How to be a scientific realist: a normative-pragmatist proposal*

Tushar Menon

Department of Linguistics and Philosophy

Massachusetts Institute of Technology,
32 Vassar Street, Cambridge, MA 02139.†

ABSTRACT

Motivated by the question about whether we should be realists about dark matter, I propose a new articulation of the debate between the scientific realist and anti-realist. I defend three claims. First, that the debate should be articulated in normative terms, where questions about normativity are understood as being questions about authority. Second, that positions in this debate should be defended using pragmatist strategies, where pragmatist strategies are understood as being agent-first strategies. Third, that the manner of implementation of a pragmatist strategy with respect to some scientific-theoretical vocabulary—such as 'dark matter'—is highly domain-specific and turns on choices made by agents about what (and whom) they recognise as authoritative.

1 *Introduction*

2 *Realism and anti-realism, alethically*

3 *The normative framework*

4 *The Semantic Question*

5 *The Epistemic Question*

6 *Realism and anti-realism, normatively*

7 *Conclusion*

A *An example of a deontic scorecard*

*Forthcoming in *Studies in History and Philosophy of Science.*
†email: mprna478@mit.edu

1 Introduction

The scientific realism debate is defective, but it doesn't have to be. Realist and anti-realist impulses are both reasonable.[1] The debate is defective because its traditional articulation fails to do justice to those reasonable impulses. To remedy this defect, we will have to confront some issues in the philosophy of language that have typically been sidelined by philosophers of science. Once we confront these issues, we make space for a superior articulation of the debate. A case study from astrophysics exemplifies this superiority.

Contemporary astrophysics appears to have given us a new iteration of what we might call an *expansion scenario*: what if all of our knowledge about the material world is knowledge of only a tiny fraction of it? In the past our world view has had to expand rapidly to encompass other continents, other planets, and other galaxies. Astrophysical data suggests that our contemporary world view might, once again, have to be expanded significantly, this time to include another kind of matter called 'dark matter'. The question of whether or not we are dealing with an expansion scenario is the question of whether or not we should be realists about dark matter.[2]

But the would-be dark matter realist has a problem: right now, there are several serious yet incompatible suggestions for what dark matter is ((Balazs et al. [2024]; Bertone and Tait [2018]; Garrett and Dūda [2011]) provide helpful overviews). Martens [2021] highlights the consequent shortcoming of a traditional form of dark matter realism: 'can one even be a realist about x, can one say that x exists, without knowing (almost anything about) what x is?' He concludes that the present status of research casts so wide a net that whatever the different dark matter research programs have in common 'is so semantically thin, so vacuous, that it barely means anything at all... To be a realist about the unknown is not to be a realist' (Martens [2021], p. 6).[3] This situation should alarm both realists and anti-realists: the realists should be concerned that their position is hollow, while the anti-realists should worry about a debate which they appear to win by fiat, rather than on the merits of their arguments.

Does this mean that the door is closed, for now, on dark matter realism? Should we just smile politely and change the topic when astrophysicists say things like '[t]he observed dis-

1 To avoid excessive repetition, I will drop the qualifier 'scientific' when talking about scientific realism and anti-realism in the rest of this paper.

2 The main focus of this paper is on developing a normative framework for realism and anti-realism. I wanted to include a case study to illustrate a particular interpretative choice point. I chose the debate over dark matter realism because its central choice point can be articulated without a great deal of technical (mathematical) exposition. I do not intend to suggest that the actual community of cosmologists treats this situation as one of genuine underdetermination; in practice dark matter realists vastly outnumber anti-realists (see e.g. Antoniou [2025]; De Baerdemaeker and Dawid [2022]; Duerr and Wolf [2023]; Vaynberg [2024]. While this fact is, strictly speaking, irrelevant to the use of the case study as an illustration of what is at stake in the debate, I will have more to say about the dominance of dark matter theorists in 5.

3 Vaynberg [2024], in my view correctly, sees Martens' objection as a special case of the 'too little/too much' objection to descriptivism identified by Psillos [2012], p. 224: 'it is the problem of how rich the description should be, given that it is not the *whole* of the theory.' Whereas Psillos' response to this challenge is to propose a hybrid causal-descriptivist strategy, the view I present in this paper abandons descriptivism entirely.

placement between the bulk of the baryons and the gravitational potential proves the presence of dark matter for the most general assumptions regarding the behaviour of gravity' (Clowe et al. [2006])? Should we paraphrase such explicitly descriptive claims in instrumentalist terms that make reference only to astrophysical observational data at extra-galactic scales? I don't think we should. To be clear, my goal in this paper is neither to argue that we should be realists about dark matter nor that we should be anti-realists. Rather, it is to construct a framework within which to articulate both realist and anti-realist views that are worth defending. Martens is right that the traditional approach to scientific realism places dark matter realism atop too semantically thin a foundation. But he is wrong to conclude that this is the best that a realist can do. The culprit is not realism, but rather the traditional articulation of realism. We need an alternative. One that balances the competing desiderata of, on the one hand, doing justice to the often well-justified realist impulses of dark matter theorists, and on the other, delivering a scientific realism worthy of the name.

To find this alternative, we begin with the notion of semantic content. The semantic content of both theoretical and observational scientific claims is best analysed in terms of what I will call 'descriptive deployments' of vocabulary, by which I mean the use of vocabulary elements to make assertions that purport both to pick out entities or structures in the world and to say something about those entities or structures. I will need to set up some semantic machinery before being able to fully spell out some general criteria to demarcate descriptive from non-descriptive deployments—I will do so in 4.2. For now though, I will rely on an example to provide an intuitive characterisation of the distinction. Consider two superficially similar assertions: (i) Cassius and Brutus are in Rome; (ii) Cassius and Brutus are in cahoots. The first sentence descriptively deploys 'Rome', and this is evident from the fact that we can ask follow-up questions like 'where is Rome?'. The second sentence, on the other hand, non-descriptively deploys 'cahoots', as is evident from the fact that it would be treated as infelicitous to ask a follow-up question of the form 'where are the cahoots?'[4]

[4] A description is generally understood as an assertoric utterance, so its natural contrast class consists of non-assertoric utterances such as an imperatives and interrogatives. A descriptive deployment, on the other hand, is not itself an utterance, rather it is a characterisation of the manner in which certain sorts of vocabulary are deployed within an utterance. An assertion often involves a descriptive deployment of some vocabulary, but not all assertions are descriptive. Non-assertoric utterances, such as imperatives, may include also descriptive deployments of vocabularies, for example, as in the deployment of 'window' in 'shut the window!' But they need not: consider the interrogative 'did you do it for her sake?' So the 'descriptive vs. non-descriptive deployment' division is orthogonal to the 'assertoric vs. non-assertoric utterance' division. In this paper, I focus exclusively on descriptive versus non-descriptive deployment, as it pertains to assertoric utterances. This is because, within the scope of the realism debate, there is no disagreement that the sentences at issue, such as 'electrons are negatively charged', are assertions. The realism debate is concerned with how best to understand the content of such assertions, not with whether or not they are assertions in the first place. The distinction between descriptive and non-descriptive deployments is intended to capture this intra-assertion divide. Since the salient contrast class to descriptive assertoric deployments, in this paper, is non-descriptive assertoric deployments, I will drop the qualifier 'assertoric'. I remain uncommitted regarding how best to understand the deployment of vocabulary in non-assertoric utterances such as 'shut the window!'

In the model-theoretic setting characteristic of Analytic Philosophy of Science, descriptive deployment is understood in terms of truth and reference. Call this the alethic–semantic stance:

> The [alethic-] semantic stance takes scientific theories at face-value, seeing them as truth-conditioned descriptions of their intended domain, both observable and unobservable... The theoretical terms featuring in theories have putative factual reference. So, if scientific theories are true, the unobservable entities they posit populate the world (Psillos [2001], p. xvii).

In what follows, I will introduce a normative framework within which to articulate the debate. I will argue that we can make sense of a descriptive deployment in ways that go beyond the alethic–semantic stance. But importantly, in the normative framework, we can still hold on to the alethic–semantic stance as a special case. This fact is worth unpacking. The normative view is not an alternative to an alethic view—it does not do away with truth or reference. It simply asks for an account of how, and under what circumstances, truth and reference play the roles that we take them to, and whether there are alternative notions that play those roles in other circumstances. I offer just such an account in this paper. My account is wholly consistent with an alethic representationalist, compositional semantics for both everyday observational vocabulary as well as some scientific vocabulary (including 'dark matter', for example). What it does not do is take the alethic treatment as the only way to achieve a meaningful, descriptive deployment of a vocabulary.[5] The choice between an alethic and a normative framework sets up three further choice points:

`Choice point 1` : How should we characterise the debate between the realist and anti-realist? I contend that it should be characterised normatively rather than alethically.

`Choice point 2` : What sort of strategy is best placed to deliver an argument for realism or for anti-realism? I contend that a pragmatist strategy is preferable to a non-pragmatist one for both (sets of) positions.

`Choice point 3` : What does a pragmatist strategy look like for scientific realism or for anti-realism? I contend that the arguments will be domain-specific claims made by agents about what authorises the descriptive or non-descriptive use of scientific–theoretical vocabulary.[6] I will demonstrate one such strategy in the context of the expansion scenario for dark matter theorists.

5 As an analogy, suppose I travel to a new country. I would be wrong to suggest that, just because people there do not pay for food in my home currency, that they do not pay for food. I should simply recognise that there are different ways to pay. Similarly, I want to argue here that there are different ways to deploy vocabulary descriptively, and we should not accuse someone of being an anti-realist simply because they do not treat their descriptive deployments alethically.

6 In Magnus and Callender's ([2004]) terminology, we should be looking for 'retail' rather than 'wholesale' arguments for realism or anti-realism.

In 2, I lay the groundwork for the first choice point by discussing the traditional (alethic–semantic) framework for realism and anti-realism. I will highlight that there are two questions that need to be answered in order to stake a position in this debate. To be sure, these are questions whose centrality has been acknowledged by commentators working within the traditional framework. But, as I will argue, the articulation of those questions within that framework is suspect. The two questions are:

`The semantic question` : Under what circumstances, are deployments of scientific–theoretical vocabulary to be understood as descriptive?

`The epistemic question` : Under what circumstances, are descriptive deployments of a scientific–theoretical vocabulary to be understood as vindicated?

In 3, borrowing some machinery from Brandom [2010], I set up the two options — alethic and normative — for the first choice point. A quick terminological clarification: as I use it in this paper (and I make no claims that this characterisation subsumes all legitimate uses of 'normative'), 'normative' means 'concerned with matters of authority' (and, less controversially, 'alethic' means 'concerned with matters of truth'). It is well-established that authority is a normative notion; here I go further, and treat it as the demarcating criterion of normativity.[7] I will argue that both the semantic question and the epistemic question are best understood primarily as questions about authority, rather than primarily as questions about reference or truth.

In the following two sections, I turn to the second choice point. I will argue that, for both the realist and the anti-realist, the most effective strategy available is a pragmatist one. As I use it in this paper (subject to the same caveats as 'normative'), 'pragmatist' means 'concerned with matters of agency.' The discussion will focus specifically on *discursive* agency: (i) what agents need to do in order to be recognised as saying things, and (ii) what, in addition, they need to do in order that those sayings are understood by other agents.

In 4 I address the semantic question, as articulated in normative terms. I introduce Brandom's pragmatist–inferentialist account of assertion, and discuss some of the details of what constitutes a descriptive deployment of vocabulary. I argue that inferentialism is better-suited to the semantic demands of scientific realism than the traditional formulation is, in large part because it allows for the articulation of a compelling account of what it is to descriptively deploy a vocabulary (and thus of what it is to 'take scientific theories at face-value', as Psillos puts it). I will argue that the authority for judging a deployment of theoretical vocabulary as descriptive is spread through the community of discursive agents, and can be systematised in a conversational scorecard.

[7] But I do not want to commit to any further substantive claim about what *constitutes* normativity. A discussion of that sort would take us well beyond the scope of this paper.

In 5, I switch focus to the epistemic question, once again articulated in normative terms. I will argue that there are two distinct sorts of authority to recognise: first, the authority to assess an endeavour as successful (or not), second, the authority to assess, as a reason for that success, a descriptive deployment of some vocabulary.

In 6, I discuss the third choice point, which brings me back to the expansion scenario for dark matter. I argue that the traditional alethic–semantic articulation does not have the resources to capture what is compelling about either a realism or an anti-realism about dark matter, but the normative–pragmatic one does.

2 Realism and anti-realism, alethically

The realist impulse is to take seriously the non-trivial descriptive claims about a perceptually-inaccessible domain of the world made by an appropriately vetted physical theory. The realist models their doxastic attitude towards the putatively unobservable entities in those theories on their doxastic attitude towards everyday observable objects. As Fine puts it:

> [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!).[8] (Fine [1984], p. 35)

But impulses are not philosophical positions. To turn the realist impulse into one, we need a framework. In the traditional alethic–semantic framework, the realist position is the conjunction of three stances (see e.g. Chakravartty [2007]; Ladyman [2012]; Psillos [2001]):

`Alethic-semantic stance` : The semantic treatment of observational terms is the same as the semantic treatment of theoretical terms. In addition, the correct semantic treatment is the alethic (i.e. model-theoretic) treatment we use for observational terms.[9]

`Epistemic stance` : 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.

`Metaphysical stance` : There exists a mind-independent world, and metaphysical status of the observable entities that populate this world is the same as the metaphysical status of the unobservable entities that populate it.

[8] Fine goes on the argue that the acceptance of scientific and everyday evidence constitutes the 'natural ontological attitude' (NOA), and that you do *not* need to be a full-blown scientific realist to be a NOA-er. This paper is neither a response to, nor a commentary on, Fine. But the NOA is a good place to start, since it captures a sound and plausible naturalistic impulse that does motivate many scientific realists, Fine's own arguments notwithstanding.

[9] To bring it more in line with the other two stances presented here, I have paraphrased Psillos' version of alethic–semantic stance that I quoted in 1.

This brings with it multiple anti-realisms, at least one for each of the three stances that is denied. I treat the semantic stance, the epistemic stance, and the relationship between them as being prior to the metaphysical stance because on pain of incoherence, we cannot make metaphysical claims about the objects of a theory if we don't know what the semantic content of that theory is. And, as I will argue, the semantic content of a theory is deeply entangled with the epistemic stance.[10] So we begin with the semantic question introduced in 1 (i.e. under what circumstances, are deployments of scientific–theoretical vocabulary to be understood as descriptive?) as rendered in the alethic framework:

`The semantic question`$_a$: Under what circumstances do scientific–theoretical terms purport to refer?

No reference to authority here, so this is a non-normative rendering of the semantic question. The alethic–semantic stance provides a non-pragmatist answer:

`The alethic-semantic stance answer` : When scientific-theo retical terms (singular terms), and scientific–theoretical properties (predicates) are mapped by the interpretation function to domain elements and sets of domain elements, respectively.

No talk of agents here, so this is a non-pragmatist answer (to an non-normative question).

Let us now turn to the epistemic stance. As with the semantic question, we start by articulating the epistemic question (i.e. under what circumstances are descriptive deployments of a scientific–theoretical vocabulary to be treated as vindicated?) in model-theoretic terms:

`The epistemic question`$_a$: Under what circumstances do we kn ow that (or have good evidence that) scientific–theoretical terms successfully refer?

To which one can give a non-pragmatist answer once again:

`The non-pragmatist epistemic stance answer` : When we know that (or have good evidence that) the scientific theory, of which that vocabulary is a constituent, is true (or approximately true).

In what follows, I will highlight the advantages of the semantic and epistemic questions in terms of authority, rather than in terms of reference and truth. I will also demonstrate that one can still be a realist, insofar as one can accept all three of its constitutive stances (modulo a small modification of the semantic stance).

10 I will not say any more about the metaphysical stance in this paper. In part because, on the pragmatist strategies that I endorse, it is unclear that there is any more to be said about metaphysics over and above the semantics and epistemology we will eventually have in place. But even if you believe there is more to be said, once we have the normative framework for the realism debate in place, it is very easy, if perhaps a little cumbersome, to articulate a metaphysical stance in those terms.

3 The normative framework

So I will not argue against the alethic–semantic stance. As I have already highlighted, there are circumstances in which it is appropriate. Instead, I will try to extol the virtues of a more capacious stance, the normative-semantic stance:

`Normative-semantic stance` : The semantic treatment of observational terms is the same as the semantic treatment of theoretical terms. In addition, the correct semantic treatment for theoretical claims is *whatever* treatment is adjudicated, by the appropriate authority, as correct for observational claims.

This still captures the semantic component of Fine's impulse: theoretical and observational terms are treated in the same way, as required. The observational vocabulary still sets the standard. But the normative semantic stance makes no further stipulation about what the correct semantic treatment is. Brandom [2010] gives us a framework within which to characterise the authority for assessments of the correctness of a the deployment of some base vocabulary: the standards of correctness of a vocabulary deployment, according to whichever authority the semantics answers to, is articulated in a metavocabulary for that base vocabulary.

So, for example, suppose God is our authority, and the standards of correctness of a deployment are understood in terms of consistency with some divine text. We could use a divine metavocabulary to specify all and only the sentences in the base vocabulary that are correct (e.g. 'The cat is on the mat' is correctly deployed if and only if The Book says that the cat is on the mat). On this view, the normative-semantic stance is satisfied just in case The Book contains both theoretical and observational sentences.

A less fanciful example. Suppose we take the unobservable domain of the world to be authoritative in the same way as we do the observable domain of the world. That is to say 'the cat is on my lap' is correctly deployed just in case 'the cat' refers to a particular cat, '...is on...' refers to a 2-place relation between two objects, and so on. Then the normative-semantic stance is satisfied just in case, for example, 'cold dark matter axions have a mass of 5 micro electron-volts' is rendered correct in the same way. This is, of course, just the alethic–semantic stance, but now transposed into a normative key. The difference is one of methodological priority: as I highlighted in 1, we can still trade in truth and reference, as long as they are understood as not being the primary drivers of semantic significance.

The semantic and epistemic questions can be expressed in normative terms, as questions about descriptive deployment:

`The semantic question`$_n$: How, and on whose authority, is a deployment of scientific–theoretical vocabulary characterised as descriptive (i.e. as having the same type of semantic content as observational vocabulary)?

`The epistemic question`$_n$: How, and on whose authority, is the descriptive deployment of some scientific–theoretical vocabulary rendered licit?

The Brandomian framework can now be used to articulate both pragmatist and non-pragmatist answers to them, depending on how the normative metavocabulary specifies standards of correctness. We have already seen two examples of non-pragmatist answers, so let me end this section with a short discussion of pragmatist answers.

On a pragmatist account of content, the semantic content of a sentence is best understood in terms of a metavocabulary that stipulates what an agent must do — i.e. what practice they must engage in, or abilities they must possess — in order to deploy that sentence correctly (contrast this with the disquotational metavocabulary for both the divine command semantics and the alethic semantics introduced above). To keep things simple, I will treat sentences as the primary bearers of semantic content, so that a phrase such as 'the deployment of a theoretical vocabulary' should be read as 'the utterance of a sentence that contains theoretical terms.'[11]

The normative-semantic stance can be given a pragmatic rendering in three stages. First, a discursive agent recognises some other agent(s), part(s) of the world, or collection of data, as authoritative with respect to the assessment of the correctness of their assertions. Next, a metavocabulary is used by that agent to specify how the vocabulary in those correct assertions was deployed. Finally, they assess whether, in this metavocabulary, the specification of theoretical vocabulary deployment is the same as the specification of practices of observational vocabulary deployment. In other words, are agents adjudged as doing the same thing when they deploy these two sorts of vocabulary? For a pragmatist, this exhausts the grounds for semantic content ascription, and so fully characterises the normative-semantic stance.

The combination of broadening the semantic stance and embracing pragmatism immediately yields a significant dividend for the realism debate. On the alethic–semantic stance, the only way for the two sorts of vocabulary, theoretical and observational, to have the same semantic treatment is for them to share what Brandom ([1994]) calls a *formal semantics*: the practice of reporting and systematising the ascriptions of certain properties such as truth, reference, satisfaction that are antecedently demarcated as the semantic ones. For example, the formal semantics of propositional calculus tells us how, given truth-value ascriptions for some base set of propositions, we can calculate the truth-values of complex propositions built, using connectives, from that base set. Contrast this with a *philosophical semantics*, which is collection of practices centred around providing explanations of, or grounds for, the ascriptions of those semantic properties in the first place.[12]

11 See Brandom [1994], Ch. 6 for a discussion of how subsentential expressions such as words and phrases can be ascribed semantic significance.

12 This is a farily standard division in the philosophy of language, although it goes by a number of different names: Burgess and Sherman's ([2014]) 'semantics' vs. ' metasemantics', Dummett's ([1975]) 'theory of meanings' vs. 'theory of meaning' and Stalnaker's ([1997]) 'descriptive semantics' vs. 'foundational seman-

On the pragmatic normative-semantic stance, these two sorts of vocabulary can be judged to have the same *pragmatic significance*. This means that we can, in principle, use different formal semantic resources to capture descriptive deployments in different domains: claims about observables can still be formalised using Tarskian model-theoretic semantics (or its more sophisticated cousins, possible worlds semantics or truthmaker semantics), but we can use alternative formal semantics for unobservables without violating the semantic stance. This is an especially potent line of attack for the realist about (elements of) the highly mathematised theories of modern physics. If such a project could be made to succeed, then there would be no need to try to reformulate mathematised claims into model-theoretic ones in order to argue for realism.[13]

This completes the discussion of the first of the three choice points that I introduced in 1: should we characterise the debate in alethic terms as described at the end of 2, or in normative terms as described here? I have left the door open for the former—as I said at the start of this section, I will not argue against the appropriateness of the alethic–semantic stance in certain circumstances. But I hope to have demonstrated why it is worth pursuing the broader normative-semantic stance instead. For the dark matter realist, for example, the pragmatic normative-semantic stance opens the door to a better way of making sense of descriptive deployment. To see how, we need to address the second choice point: what is the best type of strategy to answer the semantic and the epistemic questions? I address each in turn.

4 The Semantic Question

This section discusses the contours of a pragmatist strategy to answer the normatively-posed semantic question. We begin with Sellars ([1953], [1963b], [1974]), who offers a normative account of meaning: sentences are rendered meaningful by (and only by) their material-inferential relations to other sentences, as authorised by the community of language-users (that these inferential relations are material, rather than formal is crucial—more on this in 4.1.1). Insofar as this account tells us what grounds the ascription of the semantic property of meaning to some semantic interpretants (i.e. sentences), this account qualifies as philosophical semantics. Inferentialism (and other forms of meaning-pragmatism) is entirely consistent with an alethic model-theoretic semantics; inferentialism is a piece of philosophical semantics, while possible-worlds semantics is a piece of formal semantics.

In Brandom's hands ([2009]; [1994]), the material-inferential profile is given a pragmatic regimentation by a conversational scorecard that keeps track of 'the game of giving and asking

tics'

13 There is a good deal more to say about how such a view sits along side structural realist views that take seriously the primacy of the mathematical formalism (e.g. (Bueno and French [2011]; Ladyman et al. [2007]; Wallace [2022])). But that discussion is beyond the scope of this paper.

for reasons'. Recall from 1, that the relevant contrast class to descriptive deployments is non-descriptive deployments. The goal of this section, therefore, is to spell out precisely how Brandom's game characterises descriptive assertions, as a special case of how it characterises assertions. To that end, 4.1 focuses on what constitutes an assertoric utterance in general, and 4.2 focuses on how one might identify the proper subset of assertoric utterances that count as descriptive, in virtue of the descriptive deployment of some vocabulary.

4.1 Assertoric deployment of vocabulary

What must an agent do in order to assert? In this section, I endorse Brandom's ([1983], [2010], [2014]) claim that the ability to make material inferences is a necessary component of any practice or ability that is itself sufficient to assert. A sentence is assertoric if and only if it can serve as a premise or as a conclusion of a material inference. Subspecies of assertions, such as descriptive assertions, will correspond to a more refined set of inferential abilities.

This is decidedly non-standard, so I will try to dispel two sets of concerns that might surface immediately. Readers sympathetic to inferentialism can skip to 4.2. In 4.1.1, I provide an account of semantic intentionality. Finally, in 4.1.2, I demonstrate how the inferentialist picture fits into the broader normative–pragmatist picture introduced in 3.

4.1.1 Tethering sentences to the world

On the formalist view of inferences, the correctness of an inference is purely a matter of logical form. The inference from 'there is a rattlesnake in your path' to 'there is a reptile in your path' is valid only because it is an instance a schema with a suppressed premise ('all rattlesnakes are reptiles'). But, to know whether an inference is good in virtue of its form, but not its content, we need to already know what its content is. So any attempts to ground meaning in formal inferences are circular.

Which is why the inferentialist grounds meaning in *material*, and not formal, inferences. A material inference is just an ordered tuple of premises and conclusions that is treated as primitively valid by a community of discursive agents. Ultimately, the decision to treat such an inference as valid is vindicated by the extent to which it is a constituent of a successful discursive practice (more on this in 5).

The primacy of the material inference grounds the inferentialist's account of intentionality—i.e. the directedness of thoughts, beliefs, words and sentences towards the world. For the inferentialist, intentionality is derived from a practical involvement of agents with their surroundings. And a crucial component of this involvement is that it is dynamic—the environment responds to the actions of the agents who, in turn, respond to the changes in surroundings. There is no mystery about the mechanics of such practical intentionality—for example, I came

to learn pretty quickly that I should move away when I see a rattlesnake or duck when I see a cricket ball hurtling towards my head.

On this view semantic intentionality is derivative: the fact that some agents engage in discursive practices, among other practices, allows for practical intentionality to filter through to specific practices and performances, including discursive ones, in specific ways. There is a non-discursive practical intentionality in my understanding how to deal with snake-dodging. Now, when I want to recount my snake-dodging experience to someone, or when I want to warn someone about a nearby snake, that is when I engage in a new performance—a discursive one. I express a judgement, using a sentence, whose content is governed by the commitments and entitlements to other judgements that are materially-inferentially linked. Those judgements are about some part of the world precisely because my ability to deploy the vocabulary that expresses those judgements is a special case of my practical involvement with my surroundings (which includes my interlocutors). The semantic intentionality exhibited by my words is a special case of this thick, dynamical, interactive practical intentionality that constitutes my agency in the world. This is how the discursive practice remains tethered to the world: semantic intentionality, on this view, can be seen as a subspecies of practical intentionality (this is a line of argument pursued further in McDowell [1996], and endorsed in Brandom [2010]).

4.1.2 Making things explicit

Our ability to make the inference from 'there is a rattlesnake in your path' to 'there is a reptile in your path' is an instance of the broader abilities we possess to assertorically deploy a biological vocabulary V_b. Part of what 'rattlesnake' means, then, is constituted by its contribution to this and other material inferences. Our endorsement of this material inference reflects a commitment to endorsing one sentence expressed in V_b, namely 'there is a reptile in your path' whenever we endorse another sentence expressed in V_b, namely 'there is a rattlesnake in your path'. This is, of course, easily recognisable as an instance of our ability to deploy a conditional using V_b.

But V_b does not itself allow us to characterise what we did as deploying a conditional in V_b.' To do so, we need a pragmatic metavocabulary of conditionals. Using a (pragmatic) logical metavocabulary, call it V_L, we can say what it was we did when we endorsed the propriety of the material inference from one sentence in V_b to another: we in fact endorsed a conditional of the form if p then q. The logical metavocabulary thus specifies (a particular instance of) an ability: the ability to make a particular sort of inference in the biological vocabulary. In Brandom's terminology, the conditional *made explicit* (or 'explicitated') what was implicit in the endorsement of the material inference. This 'explicitation using the conditional' relation between the base vocabulary V_b and the metavocabulary V_L is an instance of a general feature

of Brandom's account: the metavocabulary makes explicit the norms that are implicit in the practice of deploying a base vocabulary (in this case, V_b) in order to assert.

Different elements of the metavocabulary make explicit different aspects of the practice. Negations, as we will see below, make explicit a practical commitment to an incompatibility. But this idea extends far beyond logical vocabulary. The semantic vocabulary of 'representation' can also be understood as an element of an explicitating metavocabulary for a particular set of norms that are implicit in the deployment of a base vocabulary. By identifying those norms, we can identify the sub-practice of the practice of asserting that characterises the specifically *descriptive* deployment of elements of a vocabulary.

4.2 Descriptive deployment of vocabulary

The examples above demonstrated how the logical metavocbulary of conditionals can be understood as explicitations of some of the norms that govern the practice of asserting. This idea can be extended to other sorts of metavocabulary.[14]

Most explicitations are much harder to characterise than indicative conditionals. So in order to keep track of the how subtler moves in the game of giving and asking for reasons contribute to the semantic significance of assertoric discursive performances, Brandom introduces the notion of a *deontic* scorecard. This scorecard allows speakers to keep track of dynamics of their interlocutors' normative statuses: commitments and entitlements to particular assertions.[15]

Normative statuses are only part of this account of discursive practice. The other part is the collection of normative attitudes that interlocutors acknowledge of themselves, and attribute to each other. The subtle dynamical interplay of attitudes and statuses is what, on this account, characterises the conceptual content of assertoric performances. By contrast with, for example, the traditional Lockean view of assertion as an attempt to relate something internal, such as a thought about an object, to an interlocutor, Brandom's view treats, as the fundamental driver of semantic significance, the responsibility incurred by a speaker to their interlocutor(s) in endorsing a sentence, and the subsequent normative hold exerted by that responsibility. This is what the deontic scorecard keeps track of. A conversation is a rule-governed activity. These rules, which in this section, I treat as primitive features of a community's practices,[16] govern the manner in which speakers take responsibility for the commitments and entitlements that

14 In the appendix, I use a simple example of a conversation about cricket between two agents, A and B, in order to demonstrate how metavocabularies can characterise two further sorts of implicit commitments: negations can characterise commitments to incompatibilities, and representation can characterise commitments to differences in social perspective between agents. The example deals only in sentences, and their inferential-semantic significance.

15 The scorecard is *deontic* because commitments and entitlements are the discursive renderings of the broader deontic categories of obligations and permissions. The deontic scorecard is the elaboration, in the Lewisian ([1979]) idiom, of the notion of endorsement that Sellars puts to work in Sellars [1963a].

16 But see Brandom [1994], Ch. 1 for an discussion of where those rules come from.

emanate from their assertions, and apportion this responsibility to each other.

According to Brandom, the coordination of the perspectives of different interlocutors is what is made explicit by representation-talk. The deontic scorecard then encodes *de re* propositional attitude ascriptions from A to B in terms of how assertions update the following attitudes about normative statuses: commitments attributed to B by A, entitlements attributed to B by A, and commitments acknowledged by A as being undertaken by them. When B uses the metavocbulary of negation, they make explicit their endorsement of the material incompatibility of commitments. When they use indicative conditionals, they make explicit their endorsement of certain material consequence relations. And when they make a *de re* attitude ascription to A, they make explicit the differences in social perspective between themselves and A, and express this difference in terms of the differences between the commitments attributed to A, and those acknowledged by B.[17] In the next two sections, I will link this characterisation of representation-talk — as an explicitating metavocabulary — to the notion of a descriptive deployment, in order to characterise, in normative terms, the realist and anti-realist positions.

But first, a short recapitulation. I began this section with two goals. The first was to offer a pragmatist account of assertion. I introduced and endorsed a Brandomian inferentialist account, according to which a sentence is assertoric if and only if it can justify and be justified by other sentences, by the standards of the material inferential relations in which these sentences are taken to stand by a community of speakers. I dispelled two worries that appear to stand in the way of such an account: (i) that the notion of inference cannot be prior to semantic content and (ii) that the inferential structure of the space of sentences cannot account for how sentences come to be about the world

The second goal was to offer a pragmatist account of descriptive assertion. Once again, I offered a Brandomian account, according to which a descriptive deployment of a vocabulary is achieved by uttering sentences that have the pragmatic significance of making explicit differences in social perspectives between interlocutors, in a manner captured by a specific sort of deontic scorecard dynamics. Equipped with this pragmatist–inferentialist account of descriptive deployment, we can formulate a pragmatist response to the normatively-posed semantic question.

`The semantic question`$_n$: How, and on whose authority, is a deployment of scientific–theoretical vocabulary assessed as descriptive (i.e. as having the same semantic status as observational vocabulary)?

`The pragmatist-inferentialist answer` : The assessment of a deployment as descriptive is authorised by whomever is keeping deontic score of an agent's use of such vo-

[17] For a defence of this proposal for how to understand *de re* attitude ascriptions, see Brandom [2009]; Brandom [1994], [2019]; Price et al. [2013].

cabulary. It is authorised as such when the scorekeeper attributes to that agent a *de re* propositional attitude, characterised by the recording of the dynamics of normative statuses and attitudes that reflects the difference in social perspective between that agent and the scorekeeper. In conversational English, this is regimented in the use of 'of...that...' locutions in a semantic metavocabulary, but other discursive practices will, in general, encode it in other ways. For example, if *S* utters a claim like 'the WIMPs (weakly-interacting massive particle) that constitute a cluster of dark matter are supersymmetric', it is treated as descriptive on the authority of an agent when they attribute to *S* a *de re* propositional attitude, characterised by the deontic scorecard dynamics described above.

In the next section, we will discuss what sort of authority we need to invoke in order to vindicate this deployment, i.e. to assess as correct the agent's choice to deploy a vocabulary descriptively. In general, the mechanisms of these two authorisations will be distinct, even if the sites of those two types of authority are the same, i.e. a particular agent or community.

5 The Epistemic Question

The scientific realist needs to establish more than just the claim that vocabulary is being used to represent. After all, representation-talk, understood purely in terms of *de re* attitude ascriptions, is cheap: sentences involving 'unicorn', '5' and '{Socrates}' just as easily satisfy the *de re* test described above. So there is a gap between what we might call *merely* descriptive deployment, and a thicker notion of descriptive deployment that, intuitively speaking, picks out some real-worldly entities, even if it fails to attribute the correct properties to those entities.[18],[-0.01pt][19]

In this section, I argue that the requisite thickness comes from two additional normative considerations around the vindication of practices. The first is the authority to vindicate an endeavour. The second is the authority to assess the role that a particular sort of vocabulary deployment plays in the success of that endeavour. I present my case with the help of two examples, one from the everyday context 5.1, the other from the scientific (dark matter) context 5.2.

18 This is a well-trodden issue in the pragmatist literature, often discussed under the heading of the 'bifurcation thesis' (Kraut [1990]). According to the bifurcationist, our vocabulary splits our sentences into those which should be given a descriptivist reading and those which should be given an anti-descriptivist (e.g. expressivist) reading. The account that I provide in this paper is not the only proposal on the table for how to make this division. Other proposals can be found in Blackburn [1993] and Rorty [2009].

19 This is the point at which Psillos [2012] supplements his descriptivism with an alethic–semantic causal-historical referent-fixer. I propose a normative–pragmatic alternative.

5.1 The pragmatic vindication of descriptive deployment

The pragmatic vindication of the descriptive deployment of some vocabulary as part of a successful endeavour is a two-stage process. The first is the vindication of the endeavour itself: success is assessed with respect to some standard of authority. The second is the triangulation of the load-bearing components of the discursive component of the successful endeavour. This triangulation might vindicate a descriptive deployment of some vocabulary, or it might not, depending once again, on what sort of authority is recognised, and what that authority adjudicates. At each stage, however, a different sort of authority is recognised. The pragmatist answer to the epistemic question takes both sorts of authority into account.

To illustrate the two sorts of authority, let us begin with an example from Wittgenstein:

> I send someone shopping. I give him a slip of paper marked "five red apples". He takes the slip to the shopkeeper who opens a drawer marked "apples"; then he looks up the word red in a chart and finds a colour sample next to it; then he says the series of elementary number-words... up to the word "five", and for each number-word he takes an apple of the same colour as the sample out of the drawer. (Wittgenstein [2010], §1)

The assistant, who is sent to the shop, engages in a series of actions, with the aim of acquiring apples for his boss. A discursive action — for example, an utterance to the shopkeeper of something like 'five red apples, please', or simply the reading of the slip of paper by the shopkeeper — results in the shopkeeper, in turn, engaging in a series of actions. In sum, the effect of all of these actions is the successful achievement of the assistant's goal: to have their actions of apple-acquisition vindicated by their boss, who is recognised by the assistant as having the authority to vindicate. In particular, the boss issues a vindication of what the assistant did to end up with the apples. The boss' authority *qua* endeavour-vindicator is recognised by the assistant, and that authority is constituted entirely by this recognition. But that authority ends at assessing whether or not the apple-buying endeavour was successful. In doing so, the boss need not voice a further opinion about which, if any, of the individual actions that constituted the assistant's endeavour (including the descriptive deployment of 'apples') is responsible for its success. Of course, the boss might have an opinion, but that opinion is not automatically authoritative in vindicating the descriptive deployment of the apple-vocabulary.

So there is a further question of authority here: who or what authorises the assessment of a descriptive deployment of vocabulary *as vindicated* once (i) the community of speakers has adjudged the deployment of vocabulary as descriptive (via their appropriate updates to their deontic scorecards) and (ii) the endeavour vindicator (in this case the boss) has assessed the endeavour as successful?

We can easily imagine the apple-buying mission as an instance of a fairly common sort of occurrence in the boss–assistant dynamic: that the assistant is asked to complete various tasks for the boss. In his capacity as a competent assistant, *as judged by the boss*, the assistant ties the success of his various endeavours to the descriptive deployment of certain vocabulary: the apple-acquisition, which included the descriptive deployment of apple-vocabulary, was deemed successful. So too, *mutatis mutandis* was the laundry-acquisition, the ticket-acquisition, the present-acquisition, and so on. These tasks, it turns out, were doable, but with considerably less efficiency and lower success-rates when vocabulary was not descriptively deployed (for example, when the assistant had a sore throat, so merely pointed at things instead of making utterances). And this went beyond mere acquisition tasks: repairs, for example, were successful because of other sorts of descriptive deployments, and unsuccessful or severely hindered by the decision not to deploy that vocabulary descriptively. Similarly, various types of organisational tasks.

What this suggests is that the authority to assess the *reasons* for the success of endeavours adjudicated as successful by the boss is, in the first instance, the assistant himself. The vindication of the descriptive deployment of apple-vocabulary comes from the assessment, at the hands of a skilled practitioner (whose skill is recognised in terms of their success as assessed by what or whomever they take to be authoritative) that the descriptive deployment of apple-vocabulary is the reason for their apple acquisition success. Now, this assessment is defeasible—that the assistant is an expert at apple acquisition does not mean that he is also an expert at understanding why he is an expert at acquiring apples. But this is a feature, not a bug, of the normative account. Just as the assistant might have been wrong to treat his boss as an authority with respect to apple-acquiring success (the boss might not have known the difference between an apple and a peach), he might also be wrong to treat himself as an authority with respect to the assessment of the reasons for his success. But all that means is that we need to look elsewhere for a better authority. Just as an athlete might look to hire a better coach to identify which aspect of their skills are responsible for their success, and how to improve them, the assistant might authorise someone else — another colleague, a friend, a mentor, or even his boss — to identify the reasons for his success and failures.

On the basis of this example, we can articulate a pragmatist answer to a very restricted version of the epistemic question:

`An epistemic question` : How, and on whose authority, is the descriptive deployment of apple-vocabulary rendered licit?

`The pragmatist answer` : There are two relevant sources of authority here. The first is the authority to vindicate an apple-acquisition endeavour. This authority is bestowed upon the assistant's boss by his recognition of her as authoritative. When she assesses as suc-

cessful an apple-acquisition endeavour, there is a (usually distinct) authority to assess the descriptive deployment of apple-vocabulary as the reason for that success; that authority might be the assistant himself. The interaction of these two sorts of authorisation is what renders licit the descriptive deployment of apple-vocabulary.

5.2 The pragmatic vindication of theoretical talk: dark matter realism

This dynamic can be extended, quite instructively, to the scientific domain. Let us return to the dark matter case study that began this paper. Martens and Lehmkuhl [2020] summarise the situation in terms of an empirical inconsistency:

`Gravitational laws` : Our best theory of gravitation, general relativity, is (approximately) empirically adequate.

`Luminous matter` : Most of the matter in the universe is luminous baryonic matter, i.e. it interacts with light.

Both `gravitational laws` and `luminous matter` are, separately, on solid empirical footing, but their conjunction is rendered false by astrophysical observations. Martens and Lehmkuhl call this the 'dark discrepancy.' Predictably, two communities have arisen in response, each characterised by the premise they deny. Dark matter theorists deny `luminous matter`, and attempt to build a theory of non-luminous matter which, in conjunction with unmodified general relativity, is empirically adequate. For them, the dark discrepancy points to another expansion scenario. By contrast, modified gravity theorists deny `gravitational laws`, and attempt to build a new theory of gravitation which, in conjunction with an unmodified theory of luminous matter, is empirically adequate. For them, the dark discrepancy is not evidence of an expansion scenario: on their view, no further material entities should be added to our world picture.

Recent philosophical literature has focussed primarily on assessing the pursuit-worthiness of the dark matter research program over the modified gravity program (see e.g. Duerr and Wolf [2023]), and accounting for the relative dominance of the former (see e.g. Antoniou [2025]; Vaynberg [2024]). But it is important to distinguish a commitment to the pursuit-worthiness of a research program or physical theory from (i) a commitment to *accepting*[20] a theory and (ii) the even stronger commitment of *realism* with respect to that theory's central terms. One might be overwhelmingly in favour of pursuing the string theory research program, for example, even if one does not accept it. Further, even if one accepts a theory, as is the default with standard

[20] Van Fraassen [1980], p. 88: '[t]o accept a theory is to make a commitment to the further confrontation of new phenomena within the framework of that theory, a commitment to a research programme, and a wager that all relevant phenomena can be accounted for without giving up that theory.'

model of particle physics, one can still maintain a staunch antirealism about, say, hadrons. All of this is to say that the minimal commitment to the pursuitworthiness, or even the stronger commitment to the acceptability of a dark matter theory need not entail a commitment to dark matter realism.

Nevertheless, in the case of the dark matter research program, it turns out that arguments in favour of its pursuit-worthiness are strongly tied to arguments for dark matter realism (see e.g. Martens [2021]), to the extent that an argument in favour of pursuing dark matter physics is tacitly taken to be an argument in favour of dark matter realism. This observation can be given a clean and precise rendering in the normative terms that I have been arguing for in this paper, whilst still leaving open the possibility of a position on which dark matter theories are deemed pursuit-worthy, or even acceptable, by a dark matter anti-realist. The key is in recognising that the same sort of data can be treated as authoritative with respect to different considerations, depending on various other commitments of the authorising agent or community.

Here is a simplified but helpful take on the present observational situation regarding dark matter: dark matter theories have had more success in accounting for observational data relating to very large scales — at the level of clusters of multiple galaxies (De Swart et al. [2017]; Jacquart [2021]; Trimble [1987]) while modified gravity theories have tended to be better at accounting for observational data at relatively smaller scales — to the order of individual galaxies (Famaey and McGaugh [2012]; McGaugh [2020]; Merritt [2020]). The community of dark matter realists, accordingly, assesses that the descriptive deployment of 'dark matter' and its cognates is the reason for dark matter theories' success in the larger-scale domains. This reason is treated by the community as sufficiently robust as to motivate the development of a (descriptive) dark matter theory to account for observational data at smaller galactic scales, *despite* the fact that some prominent modified gravity theories have been more successful at accounting for observational data at that scale. The descriptive authority of large-scale data underpins a methodological injunction to develop a theory that subsumes the (currently) recalcitrant small-scale galactic data into a dark matter theory.[21]

By contrast, on the authority of the community of modified gravity theorists, the rotation curves of individual galaxies are treated as the primary source of (*non*-)descriptive authority. Just as with the dark matter realist, the modified gravity theorist takes their choice of descrip-

[21] In terms of the grammar of the *de re* attitude ascriptions discussed in 5.2, a dark matter realist is someone who attributes to their interlocutor various doxastic attitudes about the entities *about which they are themselves realists*. This is an important point to clarify: suppose A includes in their deontic scorecard, 'B believes of dark matter that it is non-baryonic', the person whom we can identify as a dark matter realist is A, not B. This is because it is A, and not necessarily B, who is deploying 'dark matter' descriptively, as part of the explicitation of one of their own acknowledged commitments. Suppose, instead that C includes in their scorecard, 'B believes of dark matter, which does not exist, that it is non-baryonic.' In this case, as with the Milgrom quote above, we can identify C's acknowledged commitments by their explicit non-descriptive deployment of 'dark matter.' This aspect should not be surprising: our primary evidence for what someone's commitments are should come from what they say, not what others say about them.

tive authority to underpin an injunction to develop a theory that is empirically adequate on all scales. And a load-bearing component of this practice is the *non*-descriptive deployment of 'dark matter'. For example, according to Milgrom, an advocate of Modified Newtonian Dynamics (MOND) (which is modified-gravity research programme), 'MOND...posits that the 'dark matter' of the dark matter paradigm does not exist' (Milgrom [2020]).

This does not mean that dark matter theorists *disregard*, or in any sense treat as dispensable, the data from smaller-scale individual galaxy rotation curves, any more than MONDians and other anti-realists disregard large-scale data. Where the camps do differ is over how they understand the authority of different sorts of data: dark matter theorists take the descriptive authority of large-scale data to motivate the development of an empirically adequate theory whose central terms are descriptively deployed, *even on a small scale* rather than to allow the existence of a modified gravity theory which is successful on those smaller scales to vindicate the non-descriptive deployment of the dark matter theorist's central terms. By contrast, the anti-realist about dark matter treats small-scale galactic data as authoritative in precisely the opposite way: as vindicating a non-descriptive deployment of 'dark matter' (or, better, as vindicating the expunging of 'dark matter' from our scientific vocabulary), and motivating the development of an alternative theory to dark matter, one which accounts for large-scale observations without positing an expansion scenario.

This raises two important, and related questions: (a) why does the community of dark matter theorists treat the descriptive deployment of 'dark matter' as sufficiently robust a reason for the success of their programme (in its successful domains), that they are willing to commit to the development of a theory that continues to descriptively deploy 'dark matter' to account for observational data at the galactic scale, despite the existence of an in-principle superior theory at that scale? and (b) if the decision is really only a question about which sort of data to recognise as descriptively authoritative, then why does the community of astrophysicists skew so heavily in favour of dark matter theorists? These are ultimately questions for historians and perhaps sociologists of physics, but let me offer some brief remarks on how they might be addressed within the normative framework of this paper.

In the absence of further normative considerations, there is no reason to choose the large-scale data as authoritative over the small-scale data, with respect to theory-development. So the reason for the big discrepancy between dark matter realists and anti-realists must come down to a disagreement over which normative considerations to accept. According to Antoniou [2025], the main reasons that the dark matter paradigm has prevailed so far, at least with respect to pursuit-worthiness are (i) its problem solving potential (ii) its compatibility with established theories and (iii) its independent testability. Antoniou does not venture a hypothesis about reasons why theorists might endorse the stronger commitment to dark matter *realism*, but it isn't hard to extend his analysis to provide one. Note that in what follows, I am not endorsing

dark matter realism, but instead trying to make sense, using the normative framework, of the reasons why one might be a dark matter realist.

The compatibility of dark matter theories with established physical theories means that dark matter realism does less damage to the extant normative-inferential structure of discourse in contemporary astrophysics than anti-realism. To be a realist about dark matter is to add some new elements of theoretical vocabulary to the extant physical vocabulary, and specify their rules for deployment. A rule that specifies that the new terms are to be descriptively deployed allows for the inferential significance of well-established vocabulary — such as 'spacetime', 'curvature', 'matter' — to remain largely unchanged. Contrast this with a form of dark matter anti-realism such as MOND, on which, for instance, the functional relationships between these, and countless other, terms are significantly updated. An analogy at this point might help: recall the example from 1 of 'Brutus and Cassius are in Rome'. Suppose we want to describe the introduction of 'cahoots' into the discourse around Brutus and Cassius. A cahoots-anti-realist will suggest that 'cahoots' be treated non-descriptively. But in doing so, they change the meaning of the well-established term 'in': it now carries a non-spatiotemporal significance that it did not earlier. Of course, in the case of cahoots, this is the correct thing to do—on the authority of virtually all English-speakers, our practice vindicates the non-descriptive use of 'cahoots' and its associated updating of the meaning of 'in'. Similarly, on the authority of dark matter realists, one might argue that the practice of astrophysics vindicates the descriptive deployment of 'dark matter'.

With this in place, we can see how the pragmatist can answer the epistemic question regarding scientific–theoretical vocabulary in general:

`The epistemic question`$_n$: How, and on whose authority, is the descriptive deployment of some scientific–theoretical vocabulary (e.g. 'dark matter') rendered licit?

`The pragmatist answer` : There are two relevant sources of authority here. The first is the authority to vindicate an endeavour or a practice. This authority is bestowed upon some part of the world by its being recognised as authoritative by some agent(s). For dark matter theorists, it is data reporting large-scale cosmological behaviour; for modified gravity theorists it is data reporting smaller-scale galaxy rotation behaviour. When this authority assesses as successful an endeavour or practice, there is a (usually distinct) authority to assess the descriptive deployment of certain vocabulary as the reason for that success. The interaction dynamics of these two sorts of authorisation is what renders licit the use of some scientific–theoretical vocabulary.

6 Realism and anti-realism, normatively

We can now address the third choice point, and characterise the positions of scientific realism and anti-realism, in terms of how they employ pragmatist strategies to answer the semantic and epistemic questions:

> A scientific realist, with respect to some scientific (theoretical) vocabulary, is an agent who takes the descriptive deployment of that vocabulary to be rendered normatively licit by the role such a deployment plays in accounting for the success of a theory that uses that vocabulary, where that success is adjudicated by what or whomever is recognised, by that agent, as authoritative.

This brings with it a conjugate characterisation of scientific anti-realism:

> A scientific anti-realist, with respect to some scientific (theoretical) vocabulary, is an agent who takes the descriptive deployment of that vocabulary not to be rendered normatively licit by the role such a deployment plays in accounting for the success of a theory that uses that vocabulary, where that success is adjudicated by what or whomever is recognised, by that agent, as authoritative.[22]

One might worry that all this talk of authority leads to an unpalatable form of relativism. in the sense that there is nothing that vindicates my choice to treat as authoritative, say, the community of astrophysicists over some else's choice to treat as authoritative, say, the community of internet conspiracy theorists. If this were the case, then the normative account would decidedly fail to capture at least the scientific realist impulse.

But the scope of the scientific realism debate renders this sort of relativism implausible. On the traditional view as expressed by e.g. Van Fraassen [1980], both the realist and anti-realist are aligned on the goal of empirical adequacy, i.e. of developing predictively successful theories. The realist just has an additional goal: truth about unobservables. But the sources of candidates for those unobservables are already heavily constrained by the fact that we are looking at scientific theories. In other words, the scientific realist's opponent is the *scientific* anti-realist, not the sceptic or nihilist about science itself. And this scope-restriction applies just as well to the normative approach. Note how the candidate anti-realist in the astrophysical example was also a proponent of an astrophysical theory: MOND. Their justification for denying the descriptive deployment of 'dark matter' was the existence of a *positive* account of the behaviour of some observables. Such dark matter anti-realism is just as much a wager about the

[22] This immediately allows us to characterise three versions of anti-realism: (i) one which reserves judgement on whether the descriptive deployment is rendered normatively licit, (ii) one according to which the non-descriptive deployment is rendered normatively licit, and (iii) one according to which the descriptive deployment is rendered normatively illicit.

best way to develop science as dark matter realism. So the slide into relativism is arrested by the broad consensus on minimal success conditions for science. And this is to the credit of the normative account: it does not provide arguments for realism, it simply provides a framework within which to confront them with arguments for anti-realism.

Let me end with by highlighting one final advantage of the normative articulation over the traditional alethic articulation of the debate: that it addresses Martens' semantic thinness worry. Recall his claim that whatever practitioners in different dark matter research programs have in common 'is so semantically thin, so vacuous, that it barely means anything at all.' (Martens [2021], p. 6). This is a terrific illustration of how the alethic framework stands in the way of capturing an important realist impulse. If the only way 'dark matter' can be meaningful is if 'dark matter' successfully refers and, moreover, we know what it successfully refers to and, moreover we know enough about the properties of its referents that we can characterise dark matter in terms of those properties, then Martens is correct: no one is, or should be, a dark matter realist. Not even the dark matter theorists for whom a descriptive deployment of 'dark matter' is part of the semantic structure of their theory. All that is left, on this framework, is the thin gruel that Martens calls 'not-yet-realism'. This strikes me as a reductio on the alethic framework: it has simply failed to capture the realist impulse. On this view, claims such as '[s]ince the confirmation of dark matter's existence, a preponderance of dark matter in galaxies and clusters of galaxies has been discerned through the phenomenon of gravitational lensing...' (Riess [2024]) (even when they come from Nobel-prize winning astrophysicists!) sound lukewarm at best and flatly incorrect at worst.

By contrast, on the normative view, dark matter theorists in very different research programmes can agree to substantive realist commitments (*mutatis mutandis* anti-realist commitments). They know what they mean by 'dark matter', via its inferential profile in the theories that constitute their research programme. But even though they may all mean something quite different from each other, they can still assent to the core dark matter realist commitment: the descriptive deployment of 'dark matter.' And this is on solid footing because of what they all, pragmatically, take to be authoritative with respect to the successes of their own individual research programmes. Dark matter realism expresses a defeasible endorsement of an assessment of the role played by a particular semantic treatment of an aspect of the practices that constitute different research programme. Martens correctly highlights the unilluminatingly thin common formal semantic core. But the dark matter realist does not require a common formal semantic core; they require a sufficiently thick common *philosophical semantic* core. On the normative–pragmatist account, they have one. Meanwhile the dark matter anti-realist no longer wins by fiat, but instead, if they do, it is on the strength of their own arguments against that common philosophical semantic core.

7 Conclusion

In this paper, I looked to defend three claims: (i) the debate between the scientific realist and anti-realist should be articulated in normative terms, where questions about normativity are questions about authority; (ii) positions in this debate should be defended using pragmatist strategies, where pragmatist strategies are agent-first strategies; and (iii) the specific pragmatist strategy available with respect to some scientific–theoretical vocabulary is highly domain-specific and turns on choices made by agents about what sort of authority to recognise.

I urged that we break away from the traditional alethic–semantic articulation of the debate because of the benefits of adopting a more capacious normative framework. Within that framework, questions about realism can be approached in terms of (i) what authorises a uniformity of semantic treatment across theoretical and observational vocabularies in successful theories (the semantic question) and (ii) what authorises the assessment of a theory as successful (the epistemic question). I argued that such questions of authority are best answered in terms of the recognition of authority by agents—a pragmatist answer to normative questions. I framed these claims around the debate over whether or not we should be scientific realists about dark matter. I argued that the normative rendering of the debate, together with a pragmatist strategy for engaging in that debate, brings with it significant expressive and dialectical benefits.

Even if I have not managed to convince the reader of the value of a thoroughgoing pragmatism in response to the two central questions of scientific realism, I hope to have at least provided a clearer map of the territory than was previously available. In particular, one that allows us to characterise various discussants in the debate in terms of both how they understand what is at stake in the debate, as well as what they take to be the best strategy for arguing for their position. Sellars [1968], [1976], for example, is a realist, for whom the central questions are both normative, but for whom the answers are non-pragmatic (or at least not wholly pragmatic): the standards for normative success are understood in terms of picturing homomorphisms between sign designs and parts of the world. According to Hacking [1983], to be a realist is to give pragmatist answers, in terms of interventions, to non-normative questions. Van Fraassen [1980], [2001], by contrast, is an anti-realist for whom the questions are both non-normative, and answers non-pragmatic. There are, of course, other permutations to explore.

A An example of a deontic scorecard

Consider a conversation about cricket between two agents, A and B. Suppose A's opening salvo is the claim P_1: India won the most recent (men's) cricket world cup. B now attributes a commitment to A, in the sense that they impute to A a responsibility to provide reasons for their assertion of P_1. And this responsibility is acknowledged by A. B also (defeasibly) attributes to A a slew of other commitments that B takes to follow from P_1. Many of these

claims, however, are not acknowledged by A as one of their commitments, so will be recorded by B as unacknowledged by A.

In addition to commitments, an assertion such as A's brings with it certain entitlements, some of which are acknowledged, and others not. These are sentences which, on the basis of A's claim, B takes A to be permitted to endorse, by the discursive norms by which they both take themselves to be bound. And this situation is symmetric: both A and B have scorecards on which they keep track of each other's, and their own, commitments and entitlements, acknowledged and unacknowledged. What happens when these scorecards conflict?

Suppose B acknowledges A's commitment to P_1, but withholds the attribution to A of an entitlement to P_1. Instead B issues a challenge. They do so by asserting X_1: Australia won the most recent (men's) cricket world cup. From B's perspective, this is a challenge, because X_1 and P_1 are taken, by B, to be incompatible. But this incompatibility is only obvious because one of the material inferential consequences of X_1 that B accepts (together with the background commitment that the World Cup was not shared) is X_2: It is not the case that India won the most recent cricket world cup.

But A might soon understand that B is, in fact, issuing a challenge. Either because A recognises that X_2 is a material consequence of X_1, or because B makes their commitment to X_2 explicit by asserting it. At this point X_2 becomes an acknowledged commitment of B's. The logical vocabulary of negation is an example of the sort of explicitation introduced in the last section: B was always committed to X_2 when they asserted X_1, because they treated the transition as licit by the standards of the discursive norms. Then, in order to make explicit a commitment to the incompatibility of X_2 and P_1, B employed the vocabulary of negation, which allowed them to characterise X_2 as $\neg P_1$.

In response, A might choose to challenge B, or to withdraw or modify their assertion. As it happens, A sees exactly what's going on, and modifies P_1 to P_1': India won the most recent (men's) *T20* cricket world cup. This allows for a disambiguation between A's use of 'cricket' to mean '*T20* cricket' and B's use to mean '*one-day* cricket'. A's original claim, P_1, secretly included an unacknowledged commitment which can now be made explicit by A's endorsement of P_1': India won the most recent (men's) T20 cricket world cup. And this is entirely compatible with B's revised claim, X_1': Australia won the most recent (men's) one-day cricket world cup.

So B's full deontic scorecard for A should include their own commitments and entitlements, acknowledged and unacknowledged, as well as B's

B's Commitments		B's Entitlements	
Acknowledged	Unacknowledged	Acknowledged	Unacknowledged
X_1', X_2'		X_1'	Q_i
A's own commitments		A's own entitlements	
Acknowledged	Unacknowledged	Acknowledged	Unacknowledged
P_1'		P_1'	Q_i

The deontic scorecard is thus an amalgam of normative statuses and attitudes, and in keeping track of how successive performances updates these statuses and attitudes, a picture starts to emerge of the important moving parts of the Brandomian account of meaning. It takes norm-governed conversational dynamics as semantically significant: the meaning of a claim is nothing more than its inferential profile, encoded in the dynamics of challenges and defences of the statuses of commitments and entitlements, and the attitudes of attributions and acknowledgements.

This example highlights the utility of two sorts of metavocabulary. The first is logical, in particular, negation. B was always committed to X_2 when they asserted X_1, because they treated the transition as licit by the standards of the discursive norms. Then, in order to make explicit a commitment to the incompatibility of X_2 and P_1, B employed the logical metavocabulary of negation, which allowed them to characterise X_2 as $\neg P_1$ (the syntax makes it look like the '$\neg$' is part of the base cricket vocabulary, but it is not).

The second is the semantic metavocabulary of representation. The semantic content of P_1 can be rendered in terms of the metavocabulary of representation: A represents some entity (India, or more accurately, the Indian cricket team) as having some property. The dynamics of the ensuing disagreement demonstrated where the two interlocutors were misaligned: they each represented an entity as having a particular property, and each attributed to the other a representation incompatible with the one they each acknowledged. Call the attribution of a normative attitude of this form a *de re* attitude ascription. Paradigmatically in English, these are picked out using 'of...that...', as in 'A claims of India that they won the world cup.' The claim about what A was doing when they made a particular *de re* attitude ascriptions to B (i.e. B claims of Australia that they won the most recent men's cricket world cup, or equivalently, B represents Australia as having won the most recent men's cricket world cup), is expressible not in the vocabulary of cricket, but in the semantic metavocabulary of representation.

References

Antoniou, A. [2025]: "Why did the dark matter hypothesis supersede modified gravity in the 1980s?", *Studies in History and Philosophy of Science*, **112**, pp. 141–52.

Balazs, C., T. Bringmann, F. Kahlhoefer, and M. White [2024]: *A Primer on Dark Matter*, arXiv: 2411.05062 [astro-ph.CO].

Bertone, G. and T. M. P. Tait [2018]: "A new era in the search for dark matter", *Nature*, **562**, pp. 51–6, ISSN: 1476-4687.

Blackburn, S. [1993]: *Essays in quasi-realism*, Oxford University Press.

Brandom, R. [2009]: *Articulating Reasons: An Introduction to Inferentialism*, Harvard University Press.

Brandom, R. [1983]: "Asserting", *Noûs*, pp. 637–50.

Brandom, R. [1994]: *Making It Explicit: Reasoning, Representing, and Discursive Commitment*, Harvard University Press, ISBN: 978-0-674-54330-0.

Brandom, R. [2010]: *Between Saying and Doing: Towards an Analytic Pragmatism*, Oxford University Press.

Brandom, R. [2014]: "Intentionality and Language: A Normative, Pragmatist, Inferentialist Approach", *The Cambridge handbook of linguistic anthropology*, pp. 347–63.

Brandom, R. [2019]: *A Spirit of Trust: A Reading of Hegel's Phenomenology*, Belknap Press.

Bueno, O. and S. French [2011]: "How Theories Represent", *The British Journal for the Philosophy of Science*, **62**, pp. 857–94, ISSN: 0007-0882, JSTOR: 41410121.

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.

Clowe, D. et al. [2006]: "A Direct Empirical Proof of the Existence of Dark Matter*", *The Astrophysical Journal*, **648**, p. L109.

De Baerdemaeker, S. and R. Dawid [2022]: "MOND and meta-empirical theory assessment", *Synthese*, **200**, p. 344.

De Swart, J., G. Bertone, and J. van Dongen [2017]: "How dark matter came to matter", *Nature Astronomy*, **1**, p. 0059.

Duerr, P. M. and W. J. Wolf [2023]: "Methodological reflections on the MOND/dark matter debate", *Studies in History and Philosophy of Science*, **101**, pp. 1–23.

Dummett, M. A. E. [1975]: "What Is a Theory of Meaning?", in S. Guttenplan (*ed.*), *Mind and Language*, Oxford University Press.

Famaey, B. and S. McGaugh [2012]: "Modified Newtonian dynamics (MOND): observational phenomenology and relativistic extensions", *Living reviews in relativity*, **15**, p. 10.

Fine, A. I. [1984]: "The Natural Ontological Attitude", in *Scientific Realism*, University of California Press.

Garrett, K. and G. Dūda [2011]: "Dark Matter: A Primer", *Advances in Astronomy*, **2011**, p. 968283, available at <https://onlinelibrary.wiley.com/doi/pdf/10.1155/2011/968283>.

Hacking, I. [1983]: *Representing and intervening*, Cambridge: Cambridge University Press.

Jacquart, M. [2021]: "Dark matter and dark energy", in *The Routledge Companion to Philosophy of Physics*, Routledge, pp. 731–43.

Kraut, R. [1990]: "Varieties of Pragmatism", *Mind*, **99**, pp. 157–83, ISSN: 0026-4423, JSTOR: 2254958.

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.

Lewis, D. [1979]: "Scorekeeping in a language game", *Journal of philosophical logic*, **8**, pp. 339–59.

Magnus, P. D. and C. Callender [2004]: "Realist Ennui and the Base Rate Fallacy", *Philosophy of Science*, **71**, pp. 320–38, ISSN: 0031-8248, 1539-767X.

Martens, N. C. M. [2021]: "Dark Matter Realism", *Foundations of Physics*, **52**, p. 16, ISSN: 1572-9516.

Martens, N. C. M. and D. Lehmkuhl [2020]: "Dark Matter = Modified Gravity? Scrutinising the Spacetime–Matter Distinction through the Modified Gravity/ Dark Matter Lens", *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, **72**, pp. 237–50, ISSN: 1355-2198.

McDowell, J. [1996]: *Mind and world*, Harvard University Press.

McGaugh, S. [2020]: "Predictions and outcomes for the dynamics of rotating galaxies", *Galaxies*, **8**, p. 35.

Merritt, D. [2020]: *A philosophical approach to MOND: Assessing the Milgromian research program in cosmology*, Cambridge University Press.

Milgrom, M. [2020]: "MOND vs. Dark Matter in Light of Historical Parallels", *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, **71**, pp. 170–95, ISSN: 1355-2198.

Price, H., S. Blackburn, R. Brandom, P. Horwich, and M. Williams [2013]: *Expressivism, Pragmatism and Representationalism*, Cambridge University Press.

Psillos, S. [2001]: *Scientific Realism: How Science Tracks Truth*, Routledge.

Psillos, S. [2012]: "Causal descriptivism and the reference of theoretical terms", *Perception, realism, and the problem of reference*, pp. 212–38.

Riess, A. [2024]: "Dark Matter", *Encyclopaedia Britannica*.

Rorty, R. [2009]: *Philosophy and the Mirror of Nature*, Princeton university press.

Sellars, W. [1953]: "Inference and Meaning", *Mind*, **62**, pp. 313–38, ISSN: 0026-4423, JSTOR: 2251271.

Sellars, W. [1963a]: "Empiricism and the Philosophy of Mind", *Science, Perception, and Reality*.

Sellars, W. [1963b]: *Science, Perception, and Reality*, Ridgeview.

Sellars, W. [1968]: *Science and Metaphysics: Variations on a Kantian Theme*, Routledge and Kegan Paul.

Sellars, W. [1974]: "Meaning as Functional Classification", *Synthese*, **27**, pp. 417–37, ISSN: 1573-0964.

Sellars, W. [1976]: "Is Scientific Realism Tenable?", *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, **1976**, pp. 307–34, ISSN: 0270-8647, JSTOR: 192388.

Stalnaker, R. C. [1997]: "Reference and Necessity", in B. Hale and C. Wright (*eds.*), *A Companion to the Philosophy of Language*, Blackwell.

Trimble, V. [1987]: "Existence and nature of dark matter in the universe", *Annual review of astronomy and astrophysics*, **25**, pp. 425–72.

Van Fraassen, B. C. [1980]: *The Scientific Image*, Oxford University Press.

Van Fraassen, B. C. [2001]: "Constructive Empiricism Now", *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, **106**, pp. 151–70.

Vaynberg, E. [2024]: "Realism and the detection of dark matter", *Synthese*, **204**, p. 82.

Wallace, D. [2022]: "Stating Structural Realism: Mathematics-First Approaches to Physics and Metaphysics", *Philosophical Perspectives*, **36**, pp. 345–78.

Wittgenstein, L. [2010]: *Philosophical Investigations*, John Wiley & Sons, ISBN: 978-1-4443-0797-9.