

Spacetime functionalists should be inferentialists

Tushar Menon

ABSTRACT

This paper argues spacetime functionalism is best allied with an inferential-role metasemantics. It argues that, on making this move, 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.

1 *Introduction*

2 *Spacetime functionalism*

3 *Inferential role metasemantics*

3.1 *Spacetime functionalism and inferential role metasemantics*

4 *The Interpretational Challenge*

4.1 *Anchored-inferential spacetime functionalism*

4.2 *Pure-inferential spacetime functionalism*

4.3 *Inferentialism vs. representationalism*

5 *The Acuña-Myrvold Challenge*

6 *Conclusion*

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 different dialectical strategy. I will not attempt to argue for the general claim.

I will instead focus on a popular view in the contemporary literature on the philosophy of spacetime: spacetime functionalism. A quick terminological clarification: I will use ‘spacetime functionalism’ to name the general view that the spacetime concept should be understood functionally, and ‘Knoxian spacetime functionalism’ (Knox [2019]) to name the special case where that functional role is understood in terms of inertial structure. I choose Knoxian spacetime functionalism because it works especially well as a test-case

for the plausibility of inferentialism as a tool for understanding concepts in physics: it is a sharply articulated, naturalistic view that is applicable across a range theories of physics. The arguments I provide apply *mutatis mutandis* for other sorts of spacetime functionalism (for example, Baker's ([2020]) or Lam and Wüthrich's ([2018]; [2021])).

Becoming a spacetime functionalist means, among other things, taking a stand on issues that raise some metasemantic questions. I argue that these questions are best answered by understanding spacetime functionalism as allied with an inferential role metasemantics¹ for spacetime vocabulary. I intend that the normative injunction in the title be given a weak reading (i.e. when a spacetime functionalist engages, as they should, in metasemantics, then the inferentialist option is considerably better than its alternatives) rather than a strong reading (i.e. spacetime functionalism is a form of metasemantics, and rules out, or is inconsistent with, any non-inferentialist metasemantics.) To defend the titular claim, I demonstrate that the inferentialist reading equips us with the resources to deal with two thorny challenges for the broad project of interpreting spacetime theories, and that alternatives to inferentialism lack comparable resources.

Call the first challenge *The Interpretational Challenge*: a demand for an account of how our spacetime concept comes to be about spacetime. This challenge exists, of course, for any account of spatiotemporal vocabulary (indeed, any vocabulary).

The second challenge, call it the *Acuña-Myrvold Challenge*, was separately articulated by Acuña ([2016]) and Myrvold ([2019]), as arising from a tension in the dynamical approach as advocated by Brown and Pooley ([1999]). According to the dynamical 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. The Acuña-Myrvold challenge extends beyond the dynamical approach, to any view which posits an explanatory arrow (in either direction) between spacetime structure and dynamical symmetries. 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.²

In summary, then, the two challenges are:

The Interpretational Challenge: Account for how our spacetime concepts come to be about spacetime.

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

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

² I should point out that Acuña's and Myrvold's views are only co-extensive to first approximation. In particular, Acuña's version of the analyticity claim is indexed to special relativity, in the spirit of a Carnapian 'analytic-in-L'-style claim. For Myrvold, on the other hand, the analyticity is maintained across the board. For the purposes of this paper, this distinction is not particularly important, since, at worst, it only limits the scope of the analyticity claim. As we will see in §5, what is of interest in this paper is the separation of the semantic and metasemantic arrows of explanation, in situations where an analyticity claim is appropriate. As long as the set of such situations is non-empty (which it is), the argument I present applies. I am grateful to an anonymous reviewer for pushing me on this point, and to both Pablo Acuña and Wayne Myrvold for discussions related to it.

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 should be given an inferentialist reading 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, the first from Chalmers ([2012]), the second from Knox ([2013]). 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.

Chalmers' ([2012]) functionalism is a component of what is sometimes called the Canberra Plan. The canonical elaboration of the Plan is (Jackson [1998]). The goal of that project is to account for how facts from some putatively problematic domains of inquiry—for example, mathematical, moral or modal—can be understood in terms of facts about some unproblematic domain, for example, the physical (let us leave aside the thorny business of demarcating the physical from the non-physical). The Canberra Plan is therefore centred around either eliminating, or deriving from unproblematic physical facts, seemingly problematic, seemingly non-physical ones. It is in this context that Chalmers presents his spatial 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 causal 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.

From the perspective of this paper, what is significant about Chalmers' view is that it draws attention a very important distinction, first introduced by Price ([2004]), between two sorts of naturalism. According to Price, the so-called 'location problems' that Jackson identifies as arising for the Canberra Plan (see (Jackson [1998], Ch. 1)) are the result of an important methodological, and ultimately metasemantic, presupposition about naturalistic philosophical practice: that the important philosophical action is taking place at the level of objects in the world. Consequently, Price dubs this view 'object-naturalism.'

For the object-naturalist, claims about the world are vindicated by truthmakers: elements of the world whose existence grounds (or explains or necessitates) the truth of claims. Claims can only be vindicated by demonstrating the existence of truthmakers to which (components of) these claims refer, either directly or indirectly. Object naturalism is contrasted with 'subject-naturalism', according to which truthmakers are not required for the vindication of claims. Instead we can (and should) vindicate our locutions by understanding the appropriateness of their use. With that done, there is no further question to be answered about whether there are entities in the world which ground the truth of these claims. What makes a view naturalist, for the subject naturalist, is not that it trades only in naturalist-approved objects, but that the standards which determine whether a given pattern of use is appropriate are naturalist-approved standards. Subject-naturalism is therefore a form of pragmatism; more on this in §4.2.

Knox's functionalism has nothing to do with the Canberra Plan, and can be given

either an object- or a subject-naturalist reading; in this paper, I argue that in either case, it is best allied with a version of inferentialist metasemantics. Knox motivates her view by describing it as ‘involv[ing]... unfashionable conceptual consideration[s]’ (Knox [2019], p.4). This is important, because it demonstrates that Knox’s view aims at something more than a mere demarcation criterion. Knox wants to illuminate the nature of the spacetime concept, not just provide a cut between the spatiotemporal and the non-spatiotemporal.

Knox’s view is best understood in the context of what we might call ‘constitutive functionalism’ (Knox and Wallace [2023], p. 4), and is the extension to the domain of spacetime physics, of Dennett’s intentional stance regarding belief-attributions (Dennett [1989], [1991]). To be a constitutive functionalist about X is to think that what it is to be an X is to play some theoretically-specified function. The specific function that Knox takes to be characteristic of spatiotemporal structure is that it picks out the structure of inertial frames (Knox [2013], 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.

Knox is correct in stating that her view is not a gambit in the substantivalism–relationalism debate, (see e.g. (Knox [2019], p. 4)). The question of interest to the spacetime functionalist is both prior and orthogonal to the substantivalism–relationalism debate, which is a debate over the ontological status of structure that has already been identified as spatiotemporal. The functionalist is interested in how to make that spatiotemporal identification in the first place. Having adopted Knox’s view, the spacetime functionalist has a decision to make regarding whether to be an object- or subject-naturalist. Their choice will depend quite heavily on what motivated their functionalism in the first place.

Understood either object- or subject-naturalistically, Knox’s functionalism for spacetime structure can be 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. (Knox [2011])); (ii) it can underpin a notion of spacetime emergence in quantum gravity (see e.g. (Knox [2013])). Of course, what precisely these resolutions amount to will be different for the two sorts of naturalist.

Let us begin with the object-naturalist. For them, there are at least some truthmakers in the natural world that can indirectly ground other truths; for example, brain states can indirectly ground facts about whether or not a subject is in pain. An object-naturalist can believe that there is a large class of claims that we have good reasons to accept as true, that are not immediately understood as claims about entities directly referred to by expressions in their base set of truths (for example, claims about pain, or the claim that local spatial distances obey Pythagoras’ theorem). It might, for example, strike some object-naturalists as implausible that ‘spacetime’ directly refers to some truthmaker in the world, whose status as spacetime is just a primitive fact. But that does not rule out the possibility that facts about spatiotemporality of structures can be grounded in (or otherwise appropriately derivable from) facts about acceptable truthmakers. Thus, an object-naturalist

might find spacetime functionalism attractive for the following reason: the truthmakers of spatiotemporal claims are material bodies, and their behaviours, specifically their motion, to which we can unproblematically refer. In §4.1, I argue that the object-naturalist spacetime functionalist is best served by a species of inferentialism known as ‘anchored inferentialism.’

The subject-naturalist disputes the object-naturalist’s claim that vindication of claims proceeds via truthmakers. To help further characterise subject-naturalism, I borrow an example from (Quine [1960]). Consider a locution like ‘I am doing this for your sake’. The object-naturalist can only vindicate this claim by either accepting the existence of actual sakes (whatever those might be), or, more reasonably, by first paraphrasing it in terms of objects whose existence they are happy to concede exist, perhaps interests or mental states. The subject-naturalist, on the other hand, can vindicate it by simply pointing to what the locution allows us to do—in this case, to make my act intelligible to others. In §4.2, I argue that the subject-naturalist spacetime functionalist is best served by a species of inferentialism known as ‘pure inferentialism.’ Although Knox’s spacetime functionalism is orthogonal to the subject- vs. object-naturalism question, it is inseparable from it: spacetime functionalists do need to make a decision about what sort of naturalism their view is a species of.

3 Inferential role metasemantics

In this section, following Burgess and Sherman’s ([2014]) 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.³ 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.

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

³ For the purposes of this paper, this division can be taken as equivalent to Brandom’s ([1994]) ‘formal semantics’ vs. ‘philosophical semantics’, Dummett’s ([1975]) ‘theory of meanings’ vs. ‘theory of meaning’ and Stalnaker’s ([1997]) ‘descriptive semantics’ vs. ‘foundational semantics’

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. 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). To link this view to terminology introduced in §2, object-naturalism entails representationalism (see e.g. (Price et al. [2013], Ch. 2)).

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 plays any role in meaning, and claims instead that meanings are grounded in such roles ‘pure 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 ground meaning and representation.

I hasten to point out that, as I have set them up, pure representationalist metasemantics and pure conceptual role metasemantics are mutually exclusive. But the pure versions of these positions represent two ends of a spectrum that also contains a plausible middle-ground, in which some vocabulary is given a representationalist metasemantics, and other vocabulary a conceptual-role metasemantics (the ‘anchored inferentialism’ discussed below is an example of a view from this part of the spectrum).

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. One further terminological decision to highlight: as is standard in the literature, I will use ‘inferentialism’ to refer to any conceptual-role metasemantic view on which that role is inferential. So even the ‘mixed’ views, in the middle of the spectrum, are forms of inferentialism, even though they also include some representationalist commitments.

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 Brandom calls 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 and inferential role metasemantics

In Knox’s hands, functionalism about spacetime serves (among other things) as a demarcation criterion for spacetime vocabulary. But we might go further and understand it as underpinning a basic metasemantic claim about such vocabulary, call it ‘inertial structure metasemantics’: ‘spacetime’ means ‘the structure picked out by inertial frames.’ The

central thesis of this paper can then be restated as: Knoxian spacetime functionalism is most attractive when understood as allied with inertial structure metasemantics (and *mutatis mutandis* for other spacetime functionalisms; compare Knox's view that all the relevant implications for the behaviours of bodies are appropriately captured entirely by inertial structure with Baker's ([2020]) broader, and more non-committal conceptual role suggestions).⁴

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 (see e.g. (Read and Menon [2021])). 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 (i) that 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 [1962], Ch. 11) 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 of the continuity of meaning of spacetime vocabulary across different theoretical contexts. This view also underpins Baker's ([2020]) '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

⁴ Attention to these different forms of functionalism serves to highlight the affinity between the inferentialist functionalism that I propose in this paper, and Sellars' ([1974]) view of meaning as functional classification.

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.

The broad impulse that underlies inferentialism, as a metasemantic thesis, is that the inferential valence of expressions, as encoded by speakers' use of those expressions, 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 [2012]), 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 [non-inferentially meaningful] 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 anchored-inferential spacetime functionalism, and argue that it is the natural metasemantic position for an object-naturalist spacetime functionalist. Then, in § 4.2, I discuss pure-inferential spacetime functionalism, and argue that it is the natural metasemantic position for a subject-naturalist spacetime functionalist. Finally, in §4.3, I compare inferentialism with representationalism in the context of the interpretational question, and argue that the former is the more appropriate metasemantics for spacetime functionalism.

4.1 Anchored-inferential spacetime functionalism

In §2, I suggested that the inferential spacetime functionalist might usefully 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. To use Kraut’s ([1990]) Rortyan terminology, this bifurcates the vocabulary into two classes, and 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 (Lewis [1970]). The core of this project is the observation that often, in scientific contexts, we 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.⁵

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 (Lewis [1970])); 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,

⁵ (Gomes and Butterfield [2022]) 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 within the scientific theory, for dialectical convenience, I treat the pair as a single vocabulary.

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* (Wittgenstein? *St. Augustein?*): '[e]very word has a meaning. The meaning is correlated with the world. It is the object for which the word stands' (Wittgenstein [2009]). 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').⁶ The strongly anchored inferentialist can then establish the meanings of spacetime vocabulary via its inferential connections to these anchors.

Recall the motivation for the object-naturalist spacetime functionalism discussed in §2: that the physicalist game of locating the truthmakers of spatiotemporal claims committed them to an unattractive form naive representationalism. Their concern was not with the implausibility of any representationalist vindication of claims, just with the naive view that spatiotemporal claims were made true by some ontologically-independent worldly entity that had the primitive, unexplained property of being spacetime. For object-naturalists of this stripe, an anchored inferentialism would be a perfect fit. The anchors would just be the referents of various terms from the base set of truths, such as 'bodies' and 'velocity'.

This proposal is unacceptable to the subject-naturalist. For them, the central problem with this proposal is that it is unclear whether either our game of ostension/operationalising can be carried out with respect to all the required anchors (this is a specific case of a more general criticism of representationalism) or, if not, whether there is an alternative account of how to pick out anchors. 'Bodies' and 'velocity' might be picked out by something like ostension, but the meaning-conferring inferential links extend well beyond those anchors. 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 (although, in §4.3 I will suggest some more forceful reasons for the functionalist to reject anchored inferentialism). 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 [1997], p. 159):

⁶ Of course, describing the operationalisation procedure straightforward is a bit unfair, given the immense subtlety involved in Neumann (Neumann [1870]), Lange (Lange [1886]) and Mach's (Mach [1919]) proposals for operationalising inertial frames (and consequently the notion of 'force free') in Newtonian mechanics. For details on this project, see (Barbour [2001], Ch. 1). What I mean is that, while methodologically highly non-trivial, these efforts are metasemantically relatively straightforward.

[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. In particular, I highlighted the appropriateness of anchored inferentialism to the object-naturalist spacetime functionalist: the anchors just are the terms whose referents constitute the truthmakers of the base set of truths for the object-naturalist, and the strength of the object-naturalist’s conviction that these truthmakers can be unproblematically picked out by our words will determine whether they opt for a strongly- or weakly-anchored inferentialism.

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. In particular, I believe that the anchored-inferentialist has to deal with the following dilemma, call it the ‘naturalist dilemma’:

Naturalist dilemma: Either the functional role can be specified, at least in part, representationally (via anchors), in which case the functionalism cannot do justice to contemporary physics or the functional role cannot be representationally specified.

Of course, the anchored inferentialist can respond to this by simply highlighting my failure of imagination, and explicitly constructing an object-naturalistically-acceptable account of how certain expressions, which unproblematically refer to worldly entities—entities such as bodies, rulers, clocks, and so on—are sufficiently inferentially richly linked that they can function as anchors, not only for ‘spacetime’ and its cognates, but for all inertial vocabulary. But for now, let us explore the second horn of the dilemma: accepting the challenge of providing a non-representationalist account of the functional role.

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. The Brandomian, therefore, embraces the second horn of the naturalist dilemma, and adopts a view on which there is no representationalist content.⁷ 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 §3.

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 linguistic expressions are somewhat unique in that they can be directed towards entities, in the specific sense of being about them.⁸ 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 (Brandom [1994], 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.

⁷ Note that, in denying representationalist metasemantics, I do not commit myself to denying the existence of a relation of representation. What I deny is that this relation is what grounds meaning. In other words, a non-representationalist metasemantics does not commit its adherents to an eliminativism about representation. Inferentialism provides a reductive and non-representationalist account of representation in terms of primitive material inferences; for an explicit example of such a construction, see (Brandom [1994], Ch. 5, 8).

⁸ A terminological clarification: my use of the term 'intentionality' is intended to invoke the semantic notion of a directedness between linguistic expressions and the world. This should not be confused with the distinct but possibly related use of 'intentionality' in the phenomenological tradition of e.g. Husserl, to indicate a directedness between consciousness and conceptual content. For more on this see e.g. (Brandom [2014], §1). I am grateful to an anonymous referee for highlighting this.

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 surface grammar of '*E* means *M*' 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 international runs than...*'. The expressivist claims that, suggestive as this interpretation is, it is just a trick of the surface grammar. They 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 (Gibbard [1990])).

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?

The answer is that we do not have to re-attach anything: there was never a detachment in the first place. To see this, we note that the subject-naturalism that motivated pure-inferentialism is a form of 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, and exhausted by, what follows from their agency within the world. An important implication of this claim for the philosophy of language is the idea that semantics (or metasemantics) is answerable only to pragmatics: that the semantic content of a claim is exhausted by the significance of the use of that claim. In other words, there are no facts about semantics that are not facts about pragmatics.

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 em-

bedded 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. It is important to highlight that in order for this view to work, we need to include, under the umbrella of 'inferences', what (Sellars [1954]) refers to as language-entry and language-exit transitions : the rules for (i) how a worldly event or stimulus gets incorporated into a discursive performance (I see a snake and say 'Yikes! A snake!') and (ii) how a discursive performance leads to intentional actions (you hear my claim, and move away) are also treated as inferences.

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.

Inferentialism is a proper subset of pragmatist metasemantics. But there are plenty of other ways of being a pragmatist. For example, Davidson's ([2001]) truth-conditional semantics can be understood as pragmatist (as Davidson himself does). So you might wonder whether we really need to go as far as inferentialism to solve the problem. Perhaps the work is being done more generally by the pragmatist aspect of the metasemantics, rather than specifically the inferential aspect. My response to this is to highlight that inferentialism is the most theoretically attractive candidate for a pragmatist metasemantics for spacetime functionalists in particular, because of how inferences are baked into spacetime functionalism. We do not have to add any further machinery; we simply make machinery we already have do a little more. Any alternative pragmatist proposal will always be less theoretically parsimonious, for the spacetime functionalist, than inferentialism, since it will have to introduce some new theoretical machinery to account for how spacetime expressions come to be about the world (for example, the Davidsonian needs a primitive, substantive conception of truth to get their truth-conditional semantics off the ground). The inferential spacetime functionalist already has the requisite tools just lying around. This, of course doesn't rule out the efficacy of some alternative pragmatist metasemantics in general. But it does make inferentialism more attractive by comparison for the spacetime functionalist.⁹

⁹ I am grateful to one of the anonymous reviewers for highlighting two dimensions along which I needed

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 proposals. 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) suggestions 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.

4.3 Inferentialism vs. representationalism

When I introduced the distinction between representationalism and inferentialism, I set up the pure versions of these views as two extremes on a spectrum, whose middle-section combined aspects of each. I then further sharpened the characterisation of the middle-section of the spectrum, to leave us with, broadly speaking, four views: (i) pure representationalism, (ii) strongly-anchored inferentialism, (iii) weakly-anchored inferentialism and (iv) pure inferentialism. In this subsection, I discuss two arguments for why spacetime functionalists should embrace inferentialist metasemantics. The first appeals to conceptual motivations of Knox-style constitutive functionalism, call this the conceptual argument. The second invokes the naturalist dilemma, call this the argument from naturalism.

At its most austere, spacetime functionalism works as a demarcation criterion: given some structure, how do we identify it as spatiotemporal?¹⁰ But, as I mentioned in §2, it is motivated by much deeper considerations than mere bookkeeping. Indeed, Knox presents her functionalism about spacetime as being motivated by precisely the conceptual considerations captured in the interpretational challenge. In order for spacetime vocabulary to gain purchase on the physical world, she argues, we should understand the conceptual content of spatiotemporal claims functionally. This is to be contrasted with views according to which, for example, spacetime is characterised by its metaphysical role (i.e. being a container or a stage on which material goings-on unfold) or its mathematical form (i.e. being the thing represented formally by a rank-2 tensor of a certain form). Knox then offers the following: 'once one has analysed the role and understood what fills it, there are no further questions to be asked about the 'real' spacetime structure.' (Knox [2019], p. 4)

But wait a moment. We've just mentioned a 'metaphysical role.' Surely one can construct a functionalism that is interested, not in the spacetime concept, but in spacetime itself? Let us explore this suggestion by distinguishing between metaphysical and conceptual functionalism:

to argue in favour of inferentialist metasemantics for spacetime functionalists: (i) by comparison with other pragmatist positions, and (ii) by comparison with representationalists (whether pragmatist or not). The former forms the basis for this paragraph. The latter is the subject of §4.3.

¹⁰ Note that the object-naturalist might paraphrase this question as 'how do we identify the spatiotemporal truthmakers?' while the subject-naturalist might paraphrase it as 'how do we identify which of our vocabulary is involved in making claims about spacetime?'

Metaphysical functionalism: To be a functionalist about *F*s is to think that there is nothing more to being an *F* (or the property of *F*-ness) than to play some functional role.

Conceptual functionalism: To be a functionalist about *F*s is to think that there is nothing more to the concept of *F* than the role played by that concept.

Recall that for the pure representationalist, the referent of (or more broadly the structure represented by) '*F*' is not at all fixed by the conceptual role of '*F*'. It is fixed by something else. For example, a pure representationalist about *F*-vocabulary could believe that referent of a token of '*F*' is fixed by something like a baptism (Kripke [1972]). The way we learn about *F*s is by actually studying *F*s, not by studying the conceptual role of *F*-vocabulary as it figures in our practice. So if 'spacetime' has a conceptual role, then that role it is determined by what 'spacetime' represents. That is what the target of our inquiry should be, and we make absolutely no headway in our endeavour to understand the world by instead studying the spacetime concept. So representationalism is hostile to conceptual functionalism, and this hostility increases as one moves closer to the pure end of the spectrum, at which point there is a flat-out incompatibility.¹¹

What about inferentialism? Here, things are much more friendly. For the inferentialist, the referent of *F* is fixed (at least in part) by the role played by the *F*. To be an *F* just is to be the thing such that taking our *F*-vocabulary to be about that thing renders our inferences good. So, according to the inferentialist, we can learn about *F*s by looking at the inferential valence (i.e. the conceptual role) of our *F*-claims. So while representationalism is hostile to the conceptual functionalist, inferentialism is deeply cordial. So for the conceptual functionalist, the metasemantic choice is straightforward: inferentialism. This is the conceptual argument.

For the metaphysical functionalist, on the other hand, there is no tension with representationalism. The functional role can be understood straightforwardly in terms of the behaviour of objects that are themselves picked out by some representational relation. For a metaphysical spacetime functionalist, there is no reason to worry about the conceptual roles of our concepts (as they appear in discursive practice), a fortiori, to be an inferentialist.

The inferentialist response is in two parts. The first is to point out that Knoxian (and some related) functionalism is explicitly a conceptual functionalism. So one might read my titular claim as a conditional claim: if you are a conceptual functionalist about spacetime, then you should be an inferentialist. One might leave things there. But here, I offer second part to the argument, for the antecedent of that claim.

In §4.2, I proposed that there is no plausible representationalist account of the anchors for the Knoxian spacetime functionalist. Here, I run a the same argument against the metaphysical functionalist. The metaphysical functionalist fares no better than the anchored inferentialist, as they are vulnerable to exactly the same argument which targets the functional base: it is implausible that a representationalist account can be given of, for example in the Knoxian case, 'dynamical symmetry' and 'force-free'. So the naturalist

¹¹ Well, almost—the one situation where they are not incompatible is the deeply implausible scenario that, by chance, our discursive practices have imbued our vocabulary with precisely the conceptual role that the representation relation establishes, now and forever.

dilemma arises once again: either the metaphysical functionalist gives a representationalist, but non-naturalist grounding to their functional base, or they give a naturalist, but non-representationalist grounding to that base. And once again, the ball is in the representationalist's court to demonstrate how to deal with this dilemma; the inferentialist provides a compelling account of why we should just embrace the second horn. This is the argument from naturalism. Note that this argument does somewhat squeeze out the anchored-inferentialist; it seems implausible (though not impossible) that a representationalist resolution to the naturalist dilemma that saves the anchored inferentialist will not also save the metaphysical functionalist.

5 The Acuña-Myrvold Challenge

Earman ([1989]) 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 drag-alongs applied to absolute tensorial objects, A —which are taken to represent at least some spacetime structure—are automorphisms of the model, while dynamical symmetries are diffeomorphisms whose drag-alongs 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 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. (Myrvold [2019], p. 139)

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. (Brown and Pooley [2006], p. 84)

This sets up the tension that I identified in the introduction, which leads to the following challenge:¹²

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

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. (Myrvold [2019], p. 141)

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. (Acuña [2016], p. 9)

Here, the tension is resolved by changing the subject: the explanandum 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).

My diagnosis of the origin of the Acuña-Myrvold tension is that two separate claims have been made by proponents of the dynamical approach regarding the relationship between dynamical and spacetime symmetries. The first is that there exists an analytic connection between them. The second is that there exists an explanatory arrow between them. It is natural, then, to attribute to the dynamical approach the further claim that the arrow of explanation just is the analytic connection, and consequently, the somewhat awkward claim that an analytic arrow is explanatory. This is a mistake. There are two distinct

¹² 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.

relations that dynamical and spacetime symmetries instantiate. One relation, explanation, is uni-directional and metasemantic, while the other, analyticity, is bi-directional and semantic.

I take both Acuña and Myrvold to endorse the following basic semantic claim, in the domains in which they accept the analytic link: ‘Spacetime symmetry’ means [the same as what] ‘dynamical symmetry’ [means]. Neither of them, however, provides an account of what it is that makes this true. 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 proposal has the advantage of resolving the tension whilst accepting the original terms of the dynamical vs. geometrical debate: that there is an interesting debate to be had about the direction of the arrow of explanation between dynamical and spacetime symmetries in a physical theory. Not everyone agrees with the presupposition. Acuña focusses instead on a distinct explanatory relation. Myrvold sees no place for an explanatory arrow between dynamical and spacetime symmetries. Sus ([2019], [2023]) argues for a mutual-dependency relation that it itself derived from deeper constitutive principles of the theory. For people who hold these, and similar, views, then what I propose below is superfluous when considering the Acuña-Myrvold challenge, since their views allow them to deny one of its presuppositions.

To this I have two responses. The first is to highlight that my argument should be read conditionally: if you accept that there is a debate to be had over the direction of the explanatory arrow between dynamical and spacetime theories, then the tension between two plausible proposals can be resolved by going inferentialist. The second is that, even if you deny the antecedent of my conditional claim, the prising apart of the two arrows is of independent interest, and can even be used to supplement some of these positions, independently of concerns around the Acuña-Myrvold challenge.¹³

¹³ For example, it has been remarked upon by recent commentators (see e.g. (Dorato [2007]; Fellingine [2011]; Read [2020])) that the notion of explanation at play in the dynamical–geometrical debate is somewhat slippery. Now suppose we want to endorse, say, Dasgupta’s ([2017]) suggestion that the form of explanation at play is a constitutive explanation. Dasgupta’s claim is part of a discussion about substantivalism–relationalism, but his characterisation of relationalism, as a form of what (Friedman [1983], p 62) calls ‘Reichenbachian relationalism’ renders the dynamical approach a special case. The

The move that I propose here 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.¹⁴ 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. And since what is doing the work is the existence of a uni-directional metasemantic arrow, independently of what renders the inertial vocabulary meaningful, this account is available to both the anchored- as well as the pure-inferentialist.

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 dynamical 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 dynamics-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. Here, I present an example from string theory, but similar analyses of the dynamics of other quantum gravity theories (loop quantum gravity (Wüthrich [2017]), causal set theory (Wüthrich [2012]; Wüthrich and Huggett [2020])) suggest that the fundamental dynamical degrees of freedom explicitly resist a spatiotemporal interpretation.

In intuitive (but misleading) presentations, String Theory (more accurately, string theories) is sometimes described as a theory of a one-dimensional oscillating quantum string, sweeping out a two-dimensional ‘worldsheet’ as it evolves in spacetime, much like a particle sweeps out a one-dimensional worldline in spacetime. The more careful and accurate presentations highlight that string theories are fundamentally two-dimensional conformal quantum field theories, with no fundamental spatiotemporal properties. The most compelling arguments from both philosophers (Wüthrich and Huggett [2020]) and physicists (Polchinski [1998]) in this context implore us to treat the familiar spacetime of, say,

subject-naturalist spacetime functionalist has a compelling way to understand this explanatory relation as metasemantic: it is constitutive of the meanings of spacetime and dynamical symmetries that they are inferentially related in a particular way.

¹⁴ I take passages like the following as evidence that Myrvold is tacitly representationalist: ‘Spacetime substantivalism... is perfectly intelligible, it will be said, whether one takes it to be true or false. I can picture in my mind’s eye a stage on which the events of the world take place... To that it must be replied: pictures and metaphors are all well and good provided there is a non-metaphorical sense that can be attached to them’ (Myrvold [2019], p. 139). For Myrvold, the action is taking place at the level of structures in the world: this is an object-naturalist, a fortiori, representationalist view.

general relativity as emerging from the contingent dynamics of the worldsheet. Which is to say that there is a host of consistent string theories whose dynamics does not lead to anything like an emergent spacetime. But such string theories are clearly contentful (at least for inferentialists). And in such theories, there is a robust and well-understood notion of a dynamical symmetry; the conformal symmetry of the worldsheet is a prime example.

What this means for the inferentialist is that the talk of dynamics and dynamical degrees of freedom is rendered conceptually coherent whether or not spacetime talk is part of the associated discursive practice. In other words, meaning-conferring links exist between spacetime and dynamical concepts (such as the worldsheet) in all theories that contain spatiotemporal vocabulary. But even in those that do not contain spacetime vocabulary, claims about dynamical entities are rendered meaningful by the inferential links between dynamical vocabulary and other, non-spatiotemporal vocabulary (such as ‘worldsheet,’ ‘worldsheet metric’). This suggests that the correct metasemantic explanatory arrow, in situations in which both sorts of vocabulary appear runs from dynamical vocabulary to spacetime vocabulary: spacetime vocabulary requires dynamical vocabulary in order to be rendered meaningful, but the converse is not the case.

6 Conclusion

Spacetime functionalists should understand their position as best allied with an inferentialist metasemantics. 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 in my opinion 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 separating the 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 representationalism and embraces inferentialism.

Acknowledgements

Thanks to Rachel Fraser, Josh Hunt and Niels Martens for crucial written comments on earlier drafts. To Pablo Acuña, Bob Brandom, Nick Huggett, Wayne Myrvold, John Norton, Oliver Pooley, and James Read for extremely helpful conversations and verbal feedback. To audiences at Pittsburgh and Oxford for discussions and many searching criticisms. (I'll leave it to the reader to guess which group was more sympathetic!) And finally, to two anonymous reviewers whose detailed, specific, and actionable comments and suggestions vastly improved this paper.

*Faculty of Philosophy
University of Cambridge
Cambridge, United Kingdom
tv288@cam.ac.uk*

References

- Acuña, P. [2016]: '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**, pp. 1–12, ISSN: 13552198.
- Baker, D. J. [2020]: 'Knox's Inertial Spacetime Functionalism (and a Better Alternative)', *Synthese*, ISSN: 0039-7857, 1573-0964.
- Barbour, J. B. [2001]: *The Discovery of Dynamics: A Study from a Machian Point of View of the Discovery and the Structure of Dynamical Theories*, Oxford University Press.
- Brandom, R. [1994]: *Making It Explicit: Reasoning, Representing, and Discursive Commitment*, Harvard University Press, ISBN: 978-0-674-54330-0.
- Brandom, R. [2014]: 'Intentionality and Language: A Normative, Pragmatist, Inferentialist Approach', *The Cambridge handbook of linguistic anthropology*, pp. 347–63.
- Brown, H. R. and O. Pooley [1999]: 'The Origins of the Spacetime Metric: Bell's Lorentzian Pedagogy and Its Significance in General Relativity', in C. Callender and N. Huggett (eds.), *Physics Meets Philosophy at the Plank Scale*, Cambridge University Press, pp. 256–72.
- Brown, H. R. and O. Pooley [2006]: 'Chapter 4: Minkowski Space-Time: A Glorious Non-Entity', in D. Dieks (ed.), *The Ontology of Spacetime*, Vol. 1, Elsevier, pp. 67–89.
- Burgess, A. and B. Sherman [2014]: *Metasemantics: New Essays on the Foundations of Meaning*, Oxford University Press, ISBN: 978-0-19-966959-2.
- Chalmers, D. J. [2012]: *Constructing the World*, Oxford: Oxford University Press.
- Dasgupta, S. [2017]: 'Constitutive Explanation: Constitutive Explanation', *Philosophical Issues*, **27**, pp. 74–97, ISSN: 15336077.

- Davidson, D. [2001]: *Inquiries into Truth and Interpretation: Philosophical Essays Volume 2*, Clarendon Press.
- Dennett, D. C. [1989]: *The Intentional Stance*, MIT press.
- Dennett, D. C. [1991]: 'Real Patterns', *The Journal of Philosophy*, **88**, pp. 27–51.
- Dorato, M. [2007]: 'Relativity Theory between Structural and Dynamical Explanations', *International studies in the philosophy of science*, **21**, pp. 95–102.
- Dummett, M. A. E. [1975]: 'What Is a Theory of Meaning?', in S. Guttenplan (ed.), *Mind and Language*, Oxford University Press.
- Earman, J. [1989]: *World Enough and Spacetime*, Cambridge, MA: MIT Press.
- Felline, L. [2011]: 'Scientific Explanation between Principle and Constructive Theories', *Philosophy of Science*, **78**, pp. 989–1000, ISSN: 0031-8248, 1539-767X.
- Friedman, M. [1983]: *The Foundations of Spacetime Theories*, Princeton, NJ: Princeton University Press.
- Gibbard, A. [1990]: *Wise Choices, Apt Feelings: A Theory of Normative Judgment*, Harvard University Press.
- Gomes, H. and J. Butterfield [2022]: 'Geometrodynamics as Functionalism about Time', in C. Kiefer (ed.), *From Quantum to Classical: Essays in Honour of H.-Dieter Zeh*, Cham: Springer International Publishing, pp. 135–78, ISBN: 978-3-030-88781-0.
- Jackson, F. [1998]: *From Metaphysics to Ethics: A Defence of Conceptual Analysis*, Clarendon Press.
- Knox, E. [2011]: '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**, pp. 264–75, ISSN: 13552198.
- Knox, E. [2013]: 'Effective Spacetime Geometry', *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, **44**, pp. 346–56, ISSN: 13552198.
- Knox, E. [2019]: 'Physical Relativity from a Functionalist Perspective', *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, **67**, pp. 118–24, ISSN: 13552198.
- Knox, E. and D. Wallace [2023]: 'Functionalism Fit for Physics', *Unpublished Manuscript*.
- Kraut, R. [1990]: 'Varieties of Pragmatism', *Mind*, **99**, pp. 157–83, ISSN: 0026-4423, JSTOR: [2254958](#).
- Kripke, S. [1972]: *Naming and Necessity*, Springer New York.
- Kuhn, T. S. [1962]: *The Structure of Scientific Revolutions*, University of Chicago press.
- Lam, V. and C. Wüthrich [2018]: *Spacetime Is as Spacetime Does*, arXiv: [1803.04374 \[gr-qc, physics:physics\]](#).
- Lam, V. and C. Wüthrich [2021]: 'Spacetime Functionalism from a Realist Perspective', *Synthese*, **199**, pp. 335–53, ISSN: 1573-0964.
- Lange, L. [1886]: *Die Geschichtliche Entwicklung Des Bewegungsbegriffes Und Ihr Voraussichtliches Endergebnis*. Leipzig: Englemann.
- Lewis, D. [1970]: 'How to Define Theoretical Terms', *The Journal of Philosophy*, **67**, pp. 427–46, ISSN: 0022-362X, JSTOR: [2023861](#).
- Mach, E. [1919]: *The Science of Mechanics: A Critical and Historical Account of Its Development (2nd English (4th German) Ed.)*. Chicago: Open Court Publishing Company.

- McDowell, J. [1997]: ‘Brandom on Representation and Inference’, *Philosophy and Phenomenological Research*, **57**, pp. 157–62, ISSN: 0031-8205, JSTOR: [2953785](#).
- Myrvold, W. C. [2019]: ‘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**, pp. 137–43, ISSN: 13552198.
- Neumann, C. [1870]: *Ueber Die Principien Der Galilei-Newton’schen Theorie*. Leipzig: Teubner.
- Polchinski, J. G. [1998]: *String Theory, Volume I: An Introduction to the Bosonic String*, Cambridge university press Cambridge.
- Price, H. [2004]: ‘Naturalism without Representationalism’, in *Naturalism in Question*, Harvard University Press Cambridge, MA, pp. 71–88.
- Price, H., S. Blackburn, R. Brandom, P. Horwich and M. Williams [2013]: *Expressivism, Pragmatism and Representationalism*, Cambridge University Press.
- Quine, W. v. O. [1960]: *Word and Object*, Cambridge, MA: MIT Press.
- Read, J. [2020]: ‘Geometrical Constructivism and Modal Relationalism: Further Aspects of the Dynamical/Geometrical Debate’, *International Studies in the Philosophy of Science*, **33**, pp. 23–41, ISSN: 0269-8595, 1469-9281.
- Read, J. and T. Menon [2021]: ‘The Limitations of Inertial Frame Spacetime Functionalism’, *Synthese*, **199**, pp. 229–51, ISSN: 1573-0964.
- Sellars, W. [1954]: ‘Some Reflections on Language Games’, *Philosophy of Science*, **21**, pp. 204–28.
- Sellars, W. [1974]: ‘Meaning as Functional Classification’, *Synthese*, **27**, pp. 417–37, ISSN: 1573-0964.
- Stalnaker, R. C. [1997]: ‘Reference and Necessity’, in B. Hale and C. Wright (eds.), *A Companion to the Philosophy of Language*, Blackwell.
- Sus, A. [2019]: ‘Explanation, Analyticity and Constitutive Principles in Spacetime Theories’, *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, **65**, pp. 15–24, ISSN: 1355-2198.
- Sus, A. [2023]: ‘What Spacetime Does’, *Theoria: An International Journal for Theory, History and Foundations of Science*, **38**, pp. 67–85.
- Wittgenstein, L. [2009]: *Philosophical Investigations*, John Wiley & Sons, ISBN: 978-1-4443-0797-9.
- Wüthrich, C. [2012]: ‘The Structure of Causal Sets’, *Journal for general philosophy of science*, **43**, pp. 223–41.
- Wüthrich, C. [2017]: ‘Raiders of the Lost Spacetime’, *Towards a theory of spacetime theories*, pp. 297–335.
- Wüthrich, C. and N. Huggett [2020]: ‘Out of Nowhere: The Emergence of Spacetime from Causal Sets’, *arXiv preprint arXiv:2009.02951*, arXiv: [2009.02951](#).