Stating structural realism: mathematics-first approaches to physics and metaphysics

David Wallace

December 22, 2021

Abstract

I respond to the frequent objection that structural realism fails to sharply state an alternative to the standard predicate-logic, object / property / relation, way of doing metaphysics. The approach I propose is based on what I call a 'math-first' approach to physical theories (close to the so-called 'semantic view of theories') where the content of a physical theory is to be understood primarily in terms of its mathematical structure and the representational relations it bears to physical systems, rather than as a collection of sentences that attempt to make true claims about those systems (a 'language-first' approach). I argue that adopting the math-first approach already amounts to a form of structural realism, and that the choice between epistemic and ontic versions of structural realism is then a choice between a language-first and math-first view of *metaphysics*; I then explore the status of objects (and properties and relations) in fundamental and non-fundamental physics for both versions of math-first structural realism.

1 Introduction

> [E]very theoretical physicist who is any good knows six or seven different theoretical representations for exactly the same physics.
>
> Richard Feynman
>
> [1]

Structural realism rests on a seductive idea: that the descriptive categories of the scientific realist and the analytic metaphysician draw distinctions too fine for the scientist. Theory change might involve a radical shift of ontology, from waves as disturbances in a mechanical aether to waves as self-subsistent field states or from heat as an invisible fluid to heat as disordered motion, but the equations — the *structure* — of the theory change more continuously or not at all, and

[1] (Feynman 1967, p.162)

so knowledge of structure can be robust against theory change. Considerations of ontology and ideology might suggest a multiplying of theories, where gravity can be understood as curvature or as universal force, and where fields can be understood as extended matter or as properties of spacetime, but the equations don't care, and so underdetermination does not threaten knowledge of structure.

In the *epistemic* form of structural realism (ESR), this is a limitation on what we can know: science only tells us the structural features of reality, so any non-structural features, if knowable at all, are not knowable through science. In *ontic* structural realism (OSR), it transforms into a radical metaphysical thesis: if all science needs is structure, and if our metaphysics should conform to our science, then our metaphysics should contain nothing but structure. Both forms of structural realism, then — but especially the ontic variety — appear to challenge, even to threaten, 'traditional' analytic metaphysics, which seems committed to more than the structure that science reveals. And some of the advocates of OSR (most notably Ladyman and Ross (2007)) have been harshly critical of analytic metaphysics, so that the development of OSR has blurred into a broader campaign against supposedly unscientific philosophy.

But throughout this vibrant literature it has been frustratingly difficult to pin down just what structural realism (in either form) actually says, and attempts to precisify it seem to lead away from the original seductive idea and back to supposedly-unscientific analytic philosophy: precise versions of ESR struggle to articulate a notion of structure intermediate between full-fat scientific realism on the one-hand and instrumentalism on the other; precise versions of OSR get into surprisingly metaphysical conversations about how we can have 'relations without relata', about how individuals can be posited without essential properties, whether structural realism requires an infinite tower of individuals at one level reimagined as bundles of relations at another, and the like.

One senses frustration on both sides. Metaphysicians who engage critically with OSR (e.g., Dorr (2010b), Hawley (2010), and Sider (2020, pp.63–65)) object that the project is insufficiently clearly stated (and contrast that inclarity unfavorably with the standards of rigor in modern metaphysics). Defenders of OSR, in turn (notably Steven French and James Ladyman) lament that the resources of set theory and logic are ill-suited for OSR:

> [T]he structuralist finds herself hamstrung by the descriptive inadequacies of modern logic and set theory which retians the classical framework of individual objects represented by variables ... In lieu of a more appropriate framework for structuralist metaphysics, one has to resort to ... treating the logical variables and constants are more placeholders which allow us to define and describe the relevant relations which bear all the ontological weight. (French and Ladyman 2003)

But it has not been clear — at least not to their critics, at least not to me — how this move really gets them out of the task of explicating 'relations without relata' (Ladyman's (2020) Stanford Encyclopedia discussion of OSR continues to center this puzzle). Similarly, Ladyman and Ross (2007) make extensive appeal to Dennett's (also seductive!) notion of 'real patterns' (Dennett 1991) as the basis for a structuralist ontology — but Dennett's own discussion intentionally foregrounds concrete examples and bypasses any systematic development of the underlying metaphysics, and in any case seems committed to patterns as higher-level structures in an antecedently-understood low-level ontology, which without modification is insufficient for OSR's needs. All in all, one gets the impression that the various discussants are talking past each other.

The purpose of this paper is to define a version of structural realism — or more accurately, two, an epistemic and an ontic version — that realizes the core goals that French, Ladyman, Ross *et al* express and that is at least explicable to their metaphysical critics. My focus is not so much on *defending* this version of structural realism, as it is on simply *stating* it with sufficient clarity that all parties can achieve at least rough agreement as to what the position in question actually is.

The starting point for my account is the question of how scientific theories should be formulated, and in particular the debate between so-called 'syntactic' and 'semantic' accounts of scientific theories (though I find this terminology misleading, and will not adopt it in the main part of the paper). Several prominent structural realists (e. g. Ladyman (1998), French and Saatsi (2006)) have stressed the significance of the semantic view for structural realism, but their appeals do not seem to have penetrated the general debate about the metaphysics of OSR, and (largely in conversation, though see (Sider 2020, p.193)) I have the impression that metaphysicians mostly regard it as irrelevant to their interests: interesting to the philosopher interested in the details of scientific practice no doubt, but neutral as regards broader and deeper philosophical themes. I will argue that this is not the case: the debate between structural and standard forms of realism is inextricably bound up with the debate between different ways of conceptualizing scientific theories, and indeed the move from a 'syntactic' to a 'semantic' conception of theories is itself more or less sufficient to turn standard realism into structural realism.

The structure of the paper is as follows: In section 2 I present what I call the 'language-first' and 'math-first' views of theories (my preferred terminology for, roughly, the syntactic and semantic views) and in sections 3–5 I explain why the distinction matters, through discussion of how the two views treat the theory/evidence relation, theoretical equivalence, and inter-theoretic reduction; these sections develop and argue for ideas that are extensively used in the rest of the paper. In sections 6–7 I review standard scientific realism, and structural realism, from a language-first perspective. In section 8 I consider how scientific realism changes if we move to the math-first view of theories, and argue that this move by itself gives us a form of structural realism, 'math-first structural realism'. In sections 9–10 I develop epistemic and ontic versions of this structural realism — the difference between the two is whether we adopt a language-first or math-first view of *metaphysics*. In sections 11–12 I consider how we should understand objects — at the fundamental level and in our higher-level physical theories — within math-first structural realism. Section 13 is the conclusion.

Three disclaimers before I begin. Firstly, this account is inspired by, and deeply indebted to, the overlapping versions of OSR developed by Steven French, James Ladyman, Michael Redhead, Don Ross, and Simon Saunders. However, I don't intend it as straightforward exegesis of their respective views, and indeed in several places it appears to contradict certain details of some of these authors' own views. Insofar as I am after all succeeding in expressing in different terms their ideas, great. Insofar as there are differences, the account here should be judged on its own terms. (And insofar as there are errors, they are on me!) Secondly, structural realism has been criticized for being appropriate to physics but ill-suited for the special sciences, especially the life sciences. Beyond some brief exploratory remarks in the conclusion, I will not engage with this issue: my examples are exclusively drawn from physics, partly precisely because of its more natural fit to structural realism, partly for parochial reasons of my own technical expertise. If the versions of structural realism I present here fail even internal to physics, their extendibility to the special sciences is moot. If they provide a satisfactory account of physics but fail to generalize beyond, so be it: we will still have learned something. Finally, while I have tried to prioritize clarity in this presentation, the views I describe are inevitably only first drafts: Many details remain to be filled in, and the devil may be in those details.

2 Two views of scientific theories

What is a scientific (specifically: a physical) theory? One view, tacit in much modern metaphysics and explicit in mid-century philosophy of science, is: *theories are collections of sentences.* As a clear example of a theory of physics presented this way, we need look no further than Newton's *Principia*:

> Every body perseveres in its state of being at rest or of moving uniformly straight forward except insofar as it is compelled to change its state by forces impressed . . . A change in motion is proportional to the motive force impressed and takes place along the straight line in which that force is impressed . . . the common center of gravity of two or more bodies does not change its state whether of motion or of rest as a result of the actions of the bodies upon one another.

Examples could be multiplied: it is at the least highly defensible that the theory presented by Newton in the *Principia* is a collection of sentences in this sense.

But a modern presentation of Newtonian mechanics would look quite different, more like this:

> A model of N-particle Newtonian mechanics is specified by:
>
> 1. A list of N positive real numbers $m_1, \ldots m_N$, representing the particle masses;
> 2. A list of $N(N-1)$ smooth potential functions $V_{nm} : \mathbb{R}^3 \times \mathbb{R}^3 \to \mathbb{R}$ representing the 2-particle potential between the pairs of particles and satisfying $V_{nm} = V_{mn}$;

3. A collection of N smooth functions $x_n : \mathbb{R} \to \mathbb{R}^3$ satisfying the differential equations

$$m_n \frac{\mathrm{d}^2 x_n(t)}{\mathrm{d}t^2} = -\sum_{m=1, m\neq n}^{N} \nabla V_{nm}(|x_n - x_m|) \tag{1}$$

This is a conception of theories not as collections of sentences, but as collections[2] of *mathematical models*. Of course, I used language to describe those models to you. (How else could we have communicated? I'm not telepathic.) But what I described was a mathematical system, not (directly) the physical system that this system is intended to represent.

In this paper I will refer to the first view of theories as the *language-first* conception of theories, and the second as the *math-first* conception of theories. These are not their traditional names: the philosophy of science literature refers to them, respectively, as the *syntactic* (or 'received') view of theories, and the *semantic* view of theories. These names arise from the concomitant tendency in that literature to regard theories as *formalized*: respectively, in some formal language and in set theory. The syntactic/semantic distinction then tracks the familiar distinction between the syntax of a formal language and the model-theoretic semantics of that same language. From this perspective, the distinction between the two views of theories can seem thin, even terminological, and the supposedly 'non-linguistic' nature of the semantic view can seem superficial: any recursively-enumerable set of sentences in first-order logic closed under logical consequence will determine a class of models, and conversely, from that class of models the original set of sentences can be recovered.

But the need for, and consequences of, formalization are controversial, especially among advocates of the 'semantic' view. Patrick Suppes, one of the early advocates of the semantic view, famously[3] stated that mathematics, not metamathematics, should be used in philosophy of science — but then he quickly went on to present his version of the view in a language of set-theoretic predicates not so far removed from metamathematics. van Fraassen, one of the most prominent recent defenders of the semantic view, more robustly states that a theory

> may be described in many ways, by means of different statements in different languages, and no linguistic formulation has any privileged status. Specifically, no importance attaches as such to axiomatization, and a theory may not even be axiomatizable in any non-trivial sense. (van Fraassen 1989, p.188)

And it is common to note (see, e. g., (Barrett and Halvorson 2016, p.570), (Weatherall 2019, [part 1 pp.6-7)) that theories in modern physics are not stated

[2]The original literature on this view of theories regarded these 'collections' as unstructured sets. More recent work (e. g. (Halvorson 2012), (Weatherall 2016); for a fuller discussion and references see (Weatherall 2019, section 4)) strongly suggests that they should be understood as categories or other more-structured collections. At the level of generality of this paper, the distinction is not significant.

[3]More precisely, Van Fraassen (1980, p.65) famously quoted Suppes as saying this.

in first-order logic and that it is questionable at best if they *could* so be formulated without loss. The recent literature, however, has seen a resurgence of formal methods and a skepticism that any non-linguistic conception of theories is actually viable once we consider how it could be formally given (e. g. (Halvorson 2012; Glymour 2013; Barrett and Halvorson 2016); see again (Weatherall 2019) for a general review of this debate).

In this paper I largely eschew formalized approaches to theories, mostly through a desire to stick as closely as possible to physics practice and to theories as physicists present them. In the absence of precise definitions and the precise results they permit, I proceed ostensively, through the use of a number of concrete examples which I hope illustrate the general contours of the math-first and (to a lesser degree) the language-first approach in contemporary physics and contemporary philosophy of physics. This is to deny neither the value and fruitfulness of the more formal approach, nor the possibility of a development of the formal approach which naturally captures modern physics; however, *right now* engaging with the broader issues in philosophy of science and metaphysics of science needs an account of theories that directly applies to the theories we actually wish to discuss, and formalization does not permit this. In any case, my non-standard terminology ('math-first', 'language-first') is intended both to make clear that I do not have a formalized conception of theory in mind and to explicitly cancel any implications that might arise from talk of 'syntax' and 'semantics'. (It is also intended to allow for hybrid possibilities, in which (for instance) there are irreducibly linguistic elements to a theory nonetheless predominantly given through mathematics.)

Granting for the moment that the distinction is coherent and non-trivial, what if anything determines which view is *correct*? This cannot be an entirely descriptive matter: if we tried to determine what a theory is using just the tools of the historian or sociologist, it is hard to imagine obtaining anything other than a messy object containing a mixture of verbal content, mathematics, and pieces of tacit or explicit practice: something like Kuhn's (1970) 'disciplinary matrix'. (Sometimes just such a view is called a 'pragmatic view of theories' and constrasted with both views I dicuss here — see (Winther 2021) and references therein.) I have in mind something more like a rational reconstruction: if we tidy up the mess and missteps in (say) the development of Newtonian mechanics, or in presentations of quantum field theory, what is the actual epistemic achievement those developments granted to us: something like a set of linguistic claims, something like a collection of mathematical models, or something in between? And in answering this one has to acknowledge that physics has never been given fully through the presentation of mathematical models, and yet that for at least the last century it has never avoided some explicit presentation of such models, so that the answer requires judgement and close analysis and cannot simply be read off the textbooks.

That said, I think there is a pretty strong case that, while the language-first view may well be correct for the theories of Newton and his contemporaries, the physics of the late nineteenth century, and even more so of the twentieth and twenty-first, is much more readily interpreted via the math-first view. But

I will not directly defend that view here (I defend it indirectly through the various examples of physical practice I give in sections 3–5) and, in any case, it would be open to an advocate of the language-first approach to argue that this change marks physics as losing its way and descending into obscurity, so that extracting coherent content from physics requires substantial reconstruction along language-first lines. For instance, (Allori *et al* (2008) and Maudlin (2018) argue that any clear presentation of a physical theory must consist of a specification (i) of its ontology, explicable in non-dynamical terms, and (ii) only subsequently of its dynamics; the substantial reconstruction of extant physics required to realize this takes us some way back towards the language-first approach.

So for the most part in this paper I am less concerned with arguing why a math-first conception of theories might be *correct* as a description of contemporary practice, as I am in saying why the question is *important* for philosophy of science and for metaphysics. The answer arises from three key differences between the two conceptions: for formal notions of equivalence between theories, for inter-theoretic reduction, and for the theory/world relation. In the next three sections I explore each in turn; doing so will also serve to illustrate what I see as the core features of the language-first and math-first approaches, features which I think can be identified from scientific practice even ahead of a precise formal definition of either approach.

3 Why it matters: empirical predictions

On the language-first view of theories, the normal assumption is that a theory makes contact with empirical data by saying true things about it. At least for those statements in the theory that concern observable matters, to accept the theory is to take them literally and believe that they are true. On a language-first construal of Kepler's laws, 'Bodies in the Solar system move in ellipses with the Sun at one focus' really does assert that bodies in the Solar system move in ellipses with the Sun at one focus. The relations that a good theory's empirical statements have to the facts are those familiar from ordinary-language semantics: truth, reference, satisfaction.

On the math-first view of theories, a theory makes contact with empirical data by modelling them. A math-first construal of Kepler's laws would take them to provide certain models of planetary motion; Kepler's actual empirical data can likewise be represented mathematically (in the semantic-view literature, this is called (Suppes 1962) a 'data model'); empirical adequacy consists of a partial isomorphism of the data model into the theoretical model, probably restricted further by some interpretative constraints, so that not any old interpretation will do. The theory/world relation here is *representation*, more akin to the relation between map and territory than that between word and object.

(On neither view need "observable matters" be confined to naked-eye observations. The logical empiricists held out hope for an 'observation language' interpretable without any mediation by theory, but that hope was vain, and the predictions of modern theories of physics are usually made in terms of other physics theories. The prediction of the Higgs model that was confirmed at the LHC in Geneva was that there would be a local resonance in the proton-proton cross-section with certain characteristics; the prediction of gravity waves confirmed by LIGO was that spacetime would display a tiny pattern of expansions and contractions; the prediction of the lowest-energy gap in hydrogen made through non-relativistic quantum mechanics is that hydrogen has a spectral line at 121.57 nm. In no case is the prediction testable with the naked eye; in each case, confirming the prediction requires further input from a theory other than that being tested.[4])

The distinction between reference and representation will recur when I consider the status of scientific realism on the two views, but here I'll note two ways in whch the distinction matters even when restricted to empirical confirmation.

1. **Approximation and approximate truth.** Bodies in the Solar system do not, in fact, move in ellipses with the sun at one focus. At least for the planets of the Solar system that's a good approximation in most circumstances, but a more accurate description in Newtonian gravity would allow for the finite mass ratio between Sun and planet, and for the gravitational tugs of other planets. Even that description is not strictly correct: general relativity makes further corrections. So empirical success cannot simply be a matter of truth or falsity: Kepler's laws are false for the planets, but they are still in some sense more accurate than the claim that the planets move in spirals or rectangles, and Newtonian gravity is in the same sense more accurate than Kepler's laws, and general relativity is in the same sense more accurate than Newtonian gravity.

   But in *what* sense? In the language-first approach to theories one often hears reference to "approximate truth": Kepler's laws are false, it is said, but they are approximately true; the predictions of Newtonian gravity are also false but are an even closer approximation to being true; the predictions of general relativity are closer yet. But if "approximate truth" is some semantic notion, akin to truth *simpliciter*, it has proved frustratingly difficult to formally define (see, e. g. , Laudan 1981, section 4; Newton-Smith 1981, Ch.VIII). In practice, "it is approximately true that planets move in ellipses" just seems elliptical for "it is true that planets approximately move in ellipses", which in turn is equivalent by disquotation to "planets approximately move in ellipses". *That* notion has a number of perfectly reasonable definitions (we could, for instance, take the measure of approximation to be the distance between the predicted and actual location of the planet, time-averaged over one orbit) and those definitions have the further benefit of matching what physicists actually

[4]If you have sufficient faith in the scope of the Standard Model of Particle Physics, then *in principle* the machinery of the LHC can be described by the same theory as the Higgs boson, but that certainly isn't how the LHC works in practice. In the other two cases, the distinction is one of principle: gravity-wave astronomy uses quantum-mechanical techniques to test the predictions of classical general relativity; non-relativistic models of atomic structure know nothing of light but are tested via the frequencies of emitted photons.

do when checking the validity of approximations and predictions. But this is now really a statement in the math-first presentation of the theory, about the accuracy of fit between a model of the theory and a data model; the notion that the original statement should be true has disappeared.

To put the point another way: the notion of approximation used in physics seems too mathematically rich, too dependent on the full mathematical resources of particular theories (in the Keplerian case, spatial distances, but that will not generalise to all cases) to be captured using simply the resources of general semantics, in a theory-independent way. Truth is too brittle: without theory-specific resources, "approximately true" just seems to be polite way to say "false".

2. **Domain restriction.** Bodies in the Solar system don't always even *approximately* move in ellipses. A sufficiently close encounter between two bodies orbiting the sun can cause one or both to wildly deviate from elliptical movement: that's how gravitational slingshots are used to accelerate space probes, and how the chaotically-orbiting bodies of the early Solar system eventually sorted themselves out into their current relatively staid state. So then in what sense are Kepler's laws even *approximately* true? Is the claim that *most* bodies in the Solar system obey them? If that works, it works only accidentally: in the case of Newton's laws, it's probably not even true that *most* bodies persevere in their state of being at rest or of moving uniformly straight forward except insofar as they are compelled to change their state by forces impressed: assuming particles count as 'bodies', most bodies are quantum particles that don't strictly obey anything like Newton's laws.

   In any case (at least as applied to scientific practice) this misses the point. The value of Kepler's laws, or Newton's, isn't that some modification or weakening of them holds *universally*: it's that to a good approximation, they hold *locally*, for certain systems. Kepler's laws apply pretty accurately to the centers of masses of solar-system bodies which don't approach other bodies too closely; Newton's laws apply pretty accurately to reasonably large bodies where quantum effects can be neglected. Again, though, this shows that *truth*, even approximate truth, is not actually what we get for the empirical predictions of successful physics theories: rather (on the language-first view) what we get is empirical statements that are true *of* certain systems, and/or that are true under certain interpretations.

   This restriction, somewhat awkward on the language-first view, is fairly natural and automatic on the math-first view. There, domain-restrictedness is automatic: no-one supposes that scientific theories in general are supposed to model *all data anywhere*. It is built in to a conception of a physics theory as a structured collection of models that some models describe some systems and others describe other systems: Newtonian gravity has a 2-body sector that describes the Earth-Moon system, another 2-body sector (differing formally only by assignment of masses) that describes the Earth-

> Sun system, a 9-body sector that describes the sun and planets, a 5-body sector that describes Jupiter and its major moons, and so forth; in each case, the 'particles' represented are the centers of masses of the planets. (Someone who looked at a statement of 'N-particle Newtonian mechanics' like that in section (2) and asked 'what is N?' would be misunderstanding what is being presented.)

Both of these examples show that the naive conception of how the language-first view makes contact with empirical data faces serious obstacles and needs substantial revision if it is to do justice to scientific practice, whereas the math-first view is a more natural fit; against that, it's certainly true that we have a fuller and more carefully developed philosophical analysis of the semantic notions appealed to in the language-first view than we do of math-first-style scientific representation, although that apparent increase in clarity may rely on disregarding the approximate-truth and domain-restriction problems.

4 Why it matters: formal theoretical equivalence

Physics practice is replete with the idea that two apparently-different theories are 'equivalent'; what that means is contested, but at the least it comes with the commitment that there is no non-pragmatic scientific question as to which theory to adopt. If theories X and Y are equivalent, discussions of which is more useful in a given problem are common, but discussions of which is *correct* are not to be found. Indeed, physicists make no real distinction between 'these are two equivalent theories' and 'these are two formulations of the same theory' (cf the quote from Feynman which begins this paper).

It is implausible that theoretical equivalence is in all circumstances an entirely formal matter (on either the language-first or math-first conception of theories). Theories normally model only parts of the world and parts of those theories may come with interpretations based on other parts of the world: the term 'electron' in a theory of the structure of beryllium, say, is substantially interpreted via its use in other theories (Sklar 1982; Sider 2020, pp.178-182). And formally identical theories may be used to model radically different physical systems: the Langevin equation of statistical mechanics can model fluctuations in concrete physical systems (the position of a pollen grain), in more abstracted features of a physical system (the total energy), or a system outside the normal domain of physics entirely (the value of a stock on the stock-market).[5] At the least, equivalence requires empirical equivalence, so that those parts of the theory which make contact with empirical evidence, including empirical evidence mediated by other theories, will need to be invariant between theoretically-equivalent theories, and that requirement is unlikely to be purely formal.

[5] I have borrowed the example, with minor modifications, from (Weatherall 2019, part 1 p.8).

Nevertheless, a formal notion of equivalence seems to play at least some large role in the use of equivalence in physics. And the appropriate formal notion differs sharply depending on whether we adopt a language-first or math-first view of theories. On the language-first view, formal equivalence seems to be something like intertranslatability, or logical equivalence, or interdefinability. Equivalent theories are talking about *the same entities*, and saying *the same things* about them, just using different words or expressions. The English and Latin versions of Newton's *Principia* are not rival theories; Feynman and Gell-Mann should not be understood as having proposed rival accounts of the supposed constituents of the proton just because one called those constituents 'partons' and the other called them 'quarks'. In the context of *formal* languages, this notion of equivalence shows up in the existence of different sets of axioms that generate the same set of sentences, and in a range of formally-defined notions of theory equivalence (Glymour 2013; Barrett and Halvorson 2016), but there is no unequivocal agreed-upon answer even in that context, let alone for theories expressed in natural language.

On the math-first view, theoretical equivalence is something like equivalence by the standards of mathematics: a 1:1 transformation between models that preserves mathematical structure. Pinning that down precisely is no easier here than in the language-first context (set-theoretic isomorphism is too restrictive; categorical equivalence appears to be too permissive (Barrett and Halvorson, *ibid.*)). But it is relatively clear case-by-case, and a systematic feature of those cases is that theoretical equivalence is normally *much* more coarse-grained on the math-first than on the language-first view. I will illustrate with several examples.[6]

Euclidean space. Ordinary three-dimensional space plays a central role in many (mostly pre-relativistic) physics theories, but it can be defined mathematically in many equivalent ways: as a set coordinatized by a family of bijections into $\mathbb{R}^3$ (see Wallace 2019b), as a 3-dimensional affine space equipped with an inner product on its associated vector space, or as a 3-dimensional manifold diffeomorphic to $\mathbb{R}^3$ and equipped with a flat non-degenerate Riemannian metric, to name just three. (And the various mathematical objects named in those definitions are themselves multiply definable: there is more than one equivalent way to define a manifold or a vector space.) The math-first view regards these as equally-legitimate ways of presenting the same theory, but any plausible attempt to throw the different descriptions into language-first form (say, by describing each in first-order logical language) will realistically fail to provide any purely-formal translation between those descriptions. For instance, the affine-

[6]Some of these examples might be incorporated into the language-first view by some version of Morita equivalence (Barrett and Halvorson 2016) generalized beyond first-order logic; but not all, and in any case (to anticipate section 8) Morita equivalence seems already to move some way from the truth/reference/satisfaction model of the theory/world relation and towards the reprsentational model, so that a scientific realism based on Morita equivalence would look a lot more like math-first realism than a realism built on a more traditional notion of language-first theoretical equivalence.

space presentation of Euclidean space naturally requires one sort of variables ranging across points of space, another sort ranging across vectors, and a collection of functions relating points, vectors and real numbers; the coordinatized presentation requires one sort of variables ranging across points of space, another sort ranging across functions from those points to $\mathbb{R}^3$ (presumably handled through higher-order logic, through some fragment of set theory, or through a bespoke 3-place predicate 'maps to', where maps-to(x, y, f) is intuitively to be interpreted as $f(x) = y$). It's difficult to imagine any formal translation scheme that will relate the two.

(It would miss the point to argue that in this case we need some intrinsic characterization of Euclidean space underlying all of these accounts (a Tarskian axiomatization, say). The norms of math and physics regard the coexistence of different accounts of Euclidean space as innocuous; hence, on the math-first view of theories the 'we' who need that intrinsic characterization are not those trying to construct theories of physics. The need for that axiomatization would come, if at all, from some further *philosophical* project — of which more later.)

Non-relativistic quantum mechanics in the position representation. One way of stating non-relativistic quantum mechanics, much discussed in the metaphysics-of-physics literature, represents quantum states as complex functions on $3N$-dimensional configuration space, i. e. . the Cartesian product of N copies of Euclidean space. Even assuming we have an unequivocal mathematical presentation of the latter (written, say, as $\mathcal{E}_3$, or more generally $\mathcal{E}_n$ for n-dimensional Euclidean space), we could still describe the wavefunction as:

1. An N-place complex function on $\mathcal{E}_3$, assigning complex numbers to N-tuples of points in space (perhaps interpretable as describing a family of irreducibly nonlocal relations in which points of space can stand (Belot 2012, p.72; Wallace 2021, p.70);
2. A single-place complex function on configuration space (the N-fold Cartesian product $\mathcal{E}_3 \times \cdots \times \mathcal{E}_3$), assigning complex numbers to points in a $3N$-dimensional space (perhaps interpretable as a high-dimensional fundamental space (Albert 1996; Ney and Albert 2013);
3. A single-place function from ordered N-tuples of points to complex numbers, perhaps interpretable as a function on an *abstract*, set-theoretically-constructed space (Sider 2021).

These will look quite different expressed in the language-first view: the first will describe wavefunctions via N-place functions, the second and third via 2-place functions, but the third also requires some set-theoretic vocabulary. Mathematical practice barely distinguishes them: if (specializing to $N = 2$ for simplicity) I write $\psi : \mathcal{E}_3 \times \mathcal{E}_3 \to \mathbb{C}$, the notation doesn't tell me whether ψ is a 2-argument function both of whose arguments are in $\mathcal{E}_3$ or a 1-argument function whose argument is in $\mathcal{E}_3 \times \mathcal{E}_3$. And it further doesn't tell me whether $\mathcal{E}_3 \times \mathcal{E}_3$ should be interpreted simply as the set-theoretic Cartesian product of 2 copies of $\mathcal{E}_3 \times \mathcal{E}_3$ (in which case the distinction between the second and third possibilities is mathematically invisible) or as one of the other ways of constructing products of affine spaces (or whatever sort of space we are using for $\mathcal{E}_3$).

Of course, there are a great many other ways to represent quantum mechanics, and physicists treat those all as equivalent too; but even restricting attention to the position representation demonstrates how much coarser-grained theoretical equivalence is on the math-first view.

Newtonian particle mechanics. A model of Newtonian N-particle mechanics is given by N smooth trajectories in Euclidean space. But that statement could be precisified as (inter alia):

- N smooth (that is: infinitely-many-times differentiable) maps from the real line[7] to $\mathcal{E}_3$, satisfying such-and-such differential equation.
- N smooth curves (that is: dimension-1 submanifolds), in $\mathcal{E}_4 = \mathcal{E}_3 \times \mathcal{E}$, representing Newtonian spacetime.

The former might naturally be translated into the language-first view via some function $\mathrm{Loc}(n, t)$, giving the location of particle n at time t; the latter by some 2-place predicate $\mathrm{Occupied}(x)$ that records the points of spacetime occupied by particles. Again, the prospects of intertranslatability look dim.

Classical scalar field theory. A classical scalar field can be mathematically represented as either

- A function from Minkowski spacetime $\mathcal{M}$ to the real numbers (eliding questions, analogous to those I raised earlier about Euclidean space, as to how we represent spacetime);
- A pair of functions from a 4-dimensional differentiable manifold to, respectively, $\mathcal{M}$ and the real numbers. (Isham and Kuchar 1985; Varadarajan 2007; Wallace 2015).

These might most naturally described in language as, respectively, (1) an ontology of spacetime points and a function $\mathrm{Value}(x)$ from spacetime points to real numbers (so that there is no such *thing* as a scalar field; (2) a dual ontology of spacetime points and field parts, and two functions Value and Location, from non-extended parts of the field to, respectively, real numbers and spacetime points, encoding field strength and spacetime location. Fairly clearly there will be no straightforward translation

[7] Well, from Euclidean 1-space, really.

between the two; the two *mathematical* descriptions, however, are standardly treated as equivalent.[8]

The AdS/CFT correspondence. My last example is dramatically more complicated than the others and can only be sketched here. AdS/CFT correspondence[9] is the conjectured equivalence between any quantum theory of gravity on asymptotically N-dimensional anti-de Sitter (AdS) spacetime, and a conformal quantum field theory (CFT) on the $N-1$-dimensional conformal boundary of that spacetime. (The form of the equivalence is an isomorphism between the respective Hilbert spaces of the quantum field theories, preserving spacetime symmetry groups and with certain mappings explicitly stated; the equivalence provides a systematic method to restate questions asked of the boundary formalism in the interior formalism, and vice versa.) The natural descriptions in language of the two sides of this equivalence look as different as can be — they disagree about spacetime's geometry, whether it is dynamical, and even its dimension. But I have yet to see any suggestion in the huge physics literature on AdS/CFT that there is a substantive further question as to whether the AdS or CFT description is the *correct* description, once technical questions about establishing mathematical equivalence are set aside.[10]

A common feature of these examples is that mathematical equivalences between theories mix up ontological categories. Some terminology is helpful: 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. Then given predicate precisifications of mathematically-equivalent theories, there is often no simple relation of the objects in the first to the objects of the second; likewise the properties and relations. Mathematically-equivalent theories cannot just be construed as talking about the same entities or ascribing the same properties to them.

[8] I once submitted a paper (Wallace 2015) developing a close cousin of the second presentation — adapted for general relativity — to *Physical Review*; it was desk-rejected on the grounds that *Physical Review* does not publish reformulations of existing theories.

[9] The original references are (Maldacena 1998; Witten 1998), both of whom considered a specific 5D/4D duality; (Kaplan 2016) is a good technical presentation of the more general equivalence I consider here; see (Wallace 2018) for a (comparatively) non-technical discussion and for further routes into the (huge) literature.

[10] AdS/CFT correspondence has not been sharply stated, let alone rigorously proven, essentially because while the CFT side of the duality is fairly well understood mathematically, physicists have good mathematical characterizations only of certain approximation regimes for the AdS side, and lack a sufficiently precise way to state the AdS version of the theory in full generality — *except* as "dual through AdS/CFT correspondence to such-and-such conformal field theory". It remains a live possibility that there is *no* other way to state it, in which case the interior description would be better regarded as emergent from the boundary description rather than equivalent to it. But this turns on technical and mathematical issues; there is no suggestion in the physics literature of a substantive question about the direction of emergence even if a mathematical duality can be established.

5 Why it matters: Inter-theoretic reduction

Philosophical discussions of reduction often relate physics to the special sciences, but inter-theoretic reduction plays a crucial role internal to physics too. The reduction of thermodynamics to statistical mechanics is the classic example, but others (some more controversial than others) abound: the reduction of Newtonian gravitation to the low-velocity, low-curvature regime of general relativity; the derivation of dilute-gas dynamics from the collisional physics of particles; the derivation of the equations of fluid dynamics from molecular physics.

Accounts of inter-theoretic reduction differ between the language-first and math-first views of theories in quite similar ways to their respective accounts of theoretical equivalence: indeed, in some sense on either view reduction is just an asymmetric version of equivalence.

On the language-first view, reduction is something like derivability or definitional extension: concrete claims in the higher-level theory can be translated into claims in the lower-level theory, and the laws of the higher-level theory can be derived from those of the lower-level theory, in each case via some translation principles or "bridge laws" (see (Butterfield 2011a, 2011b; Dizadji-Bahmani, Frigg, and Hartmann 2010) and references therein for an up-to-date account). To use the most famous example (Nagel 1961, ch.11): the ideal-gas law is (supposedly) derivable from the statistical mechanics of a dilute gas via the bridge law that temperature is mean kinetic energy.

The details can get messy, though. A more thoroughgoing reduction of dilute-gas physics to statistical mechanics will probably want to refer to the *gas*, an extended body which according to dilute-gas physics has spatially-varying temperature, pressure, density, and velocity. But a language-first version of statistical mechanics presumably quantifies only over particles and, perhaps, the points of the spacetime in which they move; none of these objects can be identified with the gas or its parts. One common answer, going back to Oppenheim and Putnam's classic (1958) case for reductionism and widespread in modern metaphysics, is to enrich our ontology with composites or *mereological sums* of particles: the gas is then the mereological sum of all the atoms. (A related alternative (Butterfield 2011a) is to use set-theoretic constructions to build the gas from the particles.) This strategy still leads to awkwardness regarding the *properties* of the gas, of course: according to fluid dynamics, gases are *continuous*, with a smoothly varying density, and spread to occupy all the volume in their containers, whereas the mereological sum of the gas particles (ignoring quantum subtleties) is highly discontinous and occupies a small, fixed fraction of the volume of the gas. An instrumentalist move is tempting at this point (see, e. g. , (Sklar 2003)): the gas *appears* continuous on sufficiently large scales but is not *really* continuous. (The instrumentalist move is also available at an earlier stage, as an alternative to mereology: strictly speaking there is no gas, but it's a useful fiction to pretend there is a gas when the molecular distribution has such-and-such features.)

It is controversial whether mereology is adequate for the reduction of continuum mechanics to particle mechanics, but there are far more problematic cases. Hadronic physics (photon/neutron physics) is generally said to reduce to quantum chromodynamics (the theory of quarks): in popular accounts one learns that a proton or neutron is *composed of* three quarks, and it is natural to try reading that statement mereologically. But if one looks at the *technical* literature on chromodynamics (see, e. g. , (Campbell, Huston, and Krauss 2018)) it becomes quickly apparent that the relation between the quarks and the proton is an awkward fit at best to mereology. And in some cases mereology does not even superficially seem to fit the case. The quantum theory of vibrations in crystalline solids is normally described in terms of *phonons* — quantized particles of vibration. (For instance, the heat capacity of (insulating) solids at low temperature is calculated by treating the solid as a gas of phonons.) But there is no way to identify a quantized vibration with any atom or collection of atoms in the solid — the best low-level translation of "there are such-and-such phonons present" would be "the solid is vibrating in such-and-such a way".

The common theme here is that on the language-first view, derivability of a higher-level theory from a lower-level one, if possible at all, involves quite a holistic form of translation. The higher-level objects, properties, and relations cannot be reidentified in any piecemeal way with the objects, properties, and relations of the lower-level theory.

On the math-first view, reduction is something like *instantiation*: the realizing by some substructure of the low-level theory's models of the structure of the higher-level theory's models. In the important case of *state-space instantiation*, for instance (discussed in more detail in (Wallace 2012, ch.2)), the lower-level theory instantiates the higher-level one if (roughly) there is a map from the lower-level state space to the higher-level state space that commutes with the dynamics and leaves invariant any commonly-interpreted structures (for instance, spacetime structure) in the two theories.

As in the case of theoretical equivalence, this is a much more permissive notion of reduction than in the language-first case, and permits reduction relations which are not straightforward to analyze in that case. For instance, the phonon description of a solid can be simply seen[11] as instantiated by the atomic-level description just by showing a (spatial-symmetry-preserving) Hilbert-space isomorphism between the respective quantum theories' Hilbert spaces under an appropriate state restriction — to states of the solid not too energetic to break apart, roughly. (For further philosophical discussion of this case, see (Franklin and Knox 2018).) And one standard[12] route to the instantiation of fluid dynamics in particle statistical mechanics starts with the density distribution, analyzes it into Fourier components, and constructs an autonomous dynamics for the lowest few Fourier modes, describing the variations on large scales. The common theme here is that the higher-level degrees of freedom are identified with a (de facto) autonomous subset of the lower-level degrees of freedom — but 'degrees

[11]See (Ashcroft and Mermin 1976), or any standard graduate text on solid-state physics; to be sure, there are subtleties here, and a more sophisticated treatment — e. g. the discussion in (Anderson 1984) — would also refer to the spontaneous breaking of translation symmetry by the crystal lattice.

[12]See, e. g. , (Balescu 1997, ch.10).

of freedom' do not simply or systematically correspond to any particular ontological category, and so math-first reductions do not reliably lead to a reduction stateable in those terms. Nor is there any very obvious analog, on the math-first treatment, with the need to posit mereology to facilitate reduction.

One further relevant difference between the two approaches is their different treatment of the approximate and domain-relative nature of reduction. Generally speaking, reduction relations in physics (and beyond) hold only approximately and only under certain assumptions about the state of the lower-level system: the phonon description of vibrations only works when atoms are assembled into a solid; the fluid-dynamics description of an atomic gas works only given low densities and appropriately uncorrelated velocities; quarks only form protons in the low-energy regime; general relativity is only well approximated by Newtonian physics when curvatures and velocities are low. This is closely related to section 3's discussion of approximate truth (after all, normally superseded physics theories are retained as higher-level approximations to their successors) and the same issues basically recur: approximation is fairly unproblematic on the math-first approach because there are rich theory-specific metrics for the level of approximation being used (metrics which are widely used in practical physics when discussing when a given approximate description is appropriate) but is hard to handle on the language-first approach given the difficulty of making clear sense of 'approximate truth'. (Schaffner 1967), working in the language-first framework, proposed that reduction should be understood as a relation between a *corrected* higher-level theory and a lower-level theory, where the corrections make the deduction exact, and other proponents of language-first reduction (e. g. Dizadji-Bahmani, Frigg, and Hartmann (2010)) endorse this move, but very often in physics there *is* no way of stating the 'corrected' theory except via the lower-level theory — this is the familiar issue of *ceteris paribus* clauses in the special sciences, which normally cannot be stated in the language of those sciences (Fodor 1989).

6 Language-first scientific realism

I hope that the previous three sections have somewhat elucidated the distinction between the math-first and language-first views of theories and why that distinction might be relevant to our analysis of scientific practice, as well as giving some evidence for my earlier claim that the math-first view is a better fit at least to the contemporary practice of physics. In this section I turn to more metaphysical matters, beginning with a review of the scientific realism debate, and its implications for metaphysics, on the language-first view (which, I want to claim, plays an important background role in most of the literature on this debate). For the rest of this section, that view will be presupposed. (Here I summarise a well-established literature; see, e. g. , chapter 2 of (Ladyman and Ross 2007) and references therein for details and sources.)

The crudest statement of scientific realism is that our current best scientific theories (or, in my restricted context, our current best *physics* theories) are true, not just in their claims about observable matters but in all their claims. Almost no-one thinks anything quite that crude, given the hubris of supposing present-day physics to be the last word; a more sophisticated statement is that our current best physics theories are *approximately true* in some (underspecified) sense. (Sometimes one adds a clause to the effect that successive theories approximate the truth more closely.) There are two main arguments given for scientific realism (albeit the first is often tacit):

The negative argument: In light of the theory-ladenness of observation, there is no coherent way of making sense of scientific theories other than realism or outright Humean skepticism. (For instance, the logical-positivist and logical-empiricist alternatives are taken to rely on an illicit theory/observation distinction, and constructive empiricism is often criticized for arbitrariness.)

The positive argument: The empirical success of a successful scientific theory — or, more specifically, the success of that theory at making *novel, confirmed predictions* — has no good explanation other than its (approximate) truth; hence, the success of the theory is strong evidence of its truth. Or, applying the same argument form to science as a whole: scientific realism is the only philosophy of science that does not make science's successes miraculous (Putnam 1975). (For instance, quantum electrodynamics predicts the magnetic moment of the electron to one part in 10^{12} (Odom *et al* 2006); according to the positive argument, it is highly implausible that the theory should achieve this accuracy without being approximately true.) The form of inference relied on here is normally described as *inference to the best explanation.*

I set aside the negative argument (except to note that it relies on scientific realism itself being coherent, something which might be undermined by the difficulties with defining approximate truth) and a variety of criticisms of the inferential reasoning in the positive argument to consider the two most commonly stated objections to scientific realism: the argument from underdetermination of theory by evidence (that is, from the existence of alternative theories incompatible with, but empirically-equivalent to, our best theories), and the argument from theory change (that is, from the fact that our current theories are the successors of highly successful theories some of whose central theoretical claims are false according to our current theories). Both can be understood as attempted *reductio*s of the positive argument: a supposedly reliable inference from 'this theory makes multiple novel confirmed predictions' to 'this theory is approximately true' cannot after all be reliable if two mutually contradictory theories X and Y made those same novel confirmed predictions, whether X and Y are currently-live alternatives or X is a past theory and Y a present one.

Much of the debate on underdetermination turns on whether there really are realistic examples of underdetermination in science. Insofar as there *are* realistic examples, the standard realist response is to appeal to so-called 'extra-empirical virtues': simplicity, explanatory power, fruitfulness and the like. A commonly discussed[13] example, for instance, is the supposed underdetermination between curvature-based theories of gravity and theories where gravity is a universal force (and also responsible for stretching rods and slowing clocks in a uniform way): it is argued that considerations of simplicity and explanatory power resolve the underdetermination in favor of the curvature-based account.

As for theory change, here there is widespread consensus that there really are examples of past theories which made novel confirmed predictions but which appear to contradict present theories: commonly discussed examples are the caloric theory of heat (in which heat is a fluid rather than a form of energy), the aether theory of light (in which light is a moving disturbance in a material medium, the aether, rather than a feature of the electromagnetic field), and Newtonian gravity (in which gravity is due to instantaneous action at a distance between masses, rather than a manifestation of the curvature of spacetime). By any reasonable account, all three were highly successful theories and yet do not appear compatible with our current best theories. Realist responses to this[14] mostly come down to (i) arguments that there is more continuity than meets the eye (so that, for instance, 19th-century physicists were unknowingly talking about the quantum electromagnetic field when they spoke of the aether) and (ii) selective skepticism, where the scientific realist should only believe in the *explanatorily central* features of theories (so that, for instance, 19th-century scientific realists would not have been justified in believing in a material aether because it was not explanatorily central in the contemporary theory of light). Both responses have been criticized as *ad hoc* and as available only after the fact.

(Language-first) scientific realism fits very naturally into the view that the methodology of metaphysics is continuous with that of science (usually physics), advocated famously by Quine (e. g. (1957)) and, more recently, by metaphysicians like Dorr (2010a), Paul (2012) and Sider (2011, 2020). A partial list of the continuities:

1. Modern metaphysics often refers to the 'book of the world' (Sider 2011), the supposed theory of reality at the most fundamental level. The book, roughly, gives our completed, idealized physics; the statements of our currently-most-fundamental physics are then a tentative guide to the form of the ultimate fundamental theory.

2. In parallel with the negative argument for scientific realism, modern metaphysics generally avoids Carnap-style deflationism about the content of its theories and takes itself to be making statements whose semantics are continuous with those of (an ideal completed) physics (notwithstanding more deflationist recent arguments in metametaphysics from, e. g. , Hirsch (2011)).

3. The methodology of metaphysics (on this view) is explanationist in the sense of the positive argument: metaphysical claims, when they transcend

[13] See, e. g. , (Reichenbach 1958; Sklar 1982; Earman 1993; Norton 2008).
[14] See Psillos (1999, chs.5-6) and references therein for details.

observation and evidence, are argued for on grounds of their explanatory power.

4. In particular, metaphysical underdetermination is to be adjudicated using the same tools as scientific underdetermination: theories are to be weighed up against one another on grounds of their simplicity, explanatory power, etc.

5. Scientific claims about the unobservable are tentative in a similar way to (though to a quantitatively lesser degree than) metaphysical claims.

6. There is a sharp divide between fundamental and higher-level ontology: the former is what there really, 'strictly' is; the latter, insofar as it cannot be identified with mereological sums of the former, is at best derivative, at worst an instrumentalist fiction.

7 Language-first structural realism

John Worrall's classic (1989) paper on structural realism (or *epistemic* structural realism —ESR — as it would now be called) presents it as a way to keep the explanatory force of realism whilst avoiding the pessimistic argument from theory change: his main example is the continuity in optics and electrodynamics not at the level of ontology, but at the level of the equations. (Maxwell's equations have a recognizably similar form before and after the abandonment of the aether, and indeed that form arguably persists in quantum electrodynamics.) Translated into the framework of scientific realism I sketched above, structural realism becomes a thesis of selective skepticism: realists should believe only the *structural* claims theories make about the unobservable, and not the residual (non-structural, non-observable) claims made by theories. (One part of the appeal of structural realism of this kind is that 'structure' sounds a less *ad hoc* form of selective skepticism than others that have been offered in response to theory change.)

Making that selective skepticism precise has proven difficult. Psillos (2004, 1999, ch.7) has argued that a sensible scientific realism is already 'structural' (scientific properties and relations are understood in terms of their structural role in a theory, not via their intrinsic nature) and so a restriction to structure is no restriction at all. And the oft-discussed strategy of 'Ramsifying' a theory by replacing named theoretical terms with existentially-quantified variables seems to collapse the content of a theory to its observational claims (the so-called 'Newman objection').[15] Ultimately the difficulty is that 'structure' needs to go beyond just 'observable' and yet fall short of including the full content of a theory, and it is not clear that there is really a stable half-way house here. (Though it has been less discussed, buiding a distinction between theoretical and

[15]Originally presented in (Newman 1928), rediscovered by Demopoulos and Friedman (1985), and deployed against Worrall's structural realism by Ketland (2004); see (Ladyman 2020) and references therein for extensive discussion.

observational claims into the definition of structural realism — as the Ramsification approach does — also presumes a fairly discredited theory/observation distinction.)

James Ladyman (1998) first distinguished epistemic structural realism from *ontic structural realism* (OSR), the view not just that we should not believe a theory's non-structural claims but that theories, correctly interpreted, *do not make* non-structural claims. Attempts to make this precise within the language-first view of theories has likewise proved difficult: "structure" has generally been identified with relations and distinguished from objects and monadic properties, and (language-first) OSR has been expressed by slogans like 'relations without relata' and 'structure is ontologically prior to objects/ more fundamental than objects'. There have been many[16] attempts to flesh out and make precise these slogans (in the face of the flat-footed objection that relations are definitionally relations between objects).

From the perspective of this paper, ESR and OSR realized in these ways share a common weakness: whatever their broader epistemic and metaphysical advantages, it's far from obvious that they deliver on the original goals of structural realism. 'Structure' as defined in these approaches generally seems just as likely to be discarded on theory change as supposedly 'non-structural' content, so that ESR fails to achieve Worrall's goal of circumventing the argument from theory change. And whatever the virtues of the various attempts to spell out OSR, they are heavily metaphysical, and seem quite far from the original motivation of OSR as avoiding making overly-fine metaphysical distinctions among interpretations of physical theories.

8 Scientific realism on the math-first view

Suppose now that the math-first view, not the language-first view, actually describes theories as they are used in comtemporary physics. What are the implications for scientific realism?

For a start, it needs to be redefined. The standard scientific realist regards our best theories as (approximately) *true*, and mathematically-presented theories are not collections of sentences and so cannot strictly be said to be true or false. But there is a natural alternative: as we saw in section 3, on the language-first view we aim for theories whose empirical consequences are true, but on the math-first view we aim for theories which successfully represent empirical data. So if the move to standard scientific realism is the move from 'theories are true insofar as they make empirical claims' to 'theories are true', the move to a math-first scientific realism can naturally be understood as the move from 'theories successfully represent insofar as they are representing empirical data' to 'theories successfully represent'. That is:

Math-first scientific realism: Our successful scientific theories succeed, at

[16] For instance, (Woolff 2012; McKenzie 2014). (McKenzie 2017) and (Ladyman 2020) provide systematic overviews and references; (Sider 2020, ch.4) is an extended critical discussion.

least approximately, in representing the systems which they are used to model, including features of those systems that are unobservable.

The arguments in favor of math-first scientific realism are basically the negative and positive arguments of section 6, transferred *mutatis mutandis* to the math-first conception of theories: given the theory-ladenness of observation there is no stable alternative to realism other than full-on Humean skepticiism; in any case, the success of our theories at modelling novel observable phenomena is best explained by assuming they are representationally successful.

The standard *objections* to realism, however, are very substantially blunted on the math-first view, as I will demonstrate. Firstly, the notion of approximation in the definition of realism is much less problematic on the math-first view, for the reasons discussed in section 3: mathematically-presented theories have rich resources to handle approximation which are not readily available on the language-first view. In addition (again, as noted in section 3) scientific representation is naturally limited in scope: a math-first realist about (say) classical fluid dynamics is committed only to the fact that certain systems are well-modelled by fluid dynamics and that 'well-modelled' need not be restricted to those features of fluids that are naked-eye observable, not to some broad and hard-to-state claim about classical fluid dynamics being *true*.

Moving on to underdetermination of theory by evidence: as I noted in section 6, the strength of the objection depends on what actual examples of underdetemination can be found. Two classes of underdetermination can be set aside as unthreatening for (either sort of) realist. Firstly, so-called 'weak underdetermination': underdetermination relative to our *current* empirical data, as distinct from strong underdetermination, which is underdetermination with respect to all possible data.[17] It's weakly underdetermined what dark matter is; it's weakly underdetermined what particle physics looks like at energies above the range of the LHC; in each case, we have many theories which make incompatible predictions but we aren't currently able to test those predictions. But here the realist — like physicists themselves — should clearly just be agnostic. (Put another way: precisely because we haven't yet been able to test these theories against each other, they're not currently making the sort of novel confirmed predictions that power the positive argument.)

Secondly, contrivances, like 'everything observable happens according to T, but T isn't true', or 'T, plus there is a swarm of invisible particles that don't interact with anything in T'. Even leaving aside questions about whether 'observable' is well-defined in the first case, the contrivance is fairly clearly a *terrible* explanation, parasitic for its success on T. So the positive argument clearly tells us to accept T, not the contrivance.

The interesting class consists of *scientifically serious* examples of strong underdetermination: alternative theories which are taken seriously by scientists and make the same predictions as each other. And essentially all[18] examples

[17] The terminology is from (Newton-Smith 2000).

[18] The various alternative solutions to the quantum measurement problem are sometimes (e. g. Cordero 2001; Egg 2014; Callender 2020) given as examples of underdetermination; I

I know involve *mathematically equivalent* theories (the example of gravity-as-curvature vs. gravity-as-force has this form, for instance: the mathematics in Weinberg's (1972) force-based presentation of general relativity is equivalent[19] to that in Misner, Thorne and Wheeler's (1973) curvature-based presentation). And there are powerful reasons to think that any good examples *must* be between mathematically-equivalent theories: given the theory-ladenness of observation, it is difficult to see how one could ever *demonstrate* that two theories are empirically equivalent unless either they are mathematically equivalent or differ from mathematical equivalence only through the addition of inert surplus structure. (Norton (2008) gives a similar argument.)

As for theory change: when one successful theory of physics is replaced by another, one almost always finds that the older theory can be reduced to the newer (under certain approximations and in appropriately-restricted domains). On the math-first analysis of reduction (section 5), this reduction is an (approximate, domain-limited) instantiation of the mathematical structure of the older theory's models in the newer theory's. Classical particle mechanics is instantiated by quantum particle mechanics in the large-mass, decoherent regime; Newtonian gravity is instantiated by general relativity for low curvatures and velocities; the heat equation of classical heat-flow theory is instantiated by non-equilibrium statistical mechanics for near-equilibrium conditions.

But if the lower-level theory (i) successfully represents features of a system, and (ii) instantiates the higher-level theory, then there is no contradiction with the claim that the higher-level theory *also* represents (perhaps fewer) features of a system. There is no contradiction between the claim that general relativity correctly represents the motions of the planets, and the claim that Newtonian gravity approximately does so too (with the approximation being detectably imperfect as regards the precession of Mercury's perihelion.) The supposed inconsistency between Einstein's curvature-based account of gravity and Newton's force-based account is invisible on the math-first view of theories: the mathematical structure of the Newtonian model of the solar system is instantiated in the Einsteinian model, and that's all there is to it.[20]

So there is at least a strong case that moving to a math-first account of theories bypasses the key problems for scientific realism. And indeed if a math-first view of theories is correct, it explains why traditional scientific realism runs into difficulties with underdetermination and theory change: 'theories', on the

argue against this in (Wallace 2020b), but in any case the interpretation of quantum mechanics is too contested to be a good source of general-purpose examples.

[19]Weinberg de-emphasizes global topology, but that is not an inherent feature of his approach.

[20]That might suggest that Newton's theory, too, *need not* be read as a force-based account but can be understood as curvature-based. And indeed, the so-called 'Newton-Cartan' reformulation of Newtonian gravity is explicitly curvature-based — and mathematically equivalent to the force-based formulation. Bain (2004) uses this to argue for underdetermination in the Newtonian realm; Knox (2014) and Wallace (2020a) argue that even the 'force-based' descriptions of Newtonian gravity should be understood as describing curved spacetime (hence removing the underdetermination); Weatherall (2016) argues directly for the equivalence of the two theories from a math-first viewpoint.

language-first view, are predicate precisifications of mathematically-presented theories, and so (a) theories may have logically-inequivalent precisifications (leading to underdetermination) and (b) there may be no straightforward logical relation between the precisifications of an older and a newer theory, even when there is a mathematically-well-behaved instantiation relation between them. If the language-first view gave the correct account of scientific practice, then different precisifications are different theories (so that theory choice confronts underdetermination) and theory change involves radical change of theoretical content, including of ontology. But if the math-first view is correct, there is no problem here, even for a realist. As I noted in section 5, reduction and equivalence are just much coarser-grained notions for the math-first view — and the more fine-grained notions used by the language-first view are, from the math-first perspective, picking up illusory distinctions, artifacts of a particular choice of predicate precification.

In fact, this math-first version of scientific realism seems to realize exactly Worrall's goals in presenting structural realism in the first place: by focussing on the mathematical structure of the theory and not on its description in language, it recovers exactly the sort of continuity over theory change that physicists themselves exploit in their continued use of old theories.[21] I think 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, and from truth to successful representation as the distinctive realist claim, is itself sufficient to make this form of realism 'structural'. From here on in, I will use 'math-first structural realism', or just 'math-first realism', as the name for this position: there is no meaningful distinction to draw between 'standard' and 'structural' realism once the math-first position is adopted.

9 Math-first realism and the metaphysics of physics

Math-first structural realism does not in itself require any revision of standard metaphysics, even that part of metaphysics concerned with fundamentality. We can still talk about the fundamental theory of nature, written metaphorically in the Book of the World, and raise metaphysical puzzles about how that theory might be formulated and understood. It does, however, weaken the continuity between physics and metaphysics which appears naturally in language-first scientific realism, in ways that make the methodology of metaphysics somewhat more difficult to establish and justify.

To see this, suppose that the wildest dreams of physicists come true and we finally have a completed microphysics, a Theory of Everything, ToE (whether

[21] I don't want to claim originality for this observation, which occurs in various more-or-less explicit forms at various places in the ontic-structural-realism literature, although generally without the focus I give here on the mathematical details of theory change and theory equivalence: Ladyman (1998) points out something very similar in his advocacy of the semantic view of theories, for instance.)

string theory, or loop gravity, or some other ultraviolet completion of the Standard Model, will not matter here). By assumption, ToE is a mathematically expressed theory, and so cannot be identified with the metaphysicians' Fundamental Theory (FT), which is given in some variant of predicate logic. Yet ToE constrains FT, because given math-first structural realism, ToE succeeds in representing the world as completely as any physics theory can. This seems to mean that FT must be (or at least must include) some predicate precisification of ToE — but for the usual reasons we can expect there to be a great many of these precificiations. So FT is underdetermined by ToE: our completed physics will not tell us what objects there are, or what relations they stand in, even in structural terms (if 'structural terms' means something like the *language-first* version of ontic structural realism). Even a completed physics will significantly underdetermine fundamental metaphysics.

In this sense, math-first realism combined with the standard approach to fundamental metaphysics is a sort of *epistemic* structural realism — but while language-first ESR tells us that some parts of *scientific* theories should not be believed, math-first ESR tells us that complete acceptance of a completed scientific theory still leaves us ignorant of the underlying ontology and ideology.

We could take a somewhat Kantian attitude to this limitation: if the methods of science underdetermine our fundamental metaphysics, and if we adopt the naturalist line that we have no science-transcendant routes to empirical knowledge, then we will just have to accept that things-in-themselves will remain unknown to us. (This is somewhat reminiscent of Russellian monism (cf. (Alter and Pereboom 2019) and references therein) or Langton's (2001) 'Kantian humility' but goes much further: Langton and the Russellian monist believe that we can have knowledge of structure expressed in predicate-logic form even as we are ignorant of the intrinsic nature of things, but given math-first ESR, large amounts of information even about structure in this sense are hidden from us.)

Alternately, we could apply the standard methods of metaphysics to fill the gap: we could posit and criticize various predicate precisifications of ToE, weighing them up against one another on grounds of explanatory power. Indeed, much of modern metaphysics can be interpreted as engaged in this task, with the caveat that only schematic and general features of FT can be discussed in the absence of an actual candidate ToE. But the usual motivation for regarding this method as truth-tracking is that it is continuous with scientists' own methods, and this claim is undermined by math-first ESR, as we can see by considering in more detail the way in which physicists themselves appeal to simplicity and other hallmarks of explanatory power.

We can concede, I think, that some basic appeal to explanation underlies scientific realism and arguably science itself: the no-miracles argument for scientific realism is an inference to the best explanation, after all. But it is important to note that we are not using explanatory power to select scientific realism (or the approximate truth of some specific well-confirmed theory of physics) as preferable to rivals in good standing. The scientific realist does not say that there are lots of good explanations for a theory's novel confirmed predictions but that the theory's approximate truth edges out the others; they say that there are *no other good explanations at all* (no non-miraculous explanations, if you like). Similarly, invocation of inference to the best explanation to decide between T and 'T plus something epiphenomenal' or 'not T, but everything is as if T' does not involve weighing T's explanatory strengths and weaknesses against its rivals: it involves rejecting those rivals outright as unacceptably terrible explanations. The principle in use here is not so much 'inference to the best explanation' but 'inference to the only explanation that is any good'. (See also Deutsch (1997, ch.7)). This is probably too weak a principle to adjudicate between most claims in metaphysics.

Physicists make more substantive appeal to explanatory virtues in deciding between rival physical theories in the absence of conclusive empirical data: for instance, the debate about whether dark matter or a modification of Newtonian dynamics (MOND) best explains various anomalies in astrophysics and cosmology does not at present lend itself to *decisive* empirical testing, and so the debate involves a lot of argument as to which of the various suggestive-but-not-conclusive observations supports one theory and which another. (This is just a playing out of the familiar (Duhem-Quine) observation that realistic theories only make predictions in conjuction with various auxiliary hypotheses.)

But again, this seems quite dissimilar to the choice between *metaphysical* theories. In physics, the rival theories have different empirical consequences and empirical evidence directly bears on them; the arguments about explanatory power arise because technological constraints mean that we don't have all the evidence yet, and indeed the debate constantly evolves as more data comes in. (Twenty years ago, MOND seemed a very serious possibility; at present, the evidence seems extremely strong — though not conclusive — that it is unnecessary. See Weatherall (forthcoming) for references and for arguments in support of this view of the present evidence.) Indeed, a lot of the point of the debate is to help the physics community work out which experiments or observations to fund next: feelings run high in these disputes precisely because research money is at stake! But the metaphysical questions we are considering are not like this: they are underdetermined *in principle* by the science.

The closest match between science and metaphysics would come if there were strong underdetermination: alternative serious (non-contrived) scientific theories which are in principle (not just in practice, not just at present) underdetermined by empirical evidence. If that occurred in physics, and if physicists nonetheless chose between them, then this might well license metaphysicians to do something similar; indeed, in my discussion of standard scientific realism I identified this as a key point of methodological continuity between physics and metaphysics. But we have seen that math-first realism seems to eliminate pretty much all plausible cases of *this sort* of underdetermination: what look like examples of underdetermination on the language-first view just become equivalent formulations of the same theory on the math-first view. If we leave aside the contentious case of the quantum measurement problem (where in any case I doubt philosophers would want to take methodological lessons from physicists!) there just don't seem to be examples here that the metaphysician can use as analogies.

I don't want to claim that it is *impossible* to find justifications to choose between fundamental metaphysical pictures: there are many proposed methods in metaphysics, not all of which rely on the analogy with science. It is to say that that analogy is much weaker than is commonly supposed — at least if the math-first view of theories correctly describes physics practice, and if math-first structural realism is the correct version of scientific realism.

10 Ontic structural realism as math-first metaphysics

Math-first ESR leaves the subject matter of metaphysics largely unchanged even as it complicates its methodology. But one of the main strands of *ontic* structural realism has always been that metaphysics itself must be reconceived in the light of our current science, and that some of the questions metaphysicians ask are not just ananswerable with the tools of science, but incoherent in the light of science. Is there a math-first version of OSR that realizes these goals?

To see what it might look like, note that math-first ESR is a hybrid theory. It adopts a math-first view of *physical* theories, but maintains a language-first view of *metaphysical* theories — indeed, it is the mismatch between the two that generates metaphysical underdetermination. This suggests a simple way to realize math-based OSR: simply permit a metaphysical theory to be given in directly mathematical, rather than linguistic, terms. At the level of fundamental theories, this amounts to supposing that there is no Book of the World, whose every statement is true and which contains a complete description of the world at the most fundamental level; instead, there is a Model of the World, which exactly and completely represents the world at the most fundamental level.

Call this *math-first metaphysics*, and call the combination of math-first scientific realism and math-first metaphysics, *math-first ontic structural realism*. It seems to go some way towards realizing French and Ladyman's idea (section 1) that logic and set theory are imperfectly suited to structuralist metaphysics and that we need something closer to the physics. It also fits into the increasing trend in metaphysics to downplay its relation with philosophy of language (inseperable from metaphysics for most of the 20th century). Sider (2011, p.viii) writes that

> [T]here is a growing consensus: that it is not so important for metaphysical and linguistic theory to neatly mesh. The fundamental metaphysics underlying a discourse might have a structure quite unlike that suggested by the discourse. Whereas a good linguistic theory must fit the suggested structure, good metaphysics must fit the underlying structure.

Math-first metaphysics realizes this separation (albeit in a more radical way than I think Sider has in mind). It realizes, too, Quine's old idea that no part of our conceptual scheme is completely open to revision in light of the evidence, and that even 'central' parts of that scheme like our logic could be revised. Quine had in mind a change of our rules of inference while leaving the syntax of predicate logic unchanged, but his choice to formulate his system in predicate logic is not *a priori* but based on its supposed best fit to the foundations of science. If (as we are supposing) scientific theories are ultimately best given in the math-first framework, Quine's broader methodology should counsel restringing the web of belief so that it is built on mathematised theories and not on the predicate calculus.

But is the picture coherent? I'll freely concede that the central notions in a math-first picture of theories (representation as the theory world relation; mathematical equivalence, in a less-than-precisely-stated form, as the formal part of theoretical equivalence) are significantly less well understood than the analogous notions (reference/satisfaction/truth; synonymy or intertranslatability) on a language-first view. This complaint, again from Sider (in considering somewhat-related structural-realist moves) is reasonable enough:

> [G]reat care is needed to develop the most basic framework for theorizing. Predicate logic isn't some mindless projection of our conceptual scheme. It was developed, with great labour, in a very unforgiving area, the foundations of mathematics, where errors were bound to (and did) have huge consequences. It took a long time to reach the modern viewpoint. (Sider 2020, p.64)

On the other hand, it's not as if reference or meaning are *uncontroversial* either, especially when applied to fundamental theories of metaphysics. And given 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) we need an argument that the most fundamental way to represent the deep structure of the world is through predicate logic rather than some other tool. Sider's argument is that it is very hard to find a framework adequate for theorizing and that predicate logic was developed with just this in mind: he insists (ibid) on "a demonstration that the new proposed framework is adequate to the foundations of mathematics and science". But if the math-first view of physical theories is correct then *predicate logic* is not adequate to the foundations of physics — and not because it fails to realize structuralist dreams but more straightforwardly because it is ill-suited to describe our actually-successful physical theories.

In any case, the goal of this paper is to *state* some defensible forms of structural realism that realize its main goals, not to defend them in detail. Math-first OSR seems a natural way of realizing OSR within the math-first approach to scientific theories; in the next section I will assume its coherence and explore its implications for our ordinary talk of objects, properties and relations.

11 Ontology as derivative

Ladyman and Ross's influential (2007) defense of structural realism is titled *Every Thing Must Go*, and the claim that fundamentally speaking there are no objects is made elsewhere in the ontic structural realism literature (see also, e. g. , French (2014)). It is *more or less* realized in math-first OSR, although it would be more accurate to say that the object/property/relation way of describing reality is not applicable at the most fundamental level.

But here is a hand, and here is another; so there are hands; so there are objects. So this form of OSR, like any (cf Ladyman 2020, section 4) had better have something to say about how to make *non*-fundamental sense of object, property, and relation.

In my view the most plausible way to do so follows a proposal of Simon Saunders (2003) :

> [T]he notion of object is clearest in logic, in the structure of the proposition, but the language of physics is mathematics, not the predicate calculus. . . not all, and perhaps not even the most important part of what physicists know, can accurately be put into words. We must do our best to say what there is, so there will always be a place for objects, understood as objects of predication, but I see no reason why objects in this sense should precisely line up with the constituents of reality, whatever they are, nor with what can be known of them, given that the primary vehicle for understanding reality is mathematics (interpreted mathematics). It is true that set theory can be formalized in *Begriffschrift*; I grant that mathematics, or those parts of mathematics of use to physics, can be reduced to set theory; but I do not think that thereby one will learn what physical objects really are.

(See also Saunders 2016).

I interpret Saunders thus: for a variety of reasons we wish to know the best, the most perspicuous way to describe our physical theories in predicate-logic terms, in the categories of objects and properties and relations. Any such description will be a predicate precisification of the mathematised physical theory; some such precisifications are better than others, measured in terms of intelligibility, explanatory power, simplicity and elegance, and we should adopt the *best* such description. If we then want to say what exists, and what properties there are, according to our theory, we should consult that best description — but there will be nothing more to the claim that certain objects exist and have certain properties, *other than* the fact that those claims are entailed by our best predicate precisification of the true theory. We cannot even say that those objects do not *fundamentally* exist, because claims about objecthood and object-existence have no place in the mathematised theory. The predicate precisification of our most fundamental theory is the most fundamental level at which we can speak of what exists, but it is not the most fundamental level. (Of course, there is a language in which we describe our fundamental theory

— English, at least in my case — and that has a quantifier, but we use that language to describe the mathematical structure we use to represent reality, not to describe reality directly.)

The task of finding the *best* predicate precisification of a theory (according to math-first OSR) is very similar to the task of finding the *true* predicate precisification of a theory (according to math-first ESR). In each case, we can start with the physical theory, try to describe it in language, and then weigh up those descriptions against one another on grounds of explanatory power and the like. The difference, in metaphysical terms, is one of *grounding*: in ESR, the mathematically-presented physical theory is grounded by the precisification (taken as a fundamental theory); in OSR, the reverse is true. But this makes a big difference when it comes to justifying our selection principles. We have seen that for ESR, it is difficult to say why (say) simplicity or explanatory power should lead us from our completed physics to a *true* precisification of that physics; in contrast, it is unmysterious why we might prefer predicate descriptions of the underlying mathematical theory that are simpler or lead to more helpful explanations. Indeed, flatly anthropocentric notions like intuitive appeal, intelligibility or even fit to pre-theoretic beliefs are defensible principles of theory choice in this case. For math-first OSR, the choice of metaphysical description for a theory is closer to Strawson's (1959) descriptive metaphysics, only applied to our physical theories rather than ordinary discourse.

Another difference is that when two precisifications are tied for first place, for ESR there is a fact of the matter as to which — if any — is the *true* theory, whereas for OSR there need be no such fact. The choice of the best predicate precisification of a theory can become vague or indeterminate, and it may even be that one description is better for some purposes and one for others.

To illustrate this, consider section 4's example of the AdS/CFT correspondence. As described in that section, there are strong (theoretical) reasons to expect a mathematical equivalence between a quantum theory of gravity on asymptotically N-dimensional anti-de Sitter (AdS) spacetime, and a conformal quantum field theory (CFT) on the $N-1$-dimensional conformal boundary of that spacetime. Extremely plausibly, predicate precifications of that theory will represent it *either* as a quantum theory of gravity on a spacetime of a certain dimension (say, 4), *or* as a conformal field theory on a space of lower dimension (say, 3).

According to math-based ESR, one of those is *correct.* (This view is defended more or less explicitly by Butterfield (2021).) If, say, the correct description is the conformal-field theory description, then the world is *really* 3-dimensional, spacetime is really non-dynamical, and the 4-dimensional description is derivative on the true description — maybe it is an illusion or useful fiction, maybe a valid but non-fundamental description, but in any case it is grounded by the mathematical description of AdS/CFT, which in turn is grounded by the fundamental 3-dimensional theory. If the correct description is the 4-dimensional description, the converse is true: the world is really 4-dimensional and spacetime is really dynamical. (See figure 1.)

According to math-based OSR, both are useful ways of describing the under-

Figure 1: AdS/CFT correspondence in math-first epistemic structural realism

$$\begin{tikzcd} \text{AdS} \arrow[r, "\text{grounds}"] \arrow[d, Leftrightarrow, "\text{equivalent}"'] & \text{4D description} \\ \text{CFT} & \text{3D description} \arrow[l, "\text{grounds}"'] \end{tikzcd}$$

OR

$$\begin{tikzcd} \text{AdS} \arrow[d, Leftrightarrow, "\text{equivalent}"'] & \text{4D description} \arrow[l, "\text{grounds}"'] \\ \text{CFT} \arrow[r, "\text{grounds}"] & \text{3D description} \end{tikzcd}$$

lying mathematics; neither is correct, but one is more useful in some situations and one is more useful in others. The question 'what is the dimension of spacetime at the most fundamental level' has no answer. (See figure 2.)

So far this discussion has at least notionally concerned *fundamental* physics. Let me now consider how it changes when we consider higher-level theories, and higher-level ontology: as we will see, the difference between the ESR and OSR positions — that is, between language-first and math-first metaphysics — is much less clear in this context.

12 Higher-level ontology

Suppose we have mathematically presented theories of physics LLT (a Low-Level Theory) and HLT (a High-Level Theory), and suppose that LLT mathematically instantiates HLT in some reasonably-well-understood way. For definiteness, let HLT be the physics of continuum fluids or solids, and let LLT be atomic physics. The latter instantiates the former in many different ways in different physical situations, so we should also pick a particular instantiation: let it be of some particular fluid (say, molten iron) or solid (say, solid iron). For the moment, let's proceed under the assumption — the fiction, in our example — that LLT

Figure 2: AdS/CFT correspondence in math-first ontic structural realism

$$\begin{tikzcd}
\text{AdS} \arrow[r, "\text{grounds}"] \arrow[d, Leftrightarrow, "\text{equivalent}"'] & \text{4D description} \\
\text{CFT} \arrow[r, "\text{grounds}"] & \text{3D description}
\end{tikzcd}$$

is a fundamental theory, a Theory of Everything in the sense of the last section. On either (math-first) ESR or OSR, we can seek a description of LLT in terms of objects, properties and relations: for simplicity I'll speak of an *low-level ontology*, LLO for LLT, while acknowledging that the choice goes beyond ontology to various claims about relations and properties. The ontology might be one of point particles, for instance, or it might be some weirder quantum ontology. On ESR, we hope that our ontology is the *true* ontology, which grounds LLT; on OSR, we just seek the *best* description of LLT in ontological terms.

But we naturally use ontological language to describe HLT too: we speak of entities like *fluids* and *parts of fluids*; we attribute properties to them like *density*, *viscocity* and *temperature*. So there is a reasonable question as to how this higher-level ontology HLO is to be understood, and what relation it bears to HLT, LLT and the low-level ontology LLO.

For OSR, there is a fairly simple answer. An ontology, for OSR, is the best description in language (the best predicate precisification) of a theory — and nothing about that account requires that the theory be fundamental. So we can perfectly well start with the high-level theory, taken on its own terms, and ask what is the best precisification of that theory in language, without any reference to that theory's own derivative, higher-level status. The appropriateness of (say) an ontology of fluids and parts of fluids (probably along with spacetime points and spacetime regions), and an associated ideology of temperature, viscocity, occupation of regions of spacetime by parts of fluids, and the like, comes simply from the fact that this ontology offers the best realization in language of the mathematics of fluid dynamics. (And the same holds for the lower-level theory whether or not it is fundamental, so in OSR we can drop the fiction that LLT is fundamental.)

What relation holds between the lower-level and higher-level ontology? In specific cases we *might* be able to state one directly: perhaps fluids can be understood as mereological sums of their constituents. But the fairly radical ontological shifts that often seem to accompany theory change and intertheoretic relations suggest strongly that in general any such account of the relation will

Figure 3: Higher-level ontology in math-first ontic structural realism

$$\begin{tikzcd}
\text{High-level theory} \arrow[r, "\text{grounds}"] & \text{High-level ontology} \\
\text{Low-level theory} \arrow[r, "\text{grounds}"] \arrow[u, Rightarrow, "\text{instantiates (\& grounds)}"] & \text{Low-level ontology} \arrow[u, dashed, "\text{real pattern}"']
\end{tikzcd}$$

be at best approximate. (In the fluid case, the idea that a fluid is the composite of the particles in some mereological sense gets at some of what's going on but seems only a roughly accurate account — cf my comments in section 5.) Even in situations where such an account can be given, what grounds that account is the separate facts that the HLO and LLO are respectively grounded in the HLT and LLT, and the latter instantiates the former. (See figure 3.)

Dennett (1991) offers a picture of higher-level ontology in which high-level entities are explanatorily and predictively useful structures or patterns realised in the dynamics of lower-level theories (and in doing so attempts to finesse the ontological discontinuity one often finds in different levels of description in science). The idea has been adopted in various parts of the structural-realist and philosophy-of-physics literature, e.g. (Ross 2000; Ladyman and Ross 2007; Wallace 2003, 2012; Franklin and Robertson 2021), but it has been difficult to pin down just what a pattern is, and the irreducibly pragmatic nature of Dennett's analysis has raised suspicion that the approach gives a merely instrumental account. The version of math-first OSR I give here offers a possible analysis: given ontologies X and Y, X consists of real patterns in Y iff there are theories T_X and T_Y such that T_Y instantiates T_X and X, Y are respectively the best predicate precisifications of T_X, T_Y. The dose of pragmatism is taken in the construction of the respective ontologies from the theories, but the instantiation of T_X by T_Y lacks a pragmatic component.

The ontology that thus arises from math-first OSR is disunified, being scale- and domain-relative.[22] But the overall *metaphysical picture* is not disunified: unification occurs at the level of mathematical structure. Or, more precisely, it *need not* be disunified: it's a scientific question whether two theories are related by an instantiation relation, and math-first OSR lets us discuss the metaphysics of each theory in its own terms, agnostic as to the answer to that question. In math-first OSR, given a scientifically successful theory we can meaningfully discuss its ontology without any commitment that the theory is 'fundamental':

[22] Ladyman and Ross (2007) also argue that ontic structural realism leads to scale-relativity of ontology, though for somewhat different reasons.

Figure 4: High-level ontology grounded directly in low-level ontology

$$\begin{tikzcd}
\text{High-level theory} \arrow[r, dashed, leftrightarrow, "??"] & \text{High-level ontology} \\
\text{Low-level theory} \arrow[u, Rightarrow, "\text{instantiates (\& grounds)}"] & \text{Low-level ontology} \arrow[l, "\text{grounds}"'] \arrow[u, "\text{grounds}"]
\end{tikzcd}$$

I've argued elsewhere (Wallace 2019a; Wallace 2020b) that this is a desirable feature in any analysis of theories in physics, both because physical theories are not normally treated as fundamental in physics practice and because we don't have any fundamental physical theories as yet.

What about if we assume math-first *epistemic* structural realism, so that the mathematically-presented lower-level theory LLT is grounded in a fundamental theory LLO presented in predicate logic or something similar? (Here we return to the assumption that LLT is fundamental.) I can see three ways we might understand the relation between the high-level ontology HLO and the lower-level theory.

Firstly, we might seek to ground HLO directly in LLO, say by regarding objects in HLO as mereological sums of objects in LLO. (Figure 4.) On this account there is no very direct relation between the high-level *theory* and the high-level *ontology* (again because of the ontological shifts associated with inter-theoretic reduction): the high-level theory takes on a somewhat instrumentalist character, and high-level descriptions of the world are interpreted through terms and ideas derivative directly on the low-level ontology. The so-called 'primitive ontology' view in philosophy of physics (Allori *et al* 2008, Maudlin 2010) seems to be something like this: macroscopic objects are swarms of microscopic particles, and low-level theories make contact with observation directly through our empirical access to the locations of those objects. (This picture is also close to the 'ontology-first' picture of inter-theoretic reduction discussed — though not advocated — by Guo (forthcoming).)

The difficulty here is that our understanding of high-level physics actually does go through the high-level theories. We need no specialized equipment to detect the higher viscocity of honey than of water, the greater solidity of a table than of foam, the higher friction coefficient of ice than of granite, and the higher thermal conductivity of metal than of glass, and in each case we have fairly good understandings of these concepts in our higher-level science. A higher-level ontology that gets disconnected from those ideas will fail to connect to scientific (including folk-scientific) practice. (See also the related criticisms

Figure 5: High-level ontology grounding high-level theory

$$\begin{tikzcd}
\text{High-level theory} & \text{High-level ontology} \arrow[l, "\text{grounds}"'] \\
\text{Low-level theory} \arrow[u, Rightarrow, "\text{instantiates (\& grounds)}"] & \text{Low-level ontology} \arrow[l, "\text{grounds}"'] \arrow[u, dashed, "??"']
\end{tikzcd}$$

of Sider (2021), and of David Albert, quoted in Sider, *ibid.*)

Secondly, we might just apply the ESR methodology at the high level: try to construct a high-level ontology directly from the high-level theory, and take the latter to be grounded in the former. (Figure 5.) This solves the scientific problem (the high-level ontology meshes naturally with the high-level theory), but at the price of a highly disunified metaphysics: there is no obvious relation holding between the high-level and low-level ontology. Disunified metaphysics of this kind is sometimes advocated (e. g. , by Cartwright (1999)) but normally in the context of a disunified *physics*, whereas here we are assuming a mathematical reduction of the high-level to the low-level theory. Put in terms of ground, HLT is grounded both by HLO and (indirectly) by LLO, without any grounding relation between HLO and LLO, and that seems bizarre. We could fix it by simply stipulating that LLO grounds HLO, but that seems like a metaphysically brute fact unless connected to the actual relations between HLO and HLT, between HLT and LLT, and between LLT and LLO.

Finally, we could simply adopt the OSR approach to high-order ontology: it is the best description in object/property/relation terms of the mathematically-presented high-level theory. (Figure 6.) This seems to me the only available strategy for math-first ESR that preserves both (a) the relation between low-level ontology and high-level ontology, and (b) the relation between high-level ontology and high-level theory. It does so, however, by treating fundamental ontology entirely differently from high-level ontology: if we drop the fiction that atomic physics is fundamental, for instance, then both continuum mechanics and atomic physics will have an ontology grounded in the mathematically-presented theory. The full situation looks (like a far more complex version of) Figure 7: at every level except the most fundamental, the mathematical description of a theory is prior to its description in ontology, and only at the fundamental level is the grounding relation reversed, so that the fundamental ontology grounds the fundamental physics.

This means that for any part of metaphysics *except* that concerned with the world on the most fundamental level, it does not really matter whether we

Figure 6: High-level ontology grounded in high-level theory

$$
\begin{tikzcd}
\text{High-level theory} \arrow[r, "\text{grounds}"] & \text{High-level ontology} \\
\text{Low-level theory} \arrow[u, Rightarrow, "\text{instantiates (\& grounds)}"] & \text{Low-level ontology} \arrow[l, "\text{grounds}"'] \arrow[u, dashed, "\text{real pattern}"]
\end{tikzcd}
$$

adopt a math-first or language-first approach to metaphysics (that is, whether we adopt math-first OSR or math-first ESR). Only fundamental metaphysics proceeds on the assumption that the object/property/relation way of describing the world is prior to the mathematically-presented theory; at all other levels in the metaphysics of physics, the metaphysical description is derivative on that theory, and higher-level ontology is disunified and scale-dependent in just the same ways as we saw for OSR.

To me, this makes the case for an ontological description at the most fundamental level unmotivated, and the methodology of exploring that description obscure, especially given that at the current stage of progress in physics we do not even know if there *is* a fundamental theory, and if there is we have strong reasons to think that it is ontologically discontinuous with higher-level theory regarding even such basic concepts as space and time. But I will not argue the point further here.

13 Conclusions

'Structural realism' has come to refer ambiguously to two rather different conceptions of scientific realism and scientific metaphysics. The first conception is defined by the original motivations of Worrall, French, Ladyman *et al*: to avoid the pessimistic argument from theory change while holding on to the core no-miracles argument for realism; to avoid drawing distinctions between theories on a finer grain than physics seems to need; to reconceive metaphysics in a way more closely aligned with physics. This first conception is officially committed to the inadequacy of traditional object/property/relation ways of doing metaphysics, but it's advocates have struggled to state a positive alternative in a way that critics have found coherent. The second conception attempts a structuralist philosophy of science within the confines of the object/property/relation approach (within 'a more traditional metametaphysical outlook', as Sider (2020, p.55) puts it). As philosophy of science, this amounts to a selective skepticism

Figure 7: Only fundamental ontology is prior to mathematical theory

$$\begin{tikzcd}
\text{Continuum mechanics} \arrow[r, "\text{grounds}"] & \text{Continuum ontology} \\
\text{Atomic physics} \arrow[u, Rightarrow, "\text{instantiates (\& grounds)}"] \arrow[r, "\text{grounds}"] & \text{Particle ontology} \arrow[u, dashed, "\text{real pattern}"] \\
\text{Standard Model of Particle Physics} \arrow[u, Rightarrow, "\text{instantiates (\& grounds)}"] \arrow[r, "\text{grounds}"] & \text{Quantum-field ontology} \arrow[u, dashed, "\text{real pattern}"] \\
\ldots \arrow[u, Rightarrow, "\text{instantiates (\& grounds)}"] \arrow[r, "\text{grounds}"] & \ldots \arrow[u, dashed, "\text{real pattern}"] \\
\text{Unknown Theory of Everything} \arrow[u, Rightarrow, "\text{instantiates (\& grounds)}"] & \text{Fundamental ontology} \arrow[l, "\text{grounds}"'] \arrow[u, dashed, "\text{real pattern}"]
\end{tikzcd}$$

about the unobservable; as metaphysics, it replaces a flat rejection of traditional metametaphysics with a priority thesis that treats objects as less fundamental than relations, or with an anti-individualist metaphysics. Whatever the merits of this second conception, it has grown distant from the motivations of the first: distant enough, indeed, to have enthusiastic advocates e. g. (Dasgupta 2009, Esfeld, Lazarovici, Lam, and Hubert 2017; Esfeld, Deckert, and Oldofredi 2017) who are unmoved by those motivations and are drawn to 'structural realism' for largely distinct reasons.

The source of this dichotomy is the distinction between language-first and math-first approaches to physical theories (or between syntactic and semantic approaches, in the more standard terminology). The math-first approach leads to a different conception of the theory/world relation, to a coarser-grained notion of theoretical equivalence, and to a more continuous conception of theory change and inter-theoretic reduction. The ontological discontinuities which drive standard examples of underdetermination and discontinuity of theory change are invisible at the level of a mathematical description of theories. Simply adopting this view of theories in physics, and making the concomitant changes to realism (from a search for truth to a search for successful representation) already suffices to realize the main philosophy-of-science goals of structural realism: as such (and as anticipated by structural realists like Ladyman and French) math-first scientific realism is already a form of structural realism, distinct from the second conception but very close to the aims of the first.

Within this conception of structural realism, the distinction between epistemic and ontic structural realism concerns the relative fundamentality of object/property/relation descriptions of a physical system: is that description metaphysically primary, with even a completed physics being only a coarse-grained description of underlying reality, or should we adopt a math-first conception of metaphysics where the most fundamental description of the world is via its mathematical structure and a description in the language of object, property and relation is simply the best available approximation to that structure. The former approach is closer to the conventional assumptions of metaphysics but faces serious questions about the justification of its methodology given the discontinuity it implies between physics and metaphysics; the latter is metaphysically radical and yet actually offers a better justification for a conventional methodology. In any case, the distinction between the two makes sense only at the level of yet-unknown fundamental physics: even given epistemic structural realism, the only realistic way to do the metaphysics of higher-level theories is math-first, with the ontology and ideology derivative on the mathematics.

This paper is only a sketch of a proposal, and leaves many open questions. To name just two of the most important: firstly, my account leans heavily on a notion of mathematical equivalence that I have not developed. In practice, we seem to know it when we see it, whether in algebraic topology or in quantum gravity, but it would be good to have something sharper to say. (If there is one place where I expect the 'pure' semantic view of theories to have to be developed and complicated, it is here.) And secondly, I have confined myself to physics, but there is more to science than physics, and in much of the rest of science theories are less explicitly mathematized than in physics. My hunch — my intuition, if you like — is that this is because the subject matter of the special sciences is much more naturally and unambiguously characterized in terms of objects and properties, so that the distinction between language-first and math-first views blurs, and that this in turn reflects the fact that the categories of our language were evolved to represent goings-on at medium scales on the surface of Earth and do so much more naturally than at the esoteric scales with which physics deals. But it is only a hunch, to be substantiated or refuted by those with greater mastery than I of the special sciences in question.

Even in physics the math-first view of physical theories may not be correct. (Indeed, it probably is not correct in every detail, although I would be surprised if it requires modification radical enough to upend the basic conclusions of the paper.) But whether it is right is a matter of what scientific practice is: it is a question with metaphysical implications, but it is not a choice between metaphysical theories. If scientific practice does not support a conception of theories as collections of sentences in formal or natural language, we cannot build a philosophy of science on that assumption, and any language-first *metaphysics* must be built with a clear recognition of how distant a fundamental metaphysical theory is from a theory of physics.

To me, that makes it attractive to seek for a math-first conception of metaphysics that closes that distance, and I have sketched one conception of how it might look — a conception that seems to realize the main goals of the original ontic structural realists. If advocates of a more traditional conception of metaphysics are not persuaded by its virtues, I hope it is at least stated clearly enough for them to have something to get their teeth into.

Acknowledgements

I am grateful to Gordon Belot, Bixin Guo, Eleanor Knox, Kerry Mckenzie, Tushar Menon, Jim Weatherall, and especially to James Ladyman and Simon Saunders, for useful discussions.

References

Albert, D. Z. (1996). Elementary quantum metaphysics. In J. T. Cushing, A. Fine, and S. Goldstein (Eds.), *Bohmian Mechanics and Quantum Theory: An Appraisal*, Dordrecht, pp. 277–284. Kluwer Academic Publishers.

Allori, V., S. Goldstein, R. Tumulka, and N. Zanghí (2008). On the common structure of Bohmian mechanics and the Ghirardi-Rimini-Weber theory. *British Journal for the Philosophy of Science 59*, 353–389.

Alter, T. and D. Pereboom (2019). Russellian monism. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2019 ed.). https://plato.stanford.edu/archives/fall2019/entries/russellian-monism/.

Anderson, P. W. (1984). *Basic Notions of Condensed Matter Physics*. Boulder, CO: Westview Press.

Ashcroft, N. W. and N. D. Mermin (1976). *Solid State Physics*. Orlando, FL.: Harcourt.

Bain, J. (2004). Theories of Newtonian gravity and empirical indistinguishability. *Studies In History and Philosophy of Science Part B: Studies In History and Philosophy of Modern Physics 35*, 345–376.

Balescu, R. (1997). *Statistical Dynamics: Matter out of Equilibrium*. London: World Scientific.

Barrett, T. W. and H. Halvorson (2016). Morita equivalence. *Review of Symbolic Logic 9*, 556–582.

Belot, G. (2012). Quantum states for primitive ontologists: a case study. *European Journal for Philosophy of Science 2*, 67–83.

Butterfield, J. (2011a). Emergence, reduction and supervenience: a varied landscape. *Foundations of Physics 41*, 920–959.

Butterfield, J. (2011b). Less is different: Emergence and reduction reconciled. *Foundations of Physics 41*, 1065–1135.

Butterfield, J. (2021). On dualities and equivalences between physical theories. In C. Wúthrich, B. Le Bihan, and N. Huggett (Eds.), *Philosophy Beyond Spacetime: Implications from Quantum Gravity*, pp. 41–77. Oxford: Oxford University Press.

Callender, C. (2020). Can we quarantine the quantum blight? In S. French and J. Saatsi (Eds.), *Scientific Realism and the Quantum*, pp. 57–77. Oxford: Oxford University Press.

Campbell, J., J. Huston, and F. Krauss (2018). *The Black Book of Quantum Chromodynamics: A Primer for the LHC Era*. Oxford: Oxford University Press.

Cartwright, N. (1999). *The Dappled World: A Study of the Boundaries of Science*. Cambridge: Cambridge University Press.

Cordero, A. (2001). Realism and underdetermination: Some clues from the practices-up. *Philosophy of Science 68*, S301–S312.

Dasgupta, S. (2009). Individuals: An essay in revisionary metaphysics. *Philosophical Studies 145*, 35–67.

Demopoulos, W. and M. Friedman (1985). Bertrand Russell's *The Analysis of Matter*: its historical context and contemporary interest. *Philosophy of Science 52*, 621–639.

Dennett, D. C. (1991). Real patterns. *Journal of Philosophy 87*, 27–51. Reprinted in *Brainchildren*, D. Dennett, (London: Penguin 1998) pp. 95–120.

Deutsch, D. (1997). *The Fabric of Reality*. London: Penguin.

Dizadji-Bahmani, F., R. Frigg, and S. Hartmann (2010). Who's afraid of Nagelian reduction? *Erkenntnis 73*, 393–412.

Dorr, C. (2010a). Of numbers and electrons. *Proceedings of the Aristotelian Society 110*, 133–181.

Dorr, C. (2010b). Review of Ladyman and Ross, "Every Thing Must Go: Metaphysics Naturalized". *Notre Dame Review of Books 2010*(6).

Earman, J. (1993). Underdetermination, realism, and reason. *Midwest Studies in Philosophy 18*, 19–38.

Egg, M. (2014). *Scientific realism in particle physics: a causal approach.* Berlin: de Gruyter.

Esfeld, M., D.-A. Deckert, and A. Oldofredi (2017). What is matter? the fundamental ontology of atomism and structural realism. In A. Ijjas and B. Loewer (Eds.), *Philosophy of Cosmology*. Oxford: Oxford University Press.

Esfeld, M., D. Lazarovici, V. Lam, and M. Hubert (2017). The physics and metaphysics of primitive stuff. *British Journal for the Philosophy of Science 68*, 133–161.

Feynman, R. P. (1967). *The Character of Physical Law*. Cambridge, Mass.: MIT Press.

Fodor, J. A. (1989). Making mind matter more. *Philosophical Topics 17*, 59–79.

Franklin, A. and E. Knox (2018). Emergence without limits: The case of phonons. *Studies in the History and Philosophy of Modern Physics 64*, 68–78.

Franklin, A. and K. Robertson (2021). Emerging into the rainforest: Emergence and special science ontology. Forthcoming; online at http://philsci-archive.pitt.edu/19912/.

French, S. (2014). *The Structure of the World: Metaphysics & Representation*. Oxford: Oxford University Press.

French, S. and J. Ladyman (2003). Remodelling structural realism: Quantum physics and the metaphysics of structure. *Synthese 136*, 31–56.

French, S. and J. Saatsi (2006). Realism about structure: the semantic view and nonlinguistic representations. *Philosophy of Science 73*, 548–559.

Glymour, C. (2013). Theoretical equivalence and the semantic view of theories. *Philosophy of Science 80*, 286–297.

Guo, B. (Forthcoming). Ontology-first vs. theory-first approach to reduction: A case study from statistical mechanics. In submission.

Halvorson, H. (2012). What scientific theories could not be. *Philosophy of Science 79*, 183–206.

Hawley, K. (2010). Critical notice of *Every Thing Must Go*. *Metascience 19*, 174–179.

Hirsch, E. (2011). *Quantifier Variance and Realism: Essays in Metaontology*. New York: Oxford University Press.

Isham, C. and K. Kuchar (1985). Representations of spacetime diffeomorphisms 1: Canonical parametrized field theories. *Annals of Physics 164*, 288–315.

Kaplan, J. (2016). Lectures on AdS/CFT from the bottom up. http://sites.krieger.jhu.edu/jared-kaplan/files/2016/05/AdSCFTCourseNotesCurrentPublic.pdf.

Ketland, J. (2004). Empirical adequacy and Ramsification. *British Journal for the Philosophy of Science 55*, 287–300.

Knox, E. (2014). Newtonian spacetime structure in light of the equivalence principle. *British Journal for the Philosophy of Science 65*, 863–880.

Kuhn, T. S. (1970). Postscript-1969. In *The Structure of Scientific Revolutions* (2nd enlarged ed.)., pp. 174–210. Chicago: University of Chicago Press.

Ladyman, J. (1998). What is structural realism? *Studies in History and Philosophy of Science 29*, 409–424.

Ladyman, J. (2020). Structural realism. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2020 ed.). https://plato.stanford.edu/archives/win2020/entries/structural-realism/.

Ladyman, J. and D. Ross (2007). *Every Thing Must Go: Metaphysics Naturalized*. Oxford: Oxford University Press.

Langton, R. (2001). *Kantian Humility: Our Ignorance of Things in Themselves*. Oxford: Oxford University Press.

Laudan, L. (1981). A confutation of convergent realism. *Philosophy of Science 48*, 19–49.

Maldacena, J. M. (1998). The large N limit of superconformal field theories and supergravity. *Advances in Theoretical and Mathematical Physics 2*, 231–232.

Maudlin, T. (2010). Can the world be only wavefunction? In S. Saunders, J. Barrett, A. Kent, and D. Wallace (Eds.), *Many Worlds? Everett, Quantum Theory, and Reality*. Oxford: Oxford University Press.

Maudlin, T. (2018). Ontological clarity via canonical presentation: Electromagnetism and the AharonovBohm effect. *Entropy 20*, 465. https://doi.org/10.3390/e20060465.

McKenzie, K. (2014). Priority and particle physics: Ontic structural realism as a fundamentality thesis. *British Journal for the Philosophy of Science 65*, 353–380.

McKenzie, K. (2017). Ontic structural realism. *Philosophy Compass 12*. https://doi.org/10.1111/phc3.12399.

Misner, C. W., K. S. Thorne, and J. A. Wheeler (1973). *Gravitation*. New York: W.H. Freeman and Company.

Nagel, E. (1961). *The Structure of Science*. New York: Harcourt.

Newman, M. (1928). Mr Russell's "causal theory of perception". *Mind 37*, 137–148.

Newton-Smith, W. H. (2000). Underdetermination of theory by data. In W. Newton-Smith (Ed.), *A Companion to the Philosophy of Science*, pp. 532–536. Oxford: Blackwell.

Newton-Smith, W. S. (1981). *The Rationality of Science*. London: Routledge.

Ney, A. and D. Albert (Eds.) (2013). *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*. Oxford: Oxford University Press.

Norton, J. (2008). Must evidence underdetermine theory? In M. Carrier, D. Howard, and J. Kourany (Eds.), *The Challenge of the Social and the Pressure of Practice: Science and Values Revisited*, pp. 17–44. Pittsburgh: University of Pittsburgh Press.

Odom, B., D. Hanneke, B. D'Urso, and G. Gabrielse (2006). New measurement of the electron magnetic moment using a one-electron quantum cyclotron. *Physical Review Letters 97*, 030801.

Oppenheim, P. and H. Putnam (1958). The unity of science as a working hypothesis. *Minnesota Studies in the Philosophy of Science 2*, 3–36.

Paul, L. (2012). Metaphysics as modeling: the handmaiden's tale. *Philosophical Studies 160*, 1–29.

Psillos, S. (1999). *Scientific Realism: How Science Tracks Truth*. London: Routledge.

Psillos, S. (2004). Is structural realism possible? *Philosophy of Science 68*, S13–S24.

Putnam, H. (1975). What is mathematical truth? In *Mathematics, Matter and Method*, pp. 60–78. Cambridge: Cambridge University Press.

Quine, W. (1957). The scope and language of science. *British Journal for the Philosophy of Science 8*, 1–17.

Reichenbach, H. (1958). *The Philosophy of Space and Time*. New York: Dover.

Ross, D. (2000). Rainforest realism: a Dennettian theory of existence. In D. Ross, A. Brook, and D. Thompson (Eds.), *Dennett's Philosophy: a comprehensive assessment*, pp. 147–168. Cambridge, Massachusets: MIT Press/Bradford.

Saunders, S. (2003). Structural realism, again. *Synthese 136*, 127–133.

Saunders, S. (2016). The emergence of individuals in physics. In A. Guay and T. Pradeu (Eds.), *Individuals Across the Sciences*, pp. 165–192. New York: Oxford University Press.

Schaffner, K. F. (1967). Approaches to reduction. *Philosophy of Science 34*, 137–147.

Sider, T. (2011). *Writing the Book of the World*. Oxford: Oxford University Press.

Sider, T. (2020). *The Tools of Metaphysics and the Metaphysics of Science*. Oxford: Oxford University Press.

Sider, T. (2021). 3D in high-D. Online at http://tedsider.org/papers/3D_in_high-D.pdf.

Sklar, L. (1982). Saving the noumena. *Philosophical Topics 13*, 89–110.

Sklar, L. (2003). Dappled theories in a uniform world. *Philosophy of Science 70*, 424–441.

Strawson, P. (1959). *Individuals: An Essay in Descriptive Metaphysics*. London: Methuen.

Suppes, P. (1962). Models of data. In E. Nagel, P. Suppes, and A. Tarski (Eds.), *Logic, Methodology, and Philosophy of Science: Proceedings of the 1960 International Congress*, pp. 252–261. Stanford: Stanford University Press.

van Fraassen, B. (1989). *Laws and Symmetry*. New York: Oxford University Press.

Van Fraassen, B. C. (1980). *The Scientific Image*. Oxford: Oxford University Press.

Varadarajan, M. (2007). Dirac quantization of parametrized field theory. *Physical Review D 75*, 044018.

Wallace, D. (2003). Everett and Structure. *Studies in the History and Philosophy of Modern Physics 34*, 87–105.

Wallace, D. (2012). *The Emergent Multiverse: Quantum Theory according to the Everett Interpretation*. Oxford: Oxford University Press.

Wallace, D. (2015). Fields as bodies: a unified presentation of spacetime and internal gauge symmetry. https://arxiv.org/abs/1502.06539.

Wallace, D. (2018). The case for black hole thermodynamics, part II: statistical mechanics. *Studies in the History and Philosophy of Modern Physics 66*, 103–118.

Wallace, D. (2019a). Isolated systems and their symmetries, part I: General framework and particle-mechanics examples. Forthcoming in *Studies in the History and Philosophy of Science*.

Wallace, D. (2019b). Who's afraid of coordinate systems? An essay on representation of spacetime structure. *Studies in the History and Philosophy of modern Physics 67*, 125–136.

Wallace, D. (2020a). Fundamental and emergent geometry in Newtonian dynamics. *British Journal for the Philosophy of Science 71*, 1–32.

Wallace, D. (2020b). On the plurality of quantum theories: Quantum theory as a framework, and its implications for the quantum measurement problem. In J. Saatsi and S. French (Eds.), *Scientific Realism and the Quantum*, pp. 78–102. Oxford: Oxford University Press.

Wallace, D. (2021). Against wavefunction realism. In S. Dasgupta, R. Dotan, and B. Weslake (Eds.), *Current Controversies in Philosophy of Science*, pp. 63–74. New York: Routledge.

Weatherall, J. O. (2016). Are Newtonian gravitation and geometrized Newtonian gravitation theoretically equivalent? *Erkenntnis 81*, 1073–1091.

Weatherall, J. O. (2019). Theoretical equivalence in physics. *Philosophy Compass 14*. In two parts; https://doi.org/10.1111/phc3.12591 (part 1), https://doi.org/10.1111/phc3.12591 (part 2).

Weatherall, J. O. (forthcoming). The philosophy behind dark matter. Forthcoming.

Weinberg, S. (1972). *Gravitation and Cosmology: Principles and Applications of the General Theory of Relativity*. New York: John Wiley and sons.

Winther, R. G. (2021). The structure of scientific theories. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2021 ed.). https://plato.stanford.edu/archives/spr2021/entries/structure-scientific-theories/.

Witten, E. (1998). Anti de Sitter space and holography. *Advances in Theoretical and Mathematical Physics 2*, 253–291.

Woolff, J. (2012). Do objects depend on structures? *British Journal for the Philosophy of Science 63*(63), 607–625.

Worrall, J. (1989). Structural realism: the best of both worlds? *Dialectica 43*, 99–124.