that ‘electron’ refers to an unobservable selectron or belectron or...

An immediate response suggests itself: since the semantic realist, as per Psillos’ characterisation, only requires that theoretical terms have *putative* factual reference, the constructive empiricist need not assent to (E\*), but need only assent to the weaker claim:

(E\*\*): ‘electron’ purports refer to an unobservable electron (and *mutatis mutandis* ‘negatively charged’) but I don’t know that ‘electron’ refers to an unobservable electron (and *mutatis mutandis* ‘negatively charged’).

where ‘purporting to refer’ can be cashed out in various ways. For example, by appeal to a Fregean notion of sense.

This certainly is an option for the constructive empiricist. However, it is not an option for the *object-naturalist* semantic realist, for whom the truthmakers imbue ‘electron’ with meaning only if reference is successful; compare ‘electron’ with ‘phlogiston’. So the constructive empiricist position remains viable only if it is understood subject-naturalistically.

If read inferentially, and in particular, if given an i-coordinative reading of the meanings of theoretical terms, the constructive empiricist's epistemic anti-realism can be given a secure semantic footing: 'electron' is to be understood in terms of how its use updates speakers' deontic scores. This exhausts its conceptual content, and establishes its i-coordinative profile. This is enough for speakers to coordinate on the meaning of 'electron', and then engage in the nontrivial endeavour of deciding how positive one's doxastic attitudes towards its e-representational profile should be. van Fraassen refrains from recommending a positive global attitude, while scientific realists like Psillos ([2005]) endorse such an attitude.

For the inferential scientific vindicationist, the constructive empiricist's epistemic anti-realism can be articulated as the claim that we are not rationally mandated to believe that i-coordination of (some) scientific terms (that come from u-talk) will ever amount to e-representation. And the scientific realist's optimism as the claim that, for the appropriately delineated 'mature' sciences at least, i-coordination of (some) scientific terms determines e-coordination for those terms. Importantly, the machinery of inferentialism delineates the circumstances in which i-coordination and e-coordination coincide: when the inferential web is sufficiently rich, in the manner described in §7. The question of whether or not we should be realists thus becomes a question of the evidence we have for the claim that u-talk's talk i-coordinative profile is the same as (or sufficiently similar to, or such that we believe it will converge to) its e-coordinative profile. Equipped with the machinery of inferentialism, the constructive empiricist can realise their epistemic anti-realism in a manner that doesn't invoke an unobservable/observable distinction: to be (an inferential) constructive empiricist is to believe that the discursive practices of physics will always be such that we are not rationally mandated to accept that their inferential richness will endow at least some elements of u-talk with an e-representational profile. And to be a scientific realist is to believe that the (appropriately delineated, 'mature') discursive practices of physics will always be such that their inferential richness will endow at least some aspects of u-talk with an e-representational profile

Constructive empiricism is intended as an expression of an admirable epistemic humility. But we should not allow our epistemic humility misguide us into underestimating how epistemically successful we have been, and might be in the future. The constructive empiricist's hypothesis that we never have been and never will be in a situation where we are rationally mandated to accept, as knowledge, claims about unobservables stretches this epistemic humility too far. The constructive empiricist accuses the scientific realist of an unwarranted epistemic optimism, but suggests in its stead an unattractive epistemic pessimism.

The inferential scientific vindicationist, on the other hand, occupies the middle ground between these two extremes. They can make sense of these sorts of global assessments of the descriptive capacities of scientific theories. But they can go much further. Inferential scientific vindicationism is not committed to an a priori global view regarding epistemic optimism. The inferential scientific vindicationist can demarcate practices about which we should be epistemically negative or agnostic from those about which we should not be, on the basis of local facts about what sorts of aims a given practice can realise. Importantly, this epistemic assessment

is made in the same way, whether the associated knowledge claims are about OED or u-talk, namely by an assessment of how closely i-coordination tracks e-coordination. The important gap is not between perceptual abilities regarding observables and our perceptual inabilities regarding unobservables, it is between our discursive abilities being insufficiently rich to close the gap between i-coordination and e-coordination. Consequently, and unlike the constructive empiricist, the inferential scientific vindicationist can accept the epistemic uniformity thesis (and with it, the entire uniformity thesis).

## 9 Conclusion

I began this paper by highlighting the naturalistic impulse that made scientific realism attractive. The more closely one looks at the standard ways in which the semantics of scientific theories is discussed in the philosophy of science literature, the more clearly one sees a disconnect between the purported goal of taking scientists seriously, and the actual philosophical practice of taking scientists seriously only insofar as their linguistic practices can be fit into the Procrustean bed of the semantic resources appropriate to our everyday domain of medium sized dry goods.

I proposed that we needed better tools in order to properly do justice to the naturalistic impulse, which I took to be usefully articulated by Fine under the heading of the Natural Ontological Attitude. In what followed, I systematically added tools in such a way as to discover the appropriate semantics, by considering a series of metasemantic arguments. Price's subject-naturalism was the first component, whose adoption was motivated by the observation that our semantic vocabulary itself needs to be justified naturalistically, rather than stipulated.

Subject-naturalism is a form of pragmatism, where the latter is understood as a set of views unified by a commitment to the inseparability of knowledge of the world from agency within it. This can manifest itself metasemantically in any number of ways. So the second component was normative pragmatics, according to which claims about the world were to be understood not as automatically assertoric, but instead as expressing a speaker's endorsement of a set of practical norms that mediate their interactions with the world and with other speakers.

The final component was Brandomian inferentialism, which recovers the requisite semantic uniformity by endorsing an inferentialist metasemantics, according to which (i) discursive practices are demarcated as such by being characterisable as moves in the game of giving and asking for reasons, and (ii) the meaning of a claim is its contribution to good material inferences. As a package, inferential scientific vindicationism provides a prescription for taking science seriously, in a manner that does justice to the varied and nuanced ways in which our discursive practices, scientific or otherwise, come to represent, symbolise or evoke aspects of the world. On top of that, it both identifies and accounts for the circumstances in which our philosophical attitudes towards science most closely resemble canonical views, such as constructive empiricism and traditional scientific realism, and those in which it deviates. It earns its stripes as an attractive alternative to these views by satisfying three plausible naturalistic desiderata that form an inconsistent triad for the canonical views. Inferential scientific vindicationism thus



resolves some deep problems for the scientific realist, while arguably doing even more justice to the naturalistic impulse that motivated that realism in the first place.

*Dianoia Institute of Philosophy,  
Australian Catholic University,  
Melbourne, Australia,  
tm399.acu@pm.me*

## References

- Brandom, R. [1994]: *Making It Explicit: Reasoning, Representing, and Discursive Commitment*, Harvard University Press, ISBN: 978-0-674-54330-0.
- Brandom, R. [2009]: *Articulating Reasons: An Introduction to Inferentialism*, Harvard University Press.
- Burgess, A. and B. Sherman [2014]: *Metasemantics: New Essays on the Foundations of Meaning*, Oxford University Press, ISBN: 978-0-19-966959-2.
- Chakravartty, A. [2007]: *A Metaphysics for Scientific Realism: Knowing the Unobservable*, Cambridge University Press.
- Feigl, H. [1950]: ‘De Principiis Non Disputandum. . . ?’, in *Philosophical Analysis, A Collection of Essays*. Max Black (Ed.) Ithaca, NY: Cornell University Press, pp. 119–56.
- Fine, A. I. [1984]: ‘The Natural Ontological Attitude’, in *Scientific Realism*, University of California Press.
- Jackson, F. [1998]: *From Metaphysics to Ethics: A Defence of Conceptual Analysis*, Clarendon Press.
- Kitcher, P. [1993]: *The Advancement of Science: Science without Legend, Objectivity without Illusions*, Oxford University Press, USA.
- Ladyman, J. [1998]: ‘What is structural realism?’, *Studies in History and Philosophy of Science Part A*, **29**, pp. 409–24.
- Ladyman, J. [2012]: *Understanding Philosophy of Science*, Routledge.
- Ladyman, J., D. Ross and D. Spurrett [2007]: *Every Thing Must Go: Metaphysics Naturalized*, Oxford University Press.
- Laudan, L. [1981]: ‘A Confutation of Convergent Realism’, *Philosophy of science*, **48**, pp. 19–49.
- Moore, G. E. [1942]: ‘A Reply to My Critics.’, in *The Philosophy of G. E. Moore*. Ed. P. Schilpp. La Salle: Open Court, pp. 535–677.
- Nelson, M. [2023]: ‘Propositional Attitude Reports’, in E. N. Zalta and U. Nodelman (eds.), *The Stanford Encyclopedia of Philosophy*, Spring 2023, Metaphysics Research Lab, Stanford University.
- Price, H. [2004]: ‘Naturalism without Representationalism’, in *Naturalism in Question*, Harvard University Press Cambridge, MA, pp. 71–88.

- Price, H. [2010]: ‘One Cheer for Representationalism?’, *RE Auxier y LE Hahn (comps.), The Philosophy of Richard Rorty, Open Court, Chicago*, pp. 269–89.
- Price, H. [2011a]: ‘Expressivism for Two Voices’, *Pragmatism, science and naturalism*, **30**, pp. 87–114.
- Price, H. [2011b]: *Naturalism without Mirrors*, Oxford University Press.
- Price, H., S. Blackburn, R. Brandom, P. Horwich and M. Williams [2013]: *Expressivism, Pragmatism and Representationalism*, Cambridge University Press.
- Psillos, S. [2005]: *Scientific Realism: How Science Tracks Truth*, Routledge.
- Quine [1960]: *Word and Object*, Cambridge, MA: MIT Press.
- Quine, W. V. O. [1980]: *From a Logical Point of View: Nine Logico-Philosophical Essays*, Harvard University Press.
- Sellars, W. [1953]: ‘Inference and Meaning’, *Mind*, **62**, pp. 313–38, ISSN: 0026-4423, JSTOR: [2251271](#).
- Sellars, W. [1954]: ‘Some Reflections on Language Games’, *Philosophy of Science*, **21**, pp. 204–28.
- Sellars, W. [1962]: ‘Truth and "Correspondence"', *The Journal of Philosophy*, **59**, pp. 29–56.
- Sellars, W. [1974]: ‘Meaning as functional classification: A perspective on the relation of syntax to semantics’, *Synthese*, **27**, pp. 417–37.
- Van Fraassen, B. C. [1980]: *The Scientific Image*, Oxford University Press.
- Wallace, D. [2012]: *The emergent multiverse: Quantum theory according to the Everett interpretation*, Oxford University Press, USA.