Stating Maths-First Realism: Some Tentative Proposals

This paper is based on a talk presented in January at the Oxford Philosophy of Physics seminar and the Bristol Philosophy of Physics seminar. I thank those in attendance for their questions and comments. In addition, I would like to thank Henrique Gomes, Simon Saunders and David Wallace for further feedback. As the title indicates, the ideas expressed here are *tentative, and I hope to discuss them in more detail in future work.*

Galileo wrote that mathematics is the language in which God has written the book of the universe. This simple fact is the *raison d'être* for philosophy of physics, for while sentences of non-mathematical languages directly describe the world around us, the 'sentences' of mathematics only describe abstract objects. Unless the world literally *is* mathematics, we need more than just the bare text of definitions and equations to understand what the book of the world says about – the world. What we need are *interpretations* of physical theories: "If the primary task of the physicist is to construct models and theories of the world, the primary task of a philosopher of physics is to *interpret* these products of physics, be they theories, models, simulations, or experiments" (Rickles 2016). This, then, is "*the question of interpretation*: under what conditions is the theory true? What does it say the world is like?" (van Fraassen 1991, 242).[1]

Nevertheless, mathematics *is* a language. Why, then, does David Wallace (2022), in a recent paper on structural realism, pit a 'language-first' approach to theories against a 'mathematics-first' approach? It's language all the way down! Of course, this is a facetious objection; it is clear that Wallace does not mean to eschew any form of language, mathematics included, nor does he simply want to trade one language for another. Wallace rather claims that the interpretation of a theory written in the language of mathematics need not proceed via a reconstruction of that theory in an already-understood language. On his favoured maths-first[2] approach, mathematical theories *already* represent the world in a certain way. This approach is supposed to yield a version of structural realism that does away with traditional questions of ontology and ideology. But *how* does mathematics represent the world? And what does it say the world is *like*? Is there *really* no need for language?

Wallace acknowledges that his characterisation of the maths-first approach is only a first draft: "many details remain to be filled in, and the devil may be in those details." Part of this essay consists in an attempt to fill in said details. In doing so, I will inevitably introduce my own idiosyncrasies into Wallace's framework. It is my hope that this essay nevertheless clarifies the contours of the maths-first approach.

But I also want to sound a sceptical note: the maths-first approach may not succeed in discharging the duty of an interpreter to say what the world would be like if the theory were true. As I will show, all extant approaches to scientific representation – including those that adopt the so-called 'semantic view of theories' – must rely on linguistic means to answer the

[1] This is not meant to invoke the Cosmological Assumption, that a theory describes the entire universe, or the Fundamentality Assumption, that a theory is exactly true (Wallace 2019); just replace van Fraassen's question with: "what would the subsystem be like if the theory were approximately true of it?".
[2] With apologies to Wallace, I will stick to the British way of abbreviating 'mathematics'.

question of interpretation. Because of this, they cannot offer the kind of structural realism that shies away from metaphysics. Insofar as the maths-first approach is simply a re-articulation of the semantic view of theories, it collapses back into the language-first approach.

The maths-first approach must therefore offer a novel account of scientific representation. I sketch two options in this essay, but I wonder whether either can truly answer the question of interpretation. On the first option, that question is answered by the mathematics itself, which is therefore not like a foreign language in need of radical interpretation but like one that we already speak. This belies the fact that the project of interpretation *does* seem necessary to understand what claim a theory makes on the world. On the second option, the question does not require an answer at all: it suffices to say that the world is such that it *is* accurately represented by our theories – no matter what that is! This form of quietism may seem to shirk the duty of interpretation. We will see that both options necessitate a reconception of the purpose of philosophy of physics away from interpretation.

Syntax versus Semantics, or Language versus Mathematics

Let's start with the distinction between the syntactic and semantic view of theories. It is often said that on the syntactic view theories are presented as collections of *sentences*, whether in English, Esperanto, or first-order logic (FOL), whereas on the semantic approach they are presented as collections of *models*, which are abstract mathematical structures. Recent literature has shown *this* distinction to be elusive (Halvorson 2013, Lutz 2015, Hudetz 2019). On the one hand, any collection of sentences is satisfied by some collection of models; on the other hand, any collection of models satisfies some collection of sentences. Moreover, many notions employed by the semantic view have purely syntactic correlates, from empirical adequacy (Lutz 2015) to theoretical equivalence (Hudetz 2019). If this is the crucial difference, then it is only one of emphasis.

Wallace points out that such equivalences only hold for *axiomatised* theories, not for the more loosely-defined theories of real-life physics. Nevertheless, I contend that the import of the syntactic-semantic distinction is found elsewhere.

The true debate lies in the way theories make contact with the world. On the syntactic view, a theory's sentences directly describe reality. When Newton wrote that all bodies move uniformly unless acted upon, the term 'bodies' referred to material bodies, 'uniform motion' to physical movement, and 'acted upon' to real forces. The models of syntactically-formulated theories are interpretations in the logician's sense: they assign extensions to predicates. The domains of these models are populated by real-world objects. On the semantic view, meanwhile, sentences define the models themselves. The domains of such models consist of abstract objects – mere 'placeholders' (Shapiro 1997). There is not yet any contact with the world. This only comes after the models are taken to represent the world in a certain way. In Wallace's words: "The theory/world relation here is *representation*, more akin to the relation between map and territory than that between word and object." I will take *this* as the hallmark of the semantic view of theories.

Despite the criticisms mentioned above it still seems that the semantic view has many benefits, which Wallace eloquently illustrates: it allows for clearer notions of approximation, theoretical equivalence and inter-theoretic reduction. To illustrate just the first: sentences are either true or false, so it is difficult to define a notion of approximate truth. Models, by contrast, may map onto reality more or less perfectly. In addition, advocates of the semantic view have argued that it allows for a clearer demarcation of a theory's observable content (in terms of substructures) and of its structure (in terms of isomorphisms). This compares favourably to the syntactic view's failure to define a distinction between theoretical and observational vocabulary, as well as the near-triviality of the Ramsey sentence as a syntactic account of structure. But the possibility of a demarcation between the observable and unobservable, or between structure and nature, does not automatically entail a form of empiricism or structural realism. The semantic view may help one to articulate what one is a realist about, but it does not *by itself* commit one to any form of realism.

It is here that the maths-first approach seems to depart from the semantic view as standardly conceived in a more radical direction. According to Wallace, "we should see math-first scientific realism as *already* a form of structural realism, without any need for further selective scepticism: the move from language-first to math-first formulations of theories […] is itself sufficient to make this form of realism 'structural'." It is here that I want to give pause. If this extraordinary claim is true, what kind of structural realism does the maths-first approach commit one too? What is the world like, if maths-first realism is correct?

Wallace is relatively silent on the matter. He is clear on what maths-first realism is *not* committed to: "considerations of ontology and ideology", "our ordinary talk of objects, properties and relations", as well as "overly-fine metaphysical distinctions" such as those required to cash out talk of 'relations without relata'. The statement of maths-first realism itself is succinct:

> **Math-first scientific realism**: Our successful scientific theories succeed, at least approximately, in representing the systems which they are used to model, including features of those systems that are unobservable.[3]

One of the supposed advantages of this view is that it has easy responses to the standard objections to scientific realism: underdetermination and theory change. The most serious cases of underdetermination, Wallace claims, concern mathematically equivalent theories, but maths-first realism is only committed to the representational adequacy of a theory's mathematical structure anyway. And since mathematical structure is approximately preserved under theory change, the empirical success of past theories does not pose a problem for maths-first realism either.

However, the fact that a theory's models represent certain systems does not tell us what they represent them as being *like* – it does not yet answer the question of interpretation. It is useful

[3] I have so far remained neutral between realism and empiricism; for my purposes the clause "including features of those systems that are unobservable" is relatively unimportant (but see the discussion of 'descriptivism' below). Wallace also distinguishes between maths-first *epistemic* structural realism and maths-first *ontic* structural realism. On the former view, there is more to the world than the structure represented by a theory's models, but we cannot know about it. On the latter view, the structure represented by a theory's models is all there is. Since my question is what a theory's models represent the world as being like, this difference is also irrelevant for my purposes.

here to note that "theories, unlike maps, are always representations *as*" (Hughes 1997; cf. van Fraassen 2008, 23). To say that a model is a representation *of* a target simply means that it is *about* the target. But models do not *just* represent physical systems; they represent them *as being a certain way*. If this were not the case, one would in effect have a form of instrumentalism: a model represents certain systems and yields predictions about their behaviour, but it does not make any claims about what those systems are like. Chakravartty neatly summarises this point: "Theories are not merely objects which replicate or imitate the phenomena; even if they are such things, they are meant, in addition, to *tell us something substantive* about the nature of the world. The obvious question, then, is how is this achieved?" (2001, 330).

It is this question that I want to ask of the maths-first approach. How does the mathematics of a theory tell us something about the world, without considerations of objects, properties, relations, and the metaphysical distinctions that come with them?

Scientific Representation, Four Ways

In the remainder of this essay, I will sketch four ways in which theoretical models can make contact with the world. These are not meant to exhaust the range of possibilities. It may turn out that there is another, superior option, but to me these four seem the main contenders.

I tentatively classify the first two – *linguicism* and *descriptivism* – as language-first approaches. But only the first one assumes the syntactic view of theories; the latter concerns the standard account of scientific representation on the *semantic* view. Nevertheless, I believe that descriptivism does not deliver the kind of structural realism promised by the maths-first approach. If the latter really is just the semantic view of theories, then we are back to where we started.

The latter two options – *nativism* and *quietism* – *are* plausible candidates for the maths-first approach. Both seem implicit in Wallace's exposition. The positive aim in this part of the paper is to spell out in some more detail how maths-first realism could answer the question of interpretation. But I am not convinced that either approach *can* succeed. The first is at risk of collapsing back into a language-first approach, while the latter seems to reject rather than answer the demand for interpretation.

Linguicism

The first option says that *in order to interpret a mathematical theory it must be* translated *into a non-mathematical language*, such as English or FOL. I call this approach *linguicism* because it discriminates on the basis of language: non-mathematical languages are preferred over mathematical languages for the purpose of theory-interpretation. The rationale for linguicism is that the interpretation of sentences in English or FOL is relatively clear-cut, with relatively well-defined notions of reference, truth and meaning. Why settle for less?

Of course, linguicism presupposes the syntactic view of theories. If our theories aren't already presented as collections of sentences, we should modify them until they are. The objections to that view are well-known: the untenability of the distinction between theoretical and observational vocabulary; the troublesome interpretation of theoretical terms in terms of correspondence sentences; the problem of theory-individuation. More than that, the entire project of translation seems misguided, given the enormous success of our mathematically-formulated theories. The effectiveness of mathematics is perhaps unreasonable, but that is no reason to forsake it.

Is linguicism a language-first approach? It seems so. Wallace's concept of a *predicate precisification* is helpful here: "a predicate precisification of a mathematically-given physical theory is a presentation of that theory in language-first style, in terms of objects, their properties and their relations." It is unclear whether a 'presentation' here is the same as a *translation*, but they don't seem far apart. Wallace attributes to the language-first view a commitment to predicate precisifications: "'theories', on the language-first view, are predicate precisifications of mathematically-presented theories." So, one natural interpretation of the maths-first approach is that it constitutes a rejection of linguicism.

If this is the case, there is not much more to the maths-first approach than the adoption of the semantic view. This is not to say that Wallace's exposition does not contain any novel contributions; in particular, it emphasises the role of *mathematical* (rather than set-theoretic) structure on the semantic view. As I will discuss below, however, the semantic view by itself does not rid interpretation of the appeal to language. Even on a semantic conception of theories, interpretation may require talk of objects, properties and relations. If maths-first realism is to offer up a distinctive form of structuralism that moves beyond metaphysics, it has to go further.

There is one area of philosophy in which linguicism does seem prevalent: analytic metaphysics. Debates over whether the existential or universal quantifier is more fundamental, or whether my fist is a different object from my hand, seem predicated on the assumption that the book of the world is written in FOL. There are reasonable defences of this practice: perhaps subject matters of metaphysics, such as mereology and modality, *are* best expressed in FOL, unlike the subject matter of physics; or perhaps these are mere toy theories that will eventually find full their expression in a maths-first framework. Nevertheless, insofar as analytic metaphysics is committed to linguicism the maths-first approach is a welcome rejoinder.

But even if the structural realism supposedly inherent to mathematics is extended all the way to metaphysics, the question of interpretation remains unanswered: what does a mathematically-presented theory, whether of physics or metaphysics, say the world is like?

Descriptivism

The second option to answer this question is *descriptivism*: *in order to interpret a mathematical theory it must be* described *in a non-mathematical language*, such as English or FOL. The models of a theory are couched in a non-mathematical vocabulary that endows them with content. On the descriptivist view, mathematics is a foreign language but translation is not the correct approach; a kind of extended commentary is.[4] In Chakravartty's (2001) words: "a model can tell us about the nature of reality only if we are willing to assert that some aspect(s) of the model has a counterpart in reality. That is, if one wishes to be a realist, some sort of explicit statement asserting a correspondence between a description of

[4] Cf. Bailer-Jones' (2003) remarks on pictorial representation.

some aspect of a model and the world is inescapable. This requires the deployment of linguistic formulations, and interpreting these formulations in such a way as to understand what models are telling us about the world is the unavoidable cost of realism."

Crucially, descriptivism is compatible with the semantic view of theories: the descriptivist admits that theories are presented as classes of models, but adds that for those models to represent the world, a linguistic element is necessary. Chakravartty, again: "the emphasis in the semantic view on models does nothing to eliminate the currency of sentences so far as the project of determining how aspects of theories might be cashed out as literal representations of the world is concerned." French and Saatsi (2006), explicit advocates of the semantic view, write in a similar vein that "there is more to the semantic approach than pure logico-mathematical structures; after all, we speak of particular models representing the unobservable world behind *particular* phenomena, we *interpret* theories by *describing* the properties and relations the state variables in a model stand for, and so forth." Indeed, extant accounts of scientific representation on the semantic view invariably adopt some form of descriptivism. Below follow some examples.

First, consider the *similarity conception*. It is a tempting thought that models represent a system in virtue of their similarity to it. This intuition is especially clear for physical models, such as a scale model of an airplane. But the same idea may apply to mathematical models. The well-known problem with this account is that anything is similar to anything else in *some* respect. The standard solution is to let model-users propose *theoretical hypotheses* that specify the particular respect in which a model is similar to the system: "It is not the model that is doing the representing; it is the scientist using the model who is doing the representing. One way scientists do this is by picking out some specific features of the model that are then claimed to be similar to features of the designated real system to some (perhaps fairly loosely indicated) degree of fit" (Giere 2004). But a theoretical hypothesis is a linguistic entity! It describes which parts of the model are intended to represent which parts of the world. So, the similarity conception relies on descriptivism.

Secondly, the *structuralist* conception. On this view, models represent a system in virtue of some *morphism* (isomorphism, partial morphism, etc) between them. Wallace seems sympathetic to this view: "empirical adequacy [of a theory] consists of a partial isomorphism of the data model into the theoretical model." Like similarity, however, (iso)morphism is too flexible: almost any system instantiates *some* structure that is isomorphic to a chosen model. This is Putnam's infamous paradox. One can respond to this paradox in one of two ways, but both involve description. The first is to use an already-understood language to specify the intended isomorphism: *this* part of the model should map onto *that* part of the world. This is van Fraasen's (2008) favoured approach, but its reliance on linguistic resources is clear. The second option is to describe the world itself in a certain way in order to pick out the intended structure. This equally clearly involves language. The advocates of the structuralist conception are clear on the need for language here. French and Saatsi, for instance, admit that "it is indeed the linguistic complement of the semantic framework that allows us to sidestep the problem of unintended models." So, the structuralist conception also relies on description.

The same is the case for many other mainstream approaches to scientific representation. For example, Contessa's (2007) version of the *inferentialist* conception concerns an 'analytic interpretation' by which a user takes a model's elements to denote objects, properties and relations. On the *inferentialist-expressivist* account of Khalifa, Millson and Risjord (2020), too, "model elements are given interpretations using the resources of language." Likewise, Nguyen and Frigg's (2021) *extensional abstraction account* requires 'structure generating descriptions': "descriptions of physical systems that refer to physical objects and physical properties and relations." And on Frigg and Nguyen's (2020) DEKI account, the notion of a 'key' fulfils a similar role of "translation manual" between properties of the model and properties of the system. Although this does not exhaust the entire literature on scientific representation on the semantic view, I am not aware of any plausible account that is not a form of descriptivism.

Is descriptivism a maths-first or a language-first approach? Although it may not violate the letter of the maths-first approach – theories are not presented as collections of sentences – it surely violates the spirit. The descriptions of models are similar to Wallace's predicate precisifications, which he also characterises as "description in language […] of a theory". But the promise of maths-first realism is precisely that it does not need such descriptions!

Moreover, descriptivism does not deliver the brand of structural realism claimed by the maths-first approach. For descriptions of a theory's content will employ exactly the subject-predicate talk that incorporates distinctions between objects, properties and relations. French and Saatsi (2006) admit as much: "The theoretical variables for which the simultaneous values and their change are given by the structure are theoretically interpreted: they refer to physical properties and relations." The structural realist can of course try to isolate the structural part of a description from its non-structural part. Thus, French and Saatsi point out that successful descriptions only require reference to theoretical properties, such as charge, and not to entities such as electrons. But this is just the kind of approach that leads to "surprisingly metaphysical conversations about how we can have 'relations without relata'"; conversations that the maths-first approach had hoped to avoid!

Another way in which descriptivism stands in the way of the ambitions of maths-first realism is that theories that are equivalent by maths-first standards may yet receive incompatible linguistic descriptions. The threat of underdetermination of theory by data is one of the main motivations for the move to structural realism, and maths-first realism promises to answer this threat in a particularly simple way: putative cases of empirically equivalent theories are also cases of mathematically equivalent theories, so they must have the same physical content. Since these theories' linguistic components may differ, however, descriptivism cannot avoid underdetermination that easily.

I can think of two ways in which descriptivism can accommodate maths-first realism. The first is to restrict the class of descriptions on offer. It would seem that even maths-first realism has to admit of *some* descriptive components in order for the theory to connect to *observable* matters. Wallace's description of N-particle Newtonian mechanics on the maths-first approach, for instance, posits "a list of N positive real numbers m_1, ... m_N, *representing the particle masses*" (emphasis mine). The claim that the m_n represent particle masses is linguistic, but presumably such claims are necessary in order to endow the theory's models with empirical content in the first place.

If the maths-first approach were to restrict descriptions *just* to such observable matters, it would collapse into a form of instrumentalism. But it can offer a finer criterion for which descriptions are acceptable by means of the notion of mathematical equivalence: demand that equivalent theories receive equivalent descriptions. Unfortunately, this does not deliver a metaphysics-free version of structural realism either. For one, an entire class of equivalent theories may yet receive distinct descriptions that differ over the theory's ontology and ideology; substantivalist and relationist interpretations of spacetime theories are just one case in point. Moreover, the search for a common description of equivalent theories may involve the exact kind of metaphysics that maths-first realism wishes to rejects. Consider the interpretation of mathematically equivalent models of the same theory, which is not too dissimilar from that of equivalent theories themselves. In the case of spacetime symmetries, this has involved elaborate discussions of the primitive identities of spacetime points: a form of structuralism, perhaps, but one that does seem to trade in "overly-fine metaphysical distinctions". I don't see any reason to suppose that the search for a description of the common core of, say, the AdS/CFT correspondence would not utilise similar distinctions.

The second option is a division of labour. The semantic view excels at approximation, theoretical equivalence and inter-theory reduction. These aspects of interpretation don't require any linguistic means. The maths-first realism could thus keep such issues language-free, and only resort to descriptions when it comes to the question of what the theory says the world is like. This approach is suggested by French and Saatsi: "The structuralism of the semantic approach resides in its insightful emphasis on the structural complement of theories vis-à-vis the question of theory identity, inter-theory relations, and theory-data relations. Questions regarding these are nicely handled in a decidedly structural, holistic manner by focusing on the structural description of the variables open to interpretational descriptions." This is a form of structuralism about approximation and reduction, but unfortunately it is not a structural *realism*: it does not attempt to say what the world is like in purely structural terms.

On either option, then, descriptivism cannot deliver on the maths-first approach's promise of a metaphysics-free structural realism. Since descriptivism endorses the semantic view of theories, this means that Wallace is mistaken when he writes that "the move from a 'syntactic' to a 'semantic' conception of theories is itself more or less sufficient to turn standard realism into structural realism." In order to secure a maths-first structural realism, the semantic view by itself is insufficient. In the next two subsections, I will critically consider some more radical accounts of scientific representation that may offer a truly maths-first version of structural realism.

Nativism

The first option for maths-first realism is *nativism*: *a mathematical theory* already *describes reality,* without the intervention of language. On this view, mathematics is not a foreign language in need of translation or interpretation; it is one that we already can, or at least can learn to, speak. The mathematical structures defined by a theory are like the primitives of any other language in that they are *about* the world. Of course, the reference-relation here is rather different from that of English or FOL. The important point is that we can wield these structures directly to describe the world.

Wallace seems to suggest nativism when he writes that "there are other ways to represent features of the world than in language (maps, art, and the representational practices of nonlinguistic animals, for instance, even leaving aside physical theories)." Does nativism provide a natural notion of scientific representation on the maths-first approach?

I am not opposed to this proposal in principle, but as it stands it lacks details – and there is some reason to believe that those details will prove difficult to provide.

Consider the comparison to maps and art. It is true that the painter of a still life need not label the apple as 'apple' and the skull as 'skull' in order for the viewer to know that those are the objects depicted. Likewise, although it is customary for a map to have an 'I am here' label, one can also use distinctive features of the environment such as major roads and intersections to locate oneself on the map. In either case, representation is established without translation or description. Could mathematical representation work like that? The chances seem slim. Both the still life and the map irreducibly rely on visual perception (Gombrich 1961, Lopes 1996). I use the same perceptual resources to recognise an apple and a picture of an apple. But abstract mathematical structures are invisible: we cannot just 'see' what they represent.

In the terminology of Suppe (1977), some models, such as scale models, function as 'icons' of whatever they represent. It does not seem that mathematical models are iconic in this sense. But even if they were, this would not answer the question of interpretation. "The problem immediately arises," as Da Costa and French (1990) put it, "that on the realist view the mathematical model is in some sense supposed to say 'how the world is', whereas the iconic model does not, in the usual discussions of these things." So, iconic representation is not a fruitful avenue for maths-first realism.

What about the 'representational practices of nonlinguistic animals'? Here, too, the analogy is less than perfect. Donald Davidson famously defended the controversial Cartesian thesis that thought requires language, so animals cannot think. More precisely, on Davidson's view animals don't have propositional attitudes, such as belief. On some more recent accounts, animals *can* think but their thought is non-propositional (Camp 2009) or non-conceptual (Beck 2012). It seems to me that neither option is particularly helpful for maths-first realism, which presumably would want to say that theories warrant belief that has conceptual content. It is not enough to show that non-linguistic representation is possible; we also need the right *kind* of representation.

Perhaps the point here is simply to make plausible that *some* form of non-linguistic representation is conceivable, to which mathematical representation is ultimately reducible. This would fit well with the 'Gricean' view defended by Callender and Cohen (2006). Griceanism comes in two flavours: general and specific. The general view only states that all forms of representation are reducible to one fundamental form of representation. General Griceanism suffices to establish Callender and Cohen's claim that scientific representation is no different from any other kind of representation (except the fundamental one), but it does not necessarily support the maths-first approach. On the specific view, meanwhile, a specific account of the reduction base is offered, namely mental states – intentions and beliefs. Roughly, a representational device D means that p iff the user U uses D with the intention to active in a receiver R the belief that p. In the final account, these mental representations are themselves explained in entirely non-representational terms. There are various suggestions for how to achieve this, but for my purposes the relevant issue is the reduction of scientific representation to mental states.

There are a couple of potential problems with Specific Griceanism. Firstly, the reduction base. It is simply not clear that the mental states involved *are* non-linguistic. We have already seen that some philosophers, such as Davidson, believe that propositional attitudes invariably require language. Fodor's Language of Thought hypothesis likewise posits a language-like structure for mental states. Of course, such claims are highly controversial. But it remains to be seen whether the maths-first approach can give a coherent account of non-linguistic mental representations – as opposed to non-linguistic representations in general.

Secondly, the reduction itself. It is one thing to claim that all representation reduces to one fundamental form of non-linguistic representation; it is another to show how. On Callender and Cohen's account, the crucial ingredient here is *stipulation*. In their example, if I (publicly) stipulate that the salt shaker represents Madagascar, then I can use the former to produce in someone a mental state about the latter. These stipulations are themselves linguistic; Callender and Cohen already note the similarities between their view and Giere's theoretical hypotheses (although they explicitly reject a similarity-based account of representation). A language-free reduction base is of no help if the reduction itself proceeds by means of language. Again, this is not to say that this cannot be done; but it is one thing to show that *some* form of non-linguistic representation is possible, and quite another to show scientific representation reduces to it.

One possible avenue for nativism here is a *pragmatic* account of representation, on which representation is ultimately reduced to the way particular representational devices are *used*. Many accounts of representation feature some form of use-condition, but often these are just language in disguise. For example, on Giere's (2010) agent-based conception a model represents a system partly in virtue of the user's intention that the model represent the system. These intentions are themselves linguistic, however: on Giere's account, such an intention simply consists of the adoption of one or more theoretical hypothesis. It surely doesn't matter whether such hypotheses are explicitly written down, spoken aloud, or are left implicit. At the other hand, some use-based accounts move too far into the non-representational direction. On Knuuuttila's (2005) view, for example, "emphasis on representation places excessive limitations on our view of models and their epistemic value. Models should rather be thought of as epistemic artifacts through which we gain knowledge in diverse ways." Insofar as maths-first realism *is* interested in representation, this is not satisfactory either. But I can conceive of an account on which the representational capacities of mathematical models are grounded in the way they are used by the scientific community, without the need for language. The nativist's best hope may be for such an account to succeed.[5]

Let's take a step back. The question of the nature of mental representation is wide open, so it is not inconceivable that a non-linguistic type of mental representation does exist and can

[5] The best candidate I know of is the 'inferential scientific realism' currently under development by Tushar Menon (2024). Unlike Khalifa et al.'s inferentialist account, this view promises to fully ground the representational capacities of scientific theories in the inferential practices they license.

serve as a basis for scientific representation. And Wallace is of course correct when he writes that linguistic representation isn't understood perfectly well either. But notice just how many questions the maths-first realist has to answer for this approach to work: (1) whether a sufficiently thick form of nonlinguistic representation is possible; (2) whether it is a kind of representation we can employ; (3) how scientific representation reduces to this form of representation; (4) whether that reduction it itself entirely free from language; and (5) if the final concept of representation is structural in the relevant sense. To belabour the latter point, consider theoretical equivalence. It is an important part of maths-first realism that equivalent theories, in virtue of their shared mathematical structure, represent the world in the same way. This is what allows maths-first realism to avoid the problem of underdetermination. But a reduction of scientific representation to some fundamental form of non-linguistic representation does not entail that mathematically equivalent theories are *representationally* equivalent. It is quite possible that such theories are used in different ways, or that they license distinct inferences. The project of non-linguistic representation *per se* therefore does not automatically yield a *structural* account of representation. Again, the devil *is* in the details: it isn't clear whether nativism is even compatible with maths-first realism before such details are worked out.

Quietism

Is there an easier way out for maths-first realism? Yes, there is. The final option I will consider is called *quietism*, which says that *beyond the fact* that *our theories (approximately) represent (some part of) reality*, *there is no more to say on the question of interpretation*. On this view, the official statement of maths-first realism is taken seriously: our theories approximately represent the systems they are used to model, and that's that. There is no further answer to the question what those theories represent the world as being like.

Quietism does not so much reject the question of interpretation as reconceive what an answer could look like. It believes that one can – and perhaps *must* – say what the world is like *in terms of a representation of that world*. What would the world be like if theory T were true? It would look like it was accurately represented by one of T's models. There is no need for a 'semantic descend' towards a description that does not refer to representations. Compare this to the case of ordinary sentences. Call the sentence 'the particle is massive' *S*, and suppose that it is a theorem of T. Then one could say that the particle is such that S is true of it. By the standard rules of disquotation, this is just to say that the particle *is* massive. On the quietist attitude, there is a sense in which such disquotation is impossible for theoretical models. We can only describe the world in terms of our representations, never outside of it.

This may sound like capitulation. But the quietist's response is that the apparent victory of language-first approaches is illusory. Dewar (2023) puts it like this: "when we look at any specific such account, what form does it have? Well, just some more sentences – that is to say, just more theory!" On descriptivism, for example, it may seem as if some sentences can provide a successful interpretation of a theory. But those sentences just become part of the theory in total, which means that they, too, stand in need of interpretation. If this is correct (and I am not yet convinced that it is), quietism is the *only* option.

What, then, does interpretation look like for a quietist? Dewar suggests one can use the equivalence between theories to extract content from them. This is the upshot of Dewar's project of *internal interpretation.* On the received view, which Dewar calls *external interpretation*, the first step is to set up a map between the theory and the world. This map endows a theory with content, so two theories are equivalent whenever they receive the same interpretation. On the internal picture, on the other hand, "we begin by making determinations of equivalence, and use those determinations to get a fix on the theory's commitments." In other words, Dewar suggests that we can stipulate which theories are equivalent and *thereby* interpret them. This is illustrated with the case of symmetry-related models within a theory: if such models represent the same state of affairs, then quantities that vary under the relevant symmetries, such as absolute velocity, are not physically real. Thus, a stipulated equivalence determines the theory's commitments. It is not clear yet that the process of internal interpretation automatically reduces to a form of maths-first realism – Dewar himself is happy to discuss the metaphysical doctrine of haecceitism in this context, for example – but it does make a quietist approach more plausible.

In addition, the advocate of quietism may point out that much of philosophy of physics already *is* concerned with representations, rather than their interpretation. If we want to know whether a theory is deterministic, or whether it is local, or discrete, we need only look at a minimally-interpreted version of the formalism. Of course, this requires *some* interpretational work – for example, the identification of a time coordinate to define determinism. But once that is done, determinism itself is just a mathematical property of the theory. Or consider the claim that unitary quantum mechanics instantiates a branch-like structure of approximately classical worlds. This follows from technical results about quantum mechanics, such as Ehrenfest's theorem and decoherence theory; it does not require much interpretation. When we ask what such a classical world is *like*, however, quietism is silent. All it can say is that it is such that we can use classical mechanics to accurately represent it. Nevertheless, the fact that such a world is recoverable from quantum theory is non-trivial.

The advocate of quietism thus believes that the rich network of intra-theoretic and inter-theoretic relations of equivalence, approximation and reduction suffices to get a grip on the world. This may strike those who wish for a more substantive account of what the world is like as a Pyrrhic victory. I cannot settle this debate here. The point I wish to emphasise is that for maths-first realism to work, it may need to become an internal realism, too.

Conclusion

The maths-first approach may sound like good news for philosophy of physics: we can finally focus on the theories themselves, unconstrained by the metaphysicians' excessive demands for formalisation in FOL! But I am not so sure. In this essay, I hope to have shown that maths-first realism is not simply a consequence of the semantic view of theories, but requires a more radical account of theory-interpretation. Where does such an account leave philosophy of physics?

As I see it, much of contemporary philosophy of physics is incompatible with maths-first realism. This is true of the most metaphysically-inclined strands of the field, such as the primitive ontology movement, but also of those of us that are in the business of trying to describe in ordinary language what our best physical theories say about the world. I take it that most philosophers, like me, don't find it easy to understand what those theories say about the world – that is why we feel the need to interpret them in the first place!

On any account of interpretation compatible with maths-first realism, however, philosophy of physics becomes a radically different discipline. If nativism is true, mathematical models acquire content by the way in which they are used in practice. The aspirant-interpreter should therefore engage in a project of radical translation by studying the way in which *physicists* use their theories – or 'go native' and join the ranks of scientists themselves. On a quietist view, philosophers of physics need not re-train, but the scope of their inquiry is limited. Given that the theories under study *are* successful, it suffices to say that they accurately represent those parts of reality within their domain. The questions that remain concern formal aspects of theories and the relations between them. There is much work of interest in this area, for example on surplus structure and on theoretical equivalence. But it relinquishes the traditional concerns of interpretation: what does a theory say that exists, and what does it say *of* that which exists?

In the face of these consequences, the advocate of maths-first realism will say: so much the worse for traditionalism. In any case, it is inappropriate to choose one's form of realism on the basis of one's methodological preferences. Nevertheless, it is important to be aware of these consequences. The maths-first approach is not just a neat statement of what philosophy of physics was about any way; it is a radically novel vision on the interpretation of theories, for better or for worse.

Bibliography

Bailer-Jones, D. M. (2003). When Scientific Models Represent. *International Studies in the Philosophy of Science*, 17(1), 59–74. https://doi.org/10.1080/02698590305238

Beck, J. (2012). The Generality Constraint and the Structure of Thought. *Mind*, 121(483), 563–600. https://doi.org/10.1093/mind/fzs077

Callender, C., & Cohen, J. (2006). There is No Special Problem About Scientific Representation. *Theoria: Revista de Teoría, Historia y Fundamentos de La Ciencia*, 21(1), 67–85. https://doi.org/10.1387/theoria.554

Camp, E. (2009). A language of baboon thought? In R. W. Lurz (Ed.), *The Philosophy of Animal Minds* (pp. 108–127). Cambridge University Press. https://doi.org/10.1017/CBO9780511819001.007

Chakravartty, A. (2001). The Semantic or Model-Theoretic View of Theories and Scientific Realism. *Synthese*, 127(3), 325–345. https://doi.org/10.1023/a:1010359521312

Contessa, G. (2007). Scientific Representation, Interpretation, and Surrogative Reasoning*. *Philosophy of Science*, 74(1), 48–68. https://doi.org/10.1086/519478

Costa, N. C. A. da, & French, S. (1990). The Model-Theoretic Approach in the Philosophy of Science. *Philosophy of Science*, 57(2), 248–265.

Davidson, D. (1982). Rational Animals. *Dialectica*, 36(4), 317–327.

Dewar, N. (2023). Interpretation and equivalence; or, equivalence and interpretation. *Synthese*, 201(4), 119. https://doi.org/10.1007/s11229-023-04102-9

French, S., & Saatsi, J. (2006). Realism about Structure: The Semantic View and Nonlinguistic Representations. *Philosophy of Science*, 73(5), 548–559. https://doi.org/10.1086/518325

Frigg, R., & Nguyen, J. (2020). *Modelling Nature. An Opinionated Introduction to Scientific Representation.* Springer.
Giere, R. N. (2004). How Models Are Used to Represent Reality. *Philosophy of Science*, 71(5), 742–752. https://doi.org/10.1086/425063
Giere, R. N. (2010). An Agent-Based Conception of Models and Scientific Representation. *Synthese*, 172(2), 269–281. https://doi.org/10.1007/s11229-009-9506-z
Gombrich, E. H. (1961). *Art and illusion: A study in the psychology of pictorial representation.* Pantheon.
Halvorson, H. (2013). The Semantic View, If Plausible, is Syntactic. *Philosophy of Science*, 80(3), 475–478. https://doi.org/10.1086/671077
Hudetz, L. (2019). The Semantic View of Theories and Higher-Order Languages. *Synthese*, 196(3), 1131–1149. https://doi.org/10.1007/s11229-017-1502-0
Hughes, R. I. G. (1997). Models and Representation. *Philosophy of Science*, 64, S325–S336.
Khalifa, K., Millson, J., & Risjord, M. (2022). Scientific Representation: An Inferentialist-Expressivist Manifesto. *Philosophical Topics*, 50(1), 263–292.
Lopes, D. (1996). *Understanding Pictures*. Oxford University Press.
Lutz, S. (2015). Partial Model Theory as Model Theory. *Ergo: An Open Access Journal of Philosophy*, 2. https://doi.org/10.3998/ergo.12405314.0002.022
Menon, T. (2024, January 23). The inferentialist guide to scientific realism [Preprint]. https://philsci-archive.pitt.edu/23011/
Nguyen, J., & Frigg, R. (2021). Mathematics is not the only language in the book of nature. *Synthese*, 198(24), 5941–5962. https://doi.org/10.1007/s11229-017-1526-5
Rickles, D. (2016). *The Philosophy of Physics.* Polity.
Shapiro, S. (1997). *Philosophy of Mathematics: Structure and Ontology*. Oxford University Press.
Suppe, F. (1977). The Search for Philosophical Understanding of Scientific Theories. In F. Suppe (Ed.), *The Structure of Scientific Theories* (2nd Ed.). University of Illinois Press.
van Fraassen, B. C. (1991). *Quantum Mechanics: An Empiricist* View. Oxford University Press. https://doi.org/10.1093/0198239807.001.0001
van Fraassen, B. C. (2008). *Scientific Representation: Paradoxes of Perspective*. Oxford University Press.
Wallace, D. (2019). Isolated Systems and their Symmetries, Part I: General Framework and Particle-Mechanics Examples [Preprint]. http://philsci-archive.pitt.edu/16623/
Wallace, D. (2022). Stating Structural Realism: Mathematics-First Approaches to Physics and Metaphysics. *Philosophical Perspectives*, 36(1), 345–378. https://doi.org/10.1111/phpe.12172