Spacetime functionalists should be inferentialists

Tushar Menon

Dianoia Institute of Philosophy,
Australian Catholic University, Melbourne

Abstract

This paper offers a reading of spacetime functionalism as a form of inferential-role metasemantics. It argues that, on this reading, the spacetime functionalist has the resources to deal with two significant challenges in the foundations of spacetime theories: (i) how to account for the manner in which spacetime vocabulary gains purchase on the world and (ii) how to clarify the explanatory relationship between (external) dynamical symmetries and spacetime symmetries in classical spacetime theories.

Contents

1 Introduction

In a perfect world, this would be a very short paper. I would simply point out that the titular thesis is a special case of a more general thesis that everyone accepts. But this is the actual world, and alas, most people here do not accept that everyone should be an inferentialist. So there is a little more work to do here, using a slightly different dialectical strategy. I will not attempt to argue for the general claim. Instead, I will advance a reading of spacetime functionalism, and in particular of Knox's brand of inertial frame spacetime functionalism, as a kind of inferential role metasemantics[1]—one restricted to spacetime vocabulary. I will then demonstrate that this move equips us with the resources to deal with two thorny challenges for the broad project of interpreting spacetime theories, in the hope that these benefits make the position of inferentialism about spacetime vocabulary look attractive.

The first challenge is what we might call `The Interpretational Challenge`, which is simply a demand for an account of how spatiotemporal vocabulary comes to mean what it does. Or, to borrow a more evocative rendering from Myrvold, how our spacetime concepts gain purchase on the physical world. This challenge exists, of course, for any account of spatiotemporal vocabulary. I choose, in this paper, to detail how spacetime functionalists are particularly well-equipped to meet this challenge when understood as inferentialists, but this should not be taken to suggest that inferentialism requires a commitment to spacetime functionalism. My discussion of inferentialism is tied to spacetime functionalism, largely for dialectical convenience; Knoxian spacetime functionalism is a sharply-articulated position, and this allows for a sharp articulation of a plausible form of inferentialism. Other positions in the literature regarding spacetime vocabulary (for example Baker's [2] or Janssen's [15]) might just as easily be given an inferentialist underpinning, but to demonstrate this would take us beyond the scope of this paper.

The second challenge is one that was separately articulated by Acuña [1] and Myrvold [23] as arising from a tension in the dynamical approach as advocated by Brown and Pooley [5, 7]. In fact, this challenge extends beyond the dynamical approach, to anyone who posits an explanatory arrow (in either direction) between spacetime structure and dynamical symmetries. According to the dynamical

[1] In this paper, I will use 'inferential role metasemantics' interchangeably with 'inferentialism'.

approach, spacetime facts are to be understood as being explained by, and exhausted by, symmetry facts about material fields; according to the opposing geometrical approach, it is the spacetime facts that explain dynamical symmetry facts. In the context of the dynamical approach, Acuña and Myrvold highlight the tension between the commitment to a unidirectional arrow of explanation, and Brown's concurrent commitment to the analyticity (hence bidirectionality) of the relation between Minkowski spacetime geometry and the Lorentz invariance of the dynamical laws of matter fields. Call this the `Acuña-Myrvold Challenge`.

In summary, then, the two challenges are:

`The Interpretational Challenge`: Account for how our spacetime vocabulary comes to mean what it does.

`The Acuña-Myrvold challenge`: Reconcile the uni-directionality of explanation with bi-directionality of analyticity in the context of discussions of spacetime geometry.

I should stress that I am not presenting inferential role metasemantics as the correct way of reading Knox's own view of her project. Although Knox is somewhat unusual among spacetime functionalists in rejecting a strong connection to the Lewisian project of functional reduction, there is no textual evidence to support the claim that Knox instead sees spacetime functionalism as a metasemantic claim. Having said that, of course, I will argue that her motivations are well-served by this view, but this is intended as a normative, rather than descriptive claim: I think that Knox ought to be an inferentialist, but I do not think that she already is committed to being one.

The paper is structured as follows. I begin, in §2, with a brief overview of spacetime functionalism. I introduce inferential role metasemantics, and argue that spacetime functionalism can be read as a special case, in §3, Each of the two subsequent sections is then devoted to one of the challenges: §4 discusses the `Interpretational Challenge` and §5 the `Acuña-Myrvold challenge`.

2 Spacetime functionalism

A little over a decade ago, two distinct strains of functionalist accounts of spatial and spatiotemporal vocabulary emerged. Both posited that spatial or spatiotemporal

concepts were to be understood functionally, and both specified what that function was. But that's pretty much where the similarities cease.

The first strain is tied to Chalmers' *scrutability* project, expounded in his [9]. According to this form of functionalism, spatial properties are picked out by their causal/nomic function in a constituting a basis for certain sorts of experience. This sort of functionalism is a version of what is sometimes known as `casual functionalism`. Chalmers' goal is to recover (or construct) truths about space from a base set of truths about the sorts of things that play the appropriate causal/nomic roles.

The second strain was introduced by Knox in [16]. According to this form of functionalism, spatiotemporal structure is constituted by whatever theoretical structure picks out the structure of inertial frames (I will discuss the details of the proposal below). This sort of functionalism is a version of what we might call `constitutive functionalism`. Knox's goal by contrast with Chalmers', is to pick out certain theoretical structure as spatiotemporal in order to fast-track certain inferences. As Knox put it:

> [C]onsidering the inertial structure provides a shortcut that allows us to glean the empirical consequences of a theory without going into the messy details of our various measuring devices. [17, p. 347]

Knox therefore posits, as the functional role of spacetime structure, that it serves to pick out the structure of inertial frames, where the latter are characterised as follows [17, p. 348]:

1. Inertial frames are frames with respect to which force free bodies move with constant velocities.

2. The laws of physics take the same form (a particularly simple one) in all inertial frames.

3. All bodies and physical laws pick out the same equivalence class of inertial frames.

This identifies, as spacetime, the Lorentzian metric g_{ab} in general relativity (modulo certain edge cases as discussed in [26]), the Minkowski metric η_{ab} in special relativity, and the pair of tensors h_{ab} and t_{ab} in classical spacetimes. This sort of identification

can then put to various philosophical uses. For example, (i) it can be used to resolve problems of underdetermination of spatiotemporal structures by different theoretical formulations (see e.g. [16]); (ii) it can underpin a notion of spacetime emergence in quantum gravity (see e.g. [17]).

I contend that it can do even more. But those gains come at a price (that I argue we should pay): Knoxian spacetime functionalism needs to be understood not just as a way of identifying, as spatiotemporal, certain structure(s) in a theory, but as an account of how spatiotemporal *vocabulary* comes to mean what it does. In other words, as a piece of *metasemantics*.

3 Inferential role metasemantics

In this section, following Burgess and Sherman's [8] mapping of the landscape, I introduce, categorise, and discuss a number of distinct questions that fall under the broad heading of 'philosophical discourse about linguistic expressions'. Under this heading, we can usefully distinguish between the practices of semantics and metasemantics.[2] The former is the practice of reporting and systematising the ascriptions of certain properties, later identified as *semantic* properties, to linguistic expressions. The latter is the collection of practices centred around providing explanations of, or grounds for, these ascriptions. So, for example, a semantic claim might ascribe a meaning or a truth value to some sentence, or an extension to a singular term or predicate (think of an interpretation function in a Tarskian model); a metasemantic claim might offer an account of the grounds for making that semantic property ascription.

Let us make things more precise, and begin with semantic properties of linguistic expressions. Eventually, we will want to invoke some demarcation criteria for this sort of vocabulary, but in this paper, I will take as read that 'means' and 'refers to' straightforwardly count as semantic vocabulary, if anything does. So an example of a basic semantic claim is: E means M. One aspect of the practice of semantics, then, is the practice of determining what, for example, 'E or E'' means, given that E means M, and E' means M'. The practice as a whole is constituted by similar systematisations of the other semantic properties.

[2] For the purposes of this paper, this division can be taken as equivalent to Brandom's 'formal semantics' vs. 'philosophical semantics' [4], Dummett's 'theory of meanings' vs. 'theory of meaning' [10] and Stalnaker's 'descriptive semantics' vs. 'foundational semantics' [27]

But the practice of semantics is silent over the origin of, and the justification of the ascription of, these semantic properties to linguistic expressions. That's where metasemantics comes in. Burgess and Sherman helpfully identify three distinct aspects of the practice of metasemantics (associated with a basic semantic fact about meaning; similar practices can be constructed around other semantic properties like reference and truth):

`Basic metasemantics`: What facts/states of affairs determine the basic semantic facts? (e.g. facts of the form '*E* means *M*' or '*P* is true')

`Theory of meaning`: How do we characterise/understand/analyse the '...means...' relation?

`Metaphysics of semantic values`: What kind of thing is a meaning?

In what remains of this section, I will discuss two metasemantic accounts, and spell out their consequences for these three questions.

One canonical set of views regards meaning claims as being grounded in certain representational relations between linguistic expressions and worldly entities. Call any view that grounds meaning entirely in representation a `representationalist metasemantics`. An example might be helpful here. Consider a good and accurate drawing of a sheep. It bears some salient relation of resemblance to a sheep, in virtue of which it would be appropriate to put it in a box labelled 'sheep' and not in a box labelled 'automobiles'. The resemblance relation therefore grounds what sorts of things it is appropriate to do with the sheep drawing. By analogy with this, proponents of representational semantics take the representation relation that obtains between a linguistic expression and its target object to ground the sorts of things it is appropriate to do with those expressions (for example, make inferences).

Individual views will differ with respect to how they cash out the representational relation, as well as what makes it the case that this representation relation exhausts meaning. For example, consider a referentialist according to whom the meaning of a word is simply determined by the object to which it refers:

`Basic metasemantics`: The fact *E* means *M* is grounded in the primitive fact that *E* refers to *M*.

`Theory of meaning`: '...means...' is analysed as '...refers to...', together perhaps with some constraints of uniqueness imposed by some description *D*.

`Metaphysics of semantic values`: The meaning M just is the referent of E, i.e. the object in the world picked out by the reference relation.

An alternative collection of views regards meaning claims as being grounded in the role that certain linguistic expressions play in our thought or discursive practice. Call any view that denies that representation grounds meaning, and claims instead that meaning is grounded entirely in such roles `conceptual role metasemantics`. To return to our sheep drawing analogy, for the proponent of a conceptual role metasemantics, it is the collection of the sorts of things that it is appropriate to do with sheep drawing (for example, to put them in a box labelled 'sheep' rather than 'automobiles') that determine the representational valence of a sheep drawing. In the linguistic context, this amounts to the belief that the conceptual roles associated with the use of a linguistic expression grounds meaning and representation. I hasten to point out that, as I have set them up, representationalist metasemantics and conceptual role metasemantics are mutually exclusive (but not necessarily exhaustive): this is ensured by my use of 'entirely grounded in' in both cases.

As with representationalist metasemantics, individual conceptual role metasemantic views will differ over the precise details. In this paper, my interest is in inferentialism, which we can now understand as a form of conceptual role metasemantics in which those conceptual roles are inferential: the meaning of an expression is the contribution that that expression makes to the inferential valence of a claim.

It is worth contrasting a specific version of inferentialism (here I choose Brandomian inferentialism, which I will discuss in more detail in §4.2) with the the specific version of representationalism (i.e. referentialism) that I discussed above. Brandomian inferentialism is a normative approach to metasemantics, where the norms of inference determine what we might call a *deontic scorecard*. A deontic scorecard is just a way of keeping track of how the inferential commitments and entitlements of a speaker are updated by their acceptance or rejection of claims (that may be made by themselves, or by other speakers). So, for example, if some speaker S were to utter 'the ball is red', my deontic scorecard might be updated to include, as one of S's commitments 'the ball is visible' (there will be several other commitments, of course). In this case, the meaning of 'red' is exhausted by all the claims that (together with what I take S's background commitments to be) entail, and can be inferred from, such sentences, in other words, by its contribution to a speaker's deontic score.

So the Brandomian inferentialist is committed to the following:

`Basic metasemantics`: The fact that E means M is grounded in the fact that M is the inferential role that E plays in our discursive practice.

`Theory of meaning`: 'E means M' is understood, not in terms of a relation between some linguistic expression and some worldly entity, but rather as way of signalling that the use of E by some speaker S updates S's deontic scorecard in accordance with the norms that determine the inferential role M.

`Metaphysics of semantic values`: M is a shorthand for the contribution E makes to S's deontic score.

3.1 Spacetime functionalism as inferential role metasemantics

In Knox's hands, functionalism about spacetime serves as a demarcation criterion for spacetime vocabulary. But we might go further and understand it as underpinning a basic metasemantic claim about such vocabulary. After all, although nothing in this setup so far suggests that Knox is making a basic semantic claim, nothing rules it out either. So let us explore a reading of Knox on which she can be understood as making the following basic semantic claim: 'spacetime' means 'the structure picked out by inertial frames'.

The quote above suggests the following broad proposal: 'spacetime' means what it does because of the implications its use has for (among other things) the predicted behaviours of bodies. Knox's own proposal is somewhat narrower, in that it suggests that the relevant implications for the behaviours of bodies are appropriately captured entirely by inertial structure.; compare this with Baker's broader, and more non-committal conceptual role suggestions in [2].

The narrower Knoxian claim suggests the following basic metasemantic claim: the meaning of 'spacetime' is conceptually tied to the meanings of 'force-free,' 'velocity', 'body','constant' and other expressions that collectively characterise inertial frame structure. For ease of reference, let us refer to this collection of expressions as 'inertial vocabulary'. The tight tethering of 'spacetime' to inertial vocabulary is what makes Knox's claim both more powerful as well as more susceptible to counterexamples. In effect, what Knox does is stipulate that 'spacetime' occupies a specific node in an inferential web that includes inertial vocabulary. Consequently,

if any theory contains inertial vocabulary, it automatically has the resources to incorporate spacetime vocabulary.

Understood as a basic metasemantic proposal, Knoxian functionalism simply says that (i) the meanings of spatiotemporal expressions are exhausted by their conceptual role and (ii) that conceptual role is captured by the inferential web of inertial vocabulary in all theories. I end this section by highlighting two related consequences of this claim.

The first is that, while the inferential web of inertial vocabulary is theory-agnostic, there will, in general, be further inferential links between inertial and other vocabulary that are theory-dependent. For example, in Newtonian mechanics, 'spacetime' might be inferentially linked to 'Galilean group', whereas in special relativity, 'spacetime' might be inferentially linked to 'Poincaré group'. This fact reflects an important aspect of the claim that should not be overlooked: all there is to being a bit of spacetime vocabulary is that it is appropriately inferentially linked to inertial vocabulary, but in specific contexts, spacetime vocabulary will also be inferentially linked to other sorts of vocabulary. A functionalist about chairs, for example, might be committed to the claim that all there is to being a chair is to play the functional role of providing a surface for people to sit on. But this does not commit them to denying that every chair has a whole host of other properties or functional roles. The claim is merely that none of these other roles or properties impinge on the classification of an object as a chair. The same is true of spacetime vocabulary.

The second is that the meaning of 'spacetime', on this view, is theory-dependent, since the meaning-conferring inferential links differ in different theories or, more generally, across different linguistic practices or communities. So, as Kuhn famously suggested 'space' really does mean something different to the Newtonian compared to the general relativist. Nonetheless, what we have here is a prescription for facilitating the understanding and coordination of use of spacetime vocabulary across different theoretical contexts and communities. This view also underpins Baker's 'cluster concept' analysis of spacetime: the reason that Baker can identify so many roles in different theoretical contexts is precisely that the spacetime concept, imported wholesale with the inferential links to inertial vocabulary from one theoretical context to another, also picks up theory-specific inferential links.

4 The Interpretational Challenge

Recall the first challenge articulated in the introduction:

`The Interpretational Challenge`: Account for how our spacetime vocabulary comes to mean what it does.

In this section, I demonstrate how the inferentialist can deal with this challenge.

The broad impulse that underlies inferentialism, as a metasemantic thesis, is that the use of expressions by speakers determines their meaning. But there are several further decisions to be made before this impulse can be turned into a precise, well-articulated philosophical position. As a result there is a slew of related, but importantly distinct, inferentialist positions, many of which differ over precise details that are relevant to the question of how our concepts gain purchase on the world.There are, in particular, two choice points that need to be highlighted.

The first is over which inferences count as meaning-conferring. In the restricted context of inferentialism about spacetime vocabulary, given the highly mathematised nature of physical theories, the relevant inferential relations are those of mathematical deduction, so we can restrict our interest to deductive inferences. The second is over scope: so far, I have advocated an inferentialism about spacetime vocabulary, while remaining neutral over whether, and if so how far, this metasemantics should extend to other, non-spatiotemporal vocabulary. In this section I discuss two families of inferentialist positions, distinguished by the extent of the vocabulary that falls under its remit. Following Chalmers, [9] I distinguish between:

`Anchored inferentialism`: '[T]he contents of some primitive concepts are determined noninferentially,... and that the content of all other concepts are determined at least in part by their inferential relations to these concepts.'

`Pure inferentialism`: '[T]here are no privileged concepts...: the content of every concept is determined only by its place in the web, and the web as a whole is characterized only by its abstract structure.'

In §4.1, I discuss the options available to the anchored-inferential spacetime functionalist. And in § 4.2, I discuss the prospects for a pure-inferential spacetime functionalism.

4.1 Anchored-inferential spacetime functionalism

In §2, I suggested that the inferential spacetime functionalist can be read as endorsing the following basic semantic claim: 'spacetime' means 'the structure picked out by inertial vocabulary'. Immediately, this raises the question of how inertial vocabulary comes to mean what it does. After all, if meaning flows through inferential links, and 'spacetime' is the sink, we need to identify the source.

The most conservative option is to restrict the scope of the inferential metasemantic story to 'spacetime', and its immediate cognates: 'space', 'time', 'spatiotemporal'. The inferential links, then, confer meaning to these, and only these, expressions, by linking them to other expressions, call them 'anchors', which acquire their meaning non-inferentially. The anchors function as meaning sources for spacetime vocabulary. I will discuss one concrete proposal for the metasemantics of these anchors, together with some difficulties for the proposal. While these difficulties do not rule out anchored inferentialism as a way of spelling out inferential spacetime functionalism, they do suggest that it is worthwhile to explore pure-inferential accounts. These will be discussed in §4.2.

For the anchored-inferential spacetime functionalist, the anchors are the inertial vocabulary: terms like 'velocity', 'body', 'force' and so on. Our vocabulary thus splits into two collections, the first of which is given a representational metasemantics, the second an inferentialist metasemantics.

As a concrete example of such a proposal, consider the Ramsey-Carnap-Lewis (RCL) proposal for functional reduction of theoretical terms [20]. The core of this project is the observation that we often, in scientific contexts, have two sorts of theoretical discourse about the same sorts of entities or phenomena (for example, both a folk-theoretic and a neuroscientific account of pain). Consequently, there is some overlap in vocabulary between these theories. The game, then, is to identify a functional role that both sets of discourse possess the vocabulary to describe, and then to define certain terms in one of those vocabularies ('t-terms') as synonymous with expressions constructed out of terms in the other vocabulary, ('o-terms'). The required synonymy is established by the identity of the functional role as expressed in the two different vocabularies.[3]

[3] Gomes and Butterfield [13] highlight that the RCL proposal actually involves *three* vocabularies, one from the 'folk theory' and two from the 'scientific theory'. In this section, since nothing turns on the difference between the two sets of vocabulary from the scientific theory, for dialectical convenience, I treat the pair as a single vocabulary.

Two important points to note. The first is that Lewis makes no claim about what that functional role should be. The project is, after all, intended as a general prescription for defining troublesome terms across different sorts of discourse. The second is that Lewis does not take a stance on the metasemantics of the anchors—he makes no proposal (at least in the specific context of [20]; with Lewis, there is always a broader world-view lurking in the background) about the origin of meaning of o-terms. The anchored-inferentialist spacetime functionalist makes two further claims at precisely these points, at which Lewis remains neutral: (i) they specify that the functional role is an inferential role and (ii) they stipulate that the o-terms are to be understood non-inferentially, and the t-terms inferentially.

If we grant the intelligibility of the RCL proposal understood as a form of anchored inferentialism, we now have to tell a story about how the anchors come to be meaningful in a non-inferential manner. As a first attempt, consider a strong version of the anchoring thesis:

`Strongly anchored`: The representational directedness of linguistic expressions towards worldly entities is self-standing, in the sense that it is intelligible in isolation from, and therefore prior to, considerations of the inferential significance of those expressions.

The strongly anchored position is perhaps most famously exemplified by Wittgenstein's St. Augustine (Wittgenstine? St. Augustein?): '[e]very word has a meaning. The meaning is correlated with the world. It is the object for which the word stands ' [28].And this world-word relation is to be understood as grounding any further use to which the words are put. At first glance, this looks promising. Notice that many of the key expressions in our inertial vocabulary can be given meaning either by ostension (e.g. 'body') or some reasonably straightforward operational definition (e.g. 'velocity', 'force-free').[4] The strongly anchored inferentialist can then establish the meanings of spacetime vocabulary via its inferential connections to these anchors.

The central problem with this proposal is that it is unclear whether our game of ostension/operationalising can be carried out with respect to all the required anchors

[4] Of course, describing the operationalisation procedure straightforward is a bit unfair, given the immense subtlety involved in Neumann [24], Lange [19] and Mach's [21] proposals for operationalising inertial frames (and consequently the notion of 'force free') in Newtonian mechanics. For details on this project, see [3, Ch. 1]. What I mean is that, while methodologically highly non-trivial, these efforts are metasemantically relatively straightforward.

(this is a specific case of a more general criticism of representationalism). 'Bodies' and 'velocity' seem fine, but the meaning-conferring inferential links extend well beyond those anchors. But what about 'dynamical symmetry', 'force', 'curvature' and other terms that pop up in different theoretical contexts? Perhaps the strongly anchored inferentialist has a story to tell about these terms, and I'm just suffering from a failure of imagination. But if not, then it might be worth exploring the possibility of a more weakly anchored account:

`Weakly anchored`: The representational directedness of linguistic expressions towards worldly entities is not self-standing, and is only intelligible in combination with considerations of the inferential significance of those expressions.

The idea here is that there is an ineliminable representational dimension to meaning conferral, but crucially, there is also an ineliminable inferential dimension. There is no unidirectional arrow of semantic explanation between inference and representation in general; they come as a package deal and simultaneously confer meaning to expressions. Here is an example from McDowell [22, p. 159]:

> [C]onsider this explicitly relational claim: " 'Snow' and snow are related thus: concatenating the former with, e.g., 'is white' yields a sentence usable to assert a truth just in case the latter is white." This is not "Augustinian", since it affirms its relation by exploiting the concept of asserting, a move in the [inferential] language-game; and it is not representationalist in any untoward sense, if we do not pretend it would be intelligible independently of inferential relations between contents of potential assertions. Why not group it with, e.g., "Pittsburgh and Philadelphia are related thus: the former is further west than the latter"? This does not threaten our grip on the anti-representationalist insight.

On this view, then, we might posit that some expressions function as anchors, and that these anchors are identified as the expressions whose meaning are primarily (if not exclusively) determined representationally (for example, 'bodies' and 'velocity'). But this does not rule out the possibility that the meanings of some of the other inertial expressions (for example, 'dynamical symmetry' or 'force') are themselves primarily determined inferentially, in the same way as 'snow' was given meaning in McDowell's example.

Your antecedent metasemantic commitments will determine whether or not some form of anchored inferentialism about spacetime is attractive; I hope to have demonstrated, via the sketches of accounts in this section, the sorts of commitments that are consonant with anchored inferentialism. On a reading of spacetime functionalism as a form of inferentialism, one can read spacetime functionalists who adopt a Lewisian perspective, for example, Chalmers [9], Lam and Wüthrich [18], and Yates [32], as anchored-inferential spatial/spacetime functionalists.

4.2 Pure-inferential spacetime functionalism

You might worry that anchored inferentialism does not go far enough in expunging representationalism from metasemantics. I certainly do. McDowell's reading of Wittgenstein's lesson is that language games are nontrivially implicated in metasemantics; Brandom extracts a much stronger lesson from Wittgenstein: that moves in language games are wholly constitutive of meaning. In terms of the scope mentioned in §4.1, the Brandomian view is holistic and pure: all expressions derive their meanings from inferential connections to all other expressions. In what follows, I will briefly explore the prospects for a pure-inferential spacetime functionalism, as a special case of the Brandomian inferentialism that I introduced in §1.

Brandomian inferentialism, vast and intricate as it is, can be understood, at first pass, as motivated by thinking about intentionality, i.e. by thinking about the curious fact that mental states and linguistic expressions are unique in that they can be directed towards entities, in the specific sense of being *about* them. In this paper, the focus is on linguistic expressions, so we can restrict the animating intentional question to one about language. Let us use the term 'conceptual content' to describe the content of linguistic expressions in virtue of which they display intentionality.

`The intentionality question`: What makes it the case that certain linguistic expressions are about worldly entities?

The pure representationalist argues that the conceptual content of an expression, whatever it is, is prior to any inferential role. We have already encountered two examples of pure representationalism in this paper: referentialism (§1) and Augustinian representationalism (§4.1). The details differ widely across the spectrum of representationalist accounts, but what unites them all is this basic metasemantic commitment to the order of semantic explanation: representation grounds meaning and, consequently, inference.

The Brandomian (pure) inferentialist reverses this order of explanation. Here, inference grounds meaning and, consequently, representation. Although precise mechanics of how inference grounds representation go beyond the scope of this paper (but see [4, Ch. 8] for the canonical account), I will, in this section, sketch the important moving parts.

According to Brandom, the intentional content of linguistic expressions is inferentially articulated. The immediate concern with this slogan is that it appears to conjure a word-world link (e.g. reference or representation) out of nothing more than word-word relations (i.e. inference). This would, indeed, be a problem if the project was predicated on the idea that there was some privileged family of word-world 'gluing' relations that we were in the business of trying to characterise using only intralinguistic resources.

Recall, in §3, we identified as one of the three aspects of metasemantics, the `Theory of meaning`. And in that section, I described the inferentialist as asserting that'E means M' is understood, not in terms of a relation between some linguistic expression and some worldly entity, but rather as way of signalling that E updates the deontic scorecard in accordance with the norms that determine the inferential role M. This claim follows from two significant moves: (i) adopting a form of norm-expressivism and (ii) adopting, as the appropriate norms,those of inference.

The expressivist move begins with the denial of a claim that appears to follow from the surface grammar of 'E means M'. The surface grammar suggests that E and M stand in a 2-place relation of '...means', just as 'Virat Kohli' and 'Steve Smith' stand in the 2-place relation of '...has scored more runs than...' The expressivist move is to argue that certain statements, despite bearing a superficial resemblance to such declarative sentences, are, in fact, not to be understood as ascribing some properties to (or relations between) entities in the world. Instead, such statements simply express the speaker's endorsement of the contents of such claims. One common way of spelling out how this endorsement works is in terms of norms: a norm-expressivist about rationality, for example, understands 'S says that X is rationally permissible' as S's expressing a commitment to a set of norms of rationality, according to which X is permissible (this example is from [12]).

The Brandomian inferentialist is a norm-expressivist about *all* vocabulary, where the norms are those of inferential propriety, as systematised by the deontic scorecard associated with each speaker. For 'red' to mean red is for 'red' to be caught up in the right sorts of inferences. If S says 'the cricket ball is red' then S is thereby committed

to endorsing, among other things, 'the cricket ball is coloured', 'the cricket ball has a diameter greater than 750 nm', and so on. And these commitments are kept track of by *S*, and all of *S*'s interlocutors (who will, in general, disagree over where *S* should endorse those commitments, given their own differing background commitments). So the basic idea here is that 'red' has the conceptual content that it does in virtue of the complex web of inferences in which it is caught up.

But how can this be? Remember, the inferentialist severed the representationalist's tethering of word to world, in order to attempt to reverse the order of semantic explanation. How, then, can any claims about inference be conceptual *a fortiori* intentional? How, in other words, can we re-attach our language to the world?

This is where we have to introduce the last component of the account: Pragmatism. Capital-P Pragmatism refers to a collection of views based on the insight that an agent's knowledge of the world is inseparable from their agency within the world. According to the inferentialist-pragmatist, we should understand language use as a form of inferentially-inflected activity. Activities performed by agents are always constrained by the way the world is: for example, gravity and human physiology being what they are, I cannot reach terminal velocity and leave the earth simply by flapping my arms. Similarly, the world being the way it is constrains the sorts of inferences that agents embedded in the world can make.

An example might help here. Suppose I'm at a traffic signal, and my perceptual evidence leads me to conclude that the red light, and no other, is glowing. So I stop. This is because I'm embedded in a social context in which it is necessary to stop at a red light. But I can only make that judgement if I understand what 'red' means. This is where the inferential significance comes in. Here is one important inference that follows from 'the traffic light is currently glowing red': 'the traffic light is not currently glowing green'. Evidence that I endorse this inference is that I stop at the signal. The use of 'not', in this case, signals a commitment to a particular incompatibility. And the use of 'red' and 'green' helps articulate what that incompatibility is. And that is the expressive role of linguistic expressions, on this picture: to make explicit, and articulable, what is implicit in my behaviour. The world constrains my behaviour, and my behaviour has aspects that are made explicit by particular linguistic performances. This is how the word-world link is re-established.

Inferentially articulated content, when tethered to the world in this way, becomes conceptual, *a fortiori*, intentional content. And this is how the Brandomian

inferentialist answers the question of intentionality. The entire web of concepts at a speaker's disposal is constantly being updated by (i) the linguistic performances of others and (ii) cues from the world. The pure-inferential spacetime functionalist, therefore, imbues conceptual content to spatiotemporal vocabulary by highlighting the inferential links to inertial vocabulary. But unlike the anchored-inferential spacetime functionalist, they understand inertial vocabulary as conceptually contentful via global inferential links to other nodes, rather than via piggybacking on a representationalist story about anchors.

In summary, then, in this section I have sketched a few accounts of how to get inferential spacetime functionalism off the ground; more will need to be done to turn these into full-fledged semantic accounts. The main purpose of these sketches was to demonstrate, ultimately, how one might account for the ways in which spacetime concepts gain purchase on the world, in terms of the sorts of further semantic and metasemantic commitments one might have to incur. What these accounts have in common is that they provide a way of underpinning Myrvold's (in my opinion correct) suggestion that (i) careful attention needs to be paid to the conceptual role being played by spacetime vocabulary and (ii) we should be careful not to allow metaphorical uses of linguistic expression to cloud our judgements about their actual semantic significance.

5 The Acuña-Myrvold Challenge

Earman [11] introduced two symmetry principles, SP1 and SP2, as a way of maintaining good philosophical hygiene when interpreting spacetime theories. The principles, at least in the context of classical particle physics, are compelling: SP1 mandates that (our interpretations be such that) all dynamical symmetries are spacetime symmetries, and SP2 the converse. In order for these principles to function as nontrivial directives, there has to be some sense in which one can fail to adopt them. What this means is that it should be possible for spacetime symmetries to outstrip dynamical symmetries and vice versa. And on Earman's picture, this seems to be the case: spacetime symmetries on models of the form $\langle M, A, P\rangle$ are diffeomorphisms of the manifold M, whose drags-along applied to absolute tensorial objects, A—which are taken to represent spacetime—are automorphisms of the model, while dynamical symmetries are diffeomorphisms whose drags-along applied to dynamical tensorial objects, P preserve solutionhood of models. On

this setup, it seems that SP1 and SP2 can, but according to Earman should not, be violated.

Myrvold identifies an important presupposition to the intelligibility of the failure of SP1 and SP2: that two notions of 'dynamical symmetry' and 'spacetime symmetry' can be conceptually prised apart. Clearly, we can define using tensorial models of a theory, two formally distinct operations as we did above. But we should not simply christen those operations 'dynamical' and 'spacetime' symmetry transformations, without first ensuring that that choice of vocabulary reflects the meanings of 'spacetime' and 'dynamical' in the broader theoretical context. Myrvold denies that this is possible:

> Let us consider SP1. If it is not analytically true, then it makes sense (even if we take it to be false) to talk about a spacetime asymmetry with no corresponding dynamical asymmetry. Any attempt to do so, I claim, fails. To speak, for example, of a state of rest with no dynamical significance is to abandon any sense the word has in physical discourse and not replace that sense with anything else.[23]

In the rest of the paper from which this quote is taken, Myrvold argues that any attempt to prise apart 'dynamical symmetry' from 'spacetime symmetry' ultimately collapse into intuition pumps that rely on feeble metaphors. In this section, I want to focus on a consequence of Myrvold's analyticity claim: that it undermines a central claim made by proponents of the dynamical approach, that spacetime symmetries are explained by, or grounded in, dynamical symmetries. Here is a representative quote:

> It is the Lorentz covariance of the laws that underwrites the fact that the geometry of spacetime is Minkowskian.[5]

This sets up the tension that I identified in the introduction, which leads to the following challenge:[5]

`The Acuña-Myrvold challenge`: Reconcile the uni-directionality of explanation with bi-directionality of analyticity in the context of discussions of spacetime geometry.

[5] As I mentioned in the introduction, this would also be a problem for a particular sort of proponent of the geometrical approach, namely one for whom (i) dynamical symmetry structure is grounded in the structure of spacetime, and (ii) SP1 and SP2 are analytically true.

When confronted with a tension like this, if we do not want to simply deny that explanation is uni-directional, or that SP1 and SP2 are analytic, then it seems as if there are, broadly speaking, two choices. The first is to simply deny the explanatory claim *tout court*. This is Myrvold's strategy:

> Every moving rod indeed contracts because of the nature of the forces that hold it rigid, but it is not correct to say that it is not because of the spacetime environment that rods contract, if attributes of that spacetime environment are codifications of symmetries shared by all forces that could be responsible for the behaviour of moving rods... on the dynamical perspective, there is no question of explanatory priority between spacetime structure and dynamics.[6]

The second choice is to deny that the explanans and explananda are, respectively, dynamical and spacetime symmetries. This is Acuña's strategy:

> The horizontal bidirectional arrow connecting the Lorentz transformations and Minkowski spacetime structure is not explanatory. However,...Minkowski spacetime structure is not explanatorily idle. Although it does not explain Lorentz invariance (nor is it explained by it), that Minkowski spacetime is the unfolding of the chronogeometric structure of the theory certainly constitutes a source of further intelligibility of physical phenomena, not of the Lorentz invariance of physical laws.[1]

Here, the tension is resolved by changing the subject: the explanadum is 'physical phenomena', while the explanans, is *both* the Lorentz symmetries of the laws and Minkowski spacetime, since they are analytically linked (i.e. claims about Minkowski spacetime are synonymous with claims about Lorentz invariance of dynamical laws).

While these are both perfectly sensible options, they each come at a cost that, *ceteris paribus*, we should try to avoid. With this in mind, I would like to propose a *tertium quid*: that the uni-directional arrow of explanation is metasemantic, while the bi-directional arrow of analyticity is, as always, semantic. This proposal has the advantage of resolving the tension without either denying a useful explanatory claim, as Myrvold does, or changing the explanadum, as Acuña does.

I take both Acuña and Myrvold to endorse the following basic semantic claim: 'Spacetime symmetry' means [the same as] 'dynamical symmetry'[means]. Neither

of them, however, provides an account of why this is the case. In other words, neither of them commits to an underlying basic metasemantic claim. This is where inferential spacetime functionalism comes in.

The inferential spacetime functionalist's basic metasemantic proposal is that the core web of inferential, therefore meaning-conferring, links are between 'spacetime' and inertial vocabulary. In particular, 'spacetime' is conceptually tied to the meanings of expressions like 'force-free', 'dynamical symmetry' and 'material', in any theory in which all of those linguistic expressions appear. But note that the converse is not necessarily true: 'dynamical symmetry' and 'material particle/field', for example, are both meaningful even in a theoretical context with no force-free motion, or other inertial vocabulary. In other words, according to the inferential spacetime functionalist, dynamical expressions are not, in general, conceptually tied to spacetime vocabulary in the same way that spacetime vocabulary is conceptually tied to dynamical expressions.

Consequently, the basic metasemantic claim is that 'spacetime' means what it does because its conceptual role is tied specific inertial vocabulary, and not the other way around. The metasemantic claim is unidirectional, and in virtue of it being a basic metasemantic claim, is by construction, explanatory of the basic semantic claim. And none of this impinges on the analyticity of the relationship between 'spacetime symmetry' and 'dynamical symmetry'.

This move is available only to the inferentialist. The representationalist simply does not have the resources to separate the semantic arrow from the metasemantic one. After all, if representation grounds meaning, then there is no question of the conceptual roles being understood in any way other than what is fixed by the representation relation. I read Myrvold as implicitly adopting a representationalist metasemantics, which is why he is forced to drop the explanatory claim. But if we deny the representationalist assumption, then the Myrvoldian argument becomes even more compelling. Rather than stipulating that 'spacetime symmetry' and 'dynamical symmetry' mean the same thing, and then challenging their opponent to provide a compelling, non-metaphorical account of how this might not be the case, as the representationalist Myrvoldian does, the inferentialist Myrvoldian can provide an account of what grounds the analyticity claim.

So the inferentialist clearly has the edge over the representationalist here. But that does not mean the inferential spacetime functionalist is in the clear. For one might worry about a natural inferentialist rejoinder: can we not instead be material

functionalists? That is to say, can we not simply choose to reverse the metasemantic arrow of explanation to make spacetime vocabulary indispensable for conferring meaning to material/dynamical talk? Of course we can. And indeed, one might be able to construct an argument in favour of the geometrical approach on this basis. But a quick look at the practice of contemporary physics demonstrates why this proposal is less appealing than the spacetime functionalist's.

Virtually every proposal from quantum-gravitational physics employs talk of material degrees of freedom, and associated dynamical symmetries, independently of any associated spacetime structure. String theorists, for example, speak of dynamical symmetries of the worldsheet independently of any claim about that worldsheet being embedded in spacetime. Indeed, the most compelling arguments from both philosophers ([14]) and physicists ([25]) in this context suggest that we should understand spacetime as emerging from the dynamics of the worldsheet. Similar analyses of the dynamics of other quantum gravity theories (loop quantum gravity [30], causal set theory [29, 31]) suggest that the fundamental dynamical degrees of freedom explicitly resist a spatiotemporal interpretation. So if we want to take seriously the practice of contemporary physics, then we have to accept that talk of material degrees of freedom is conceptually coherent independently of spacetime talk.

6 Conclusion

Spacetime functionalists should understand their position as a form of inferentialism. If they do, then, as I have argued, they equip themselves with the resources to resolve some deep conceptual challenges in the foundations of classical spacetime theories. In this paper, I explicitly demonstrated how inferential spacetime functionalists deal with what I called (i) the interpretational challenge and (ii) the Acuña-Myrvold challenge.

The basic thought behind the proposal was the following: in addition to reading Knox as providing a demarcation criterion for what counts as spatiotemporal structure, we should read Knox as providing us with the following basic metasemantic claim: (i) the meanings of spatiotemporal expressions are exhausted by their conceptual role and (ii) that conceptual role is captured by the inferential web of inertial vocabulary in all theories.

The details of how this inferential web is characterised will depend on antecedent

semantic and metasemantic commitments. I canvassed two options. Firstly, anchored inferentialism, according to which some nodes in the inferential web are endowed with meaning non-inferentially. Secondly, and more compellingly, pure inferentialism, according to which all nodes are endowed with meaning purely by their positions with respect to all the other nodes.

The interpretational challenge was met by the pure-inferential spacetime functionalist by invoking the Pragmatist ideal, according to which knowledge of the world is inseparable from agency within it. Invoking a Brandomian argument about intentionality, I described how purely-inferentially articulated content could be tethered to the world via the functionalist's focus on the inferential links between spacetime vocabulary and inertial vocabulary, together with the manner in which the world determines the appropriateness of one's commitments to those inferences.

The Acuña-Myrvold challenge was met by the inferential spacetime functionalist by noting that inferentialism supplies us with a way of de-tethering semantic arrows of analyticity from metasemantic arrows of explanation. I argued that Myrvold's compelling argument about the analyticity of 'spacetime symmetry' and 'dynamical symmetry' could be given a more secure foundation if one gives up on representationalism and embraces inferentialism.

References

[1] Pablo Acuña. Minkowski spacetime and Lorentz invariance: The cart and the horse or two sides of a single coin? *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 55:1–12, August 2016.

[2] David John Baker. Knox's inertial spacetime functionalism (and a better alternative). *Synthese*, March 2020.

[3] Julian B Barbour. *The Discovery of Dynamics: A Study from a Machian Point of View of the Discovery and the Structure of Dynamical Theories*. Oxford University Press, 2001.

[4] Robert Brandom. *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Harvard University Press, 1994.

[5] Harvey Brown. *Physical Relativity*. Oxford University Press, Oxford, 2005.

[6] Harvey Brown and Oliver Pooley. Minkowski space-time: A glorious non-entity. *Philosophy and Foundations of Physics*, 1:67–89, 2006.

[7] Harvey R. Brown and Oliver Pooley. The Origins of the Spacetime Metric: Bell's Lorentzian Pedagogy and its Significance in General Relativity. In Craig Callender and Nick Huggett, editors, *Physics Meets Philosophy at the Plank Scale*, pages 256–72. Cambridge University Press, 1999.

[8] Alexis Burgess and Brett Sherman. *Metasemantics: New Essays on the Foundations of Meaning*. Oxford University Press, 2014.

[9] David J Chalmers. *Constructing the World*. Oxford Univeristy Press, Oxford, 2012.

[10] Michael A. E. Dummett. What is a theory of meaning? In Samuel Guttenplan, editor, *Mind and Language*. Oxford University Press, 1975.

[11] John Earman. *World Enough and Spacetime*. MIT Press, Cambridge, MA, 1989.

[12] Allan Gibbard. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Harvard University Press, 1990.

[13] Henrique Gomes and Jeremy Butterfield. Geometrodynamics as functionalism about time. In Claus Kiefer, editor, *From Quantum to Classical: Essays in Honour of H.-Dieter Zeh*, pages 135–178. Springer International Publishing, Cham, 2022.

[14] Nick Huggett and Christian Wüthrich. Out of Nowhere: The 'emergence' of spacetime in string theory, 2020.

[15] Michel Janssen. Drawing the line between kinematics and dynamics in special relativity. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 40(1):26–52, January 2009.

[16] Eleanor Knox. Newton–Cartan theory and teleparallel gravity: The force of a formulation. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 42(4):264–275, November 2011.

[17] Eleanor Knox. Effective spacetime geometry. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 44(3):346–356, August 2013.

[18] Vincent Lam and Christian Wüthrich. Spacetime is as spacetime does. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 64:39–51, November 2018.

[19] Ludwig Lange. *Die Geschichtliche Entwickelung Des Bewegungsbegriffes Und Ihr Voraussichtliches Endergebnis.* Englemann, Leipzig, 1886.

[20] David Lewis. How to Define Theoretical Terms. *The Journal of Philosophy*, 67(13):427–446, 1970.

[21] Ernst Mach. *The Science of Mechanics: A Critical and Historical Account of Its Development (2nd English (4th German) Ed.).* Open Court Publishing Company, Chicago, 1919.

[22] John McDowell. Brandom on Representation and Inference. *Philosophy and Phenomenological Research*, 57(1):157–162, 1997.

[23] Wayne C. Myrvold. How could relativity be anything other than physical? *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 67:137–143, August 2019.

[24] C Neumann. *Ueber Die Principien Der Galilei-Newton'schen Theorie.* Teubner, Leipzig, 1870.

[25] Joseph Gerard Polchinski. *String Theory, Volume I: An Introduction to the Bosonic String*. Cambridge university press Cambridge, 1998.

[26] James Read and Tushar Menon. The limitations of inertial frame spacetime functionalism. *Synthese*, 199(2):229–251, December 2021.

[27] Robert C. Stalnaker. Reference and necessity. In Bob Hale and Crispin Wright, editors, *A Companion to the Philosophy of Language*. Blackwell, 1997.

[28] Ludwig Wittgenstein. *Philosophical Investigations*. John Wiley & Sons, March 2010.

[29] Christian Wüthrich. The structure of causal sets. *Journal for general philosophy of science*, 43:223–241, 2012.

[30] Christian Wüthrich. Raiders of the lost spacetime. *Towards a theory of spacetime theories*, pages 297–335, 2017.

[31] Christian Wüthrich and Nick Huggett. Out of Nowhere: The emergence of spacetime from causal sets. *arXiv preprint arXiv:2009.02951*, 2020.

[32] David Yates. Thinking about spacetime. In *Philosophy beyond Spacetime, Eds Nick Huggett, Baptiste Le Bihan and Christian Wuthrich*. Oxford Univeristy Press, 2021.