

Real Patterns in Physics and Beyond

David Wallace*

April 20, 2024

Abstract

I apply Dennett’s ‘real patterns’ idea to the ontology of physics, and specifically to the puzzle of how to relate the very different ontologies one finds at different scales in physics (e.g. particles vs continua, or fields vs particles). I argue that real patterns provide part but not all of the answer to the puzzle, and locate the rest of the answer in the structural-realist idea that ontology in general is secondary to (mathematically-presented) structure. I make some suggestions for the application of these ideas outside physics, including in the philosophy of mind context that motivated Dennett’s original proposal.

This is a preprint version of a chapter to appear in a forthcoming volume on real patterns, edited by Tyler Millhouse, Steve Petersen, and Don Ross, to be published by MIT Press.

1 Introduction

Dennett’s seminal “Real Patterns” (Dennett 1991) is first and foremost an intervention in the philosophy of mind: its goal is to defend Dennett’s (1981, 1987) intentional-stance theory of content from the supposed dilemma that either beliefs are real (in which case there had better be some concrete *things* in the world, some collections of neural tissue or similar, with which they can be identified) or they are not (in which case they don’t exist and so cannot play a role in any science of the mind worth the name). In ‘Real Patterns’, beliefs are real patterns in the behavioral dispositions of believing systems, discernible through the intentional stance: abstracta to be sure, but useful abstracta, abstracta that earn their keep through the predictive and explanatory work they do and the shortcuts they allow us to take compared to the impossibly demanding and in any case uninformative task of working out what those systems would do *qua* systems of concrete matter. Abstracta that useful ought to be considered along with paradigm physical entities like electrons as part of what science tells there is in the world.

*Department of History and Philosophy of Science, and Department of Philosophy, University of Pittsburgh; david.wallace@pitt.edu

But science — even physics — is not a monolith, and once we dig deeper into physics, these ‘paradigm’ physical entities start looking more problematic, and ‘concrete’ matter starts looking less concrete. Physics describes the world on many levels, from quantum field theory, through the quantum-mechanical descriptions of nuclei and electrons, up to the quasi-classical world of molecular physics and thence to the comparatively familiar world of fluids and solids. And at those different levels, its ontology looks very different — at one point fields, at another particles, at still another continuous matter. How those various ontologies are related — how ontology at one level might be emergent from different ontology at a lower level — is not at all obvious or simple, and over the last twenty years I and others¹ have deployed the real-patterns idea inside physics itself, to elucidate how an ontology of one kind can be emergent from an ontology of a quite different kind.

This paper is about that project, and its relation to another major theme in recent philosophy of physics: *ontic structural realism*, the proposal (at least as I read it) that the Quinean project of reading ontological content off our best scientific theories assumes that those theories can be faithfully understood as sets of claims in natural language, and that once this claim is abandoned, the very idea of looking for the ontology of a theory becomes at best a second-order activity. In section 2 I spell out the ontological puzzle physics poses, and in section 3 I explain how the real-patterns idea transforms but does not fully resolve the puzzle. In section 4 I describe (my version of) structural realism, and then in section 5 I relate it back to real patterns and to the project of ontology. In section 6 I briefly consider how these ideas might find application beyond physics.

The physics I discuss in this paper is well established and I do not attempt to give original references.

2 Emergence of ontology in physics: molecules and fluids

My concern in this paper is with inter-theoretic relations *within* physics, and that might seem puzzling, even perverse. Don’t the more interesting questions concern the reduction of other sciences *to* physics? And insofar as there are reductions inside physics — of atoms to electrons and nuclei, of thermodynamics to statistical mechanics — aren’t they supposed to be the paradigm cases of well-understood reduction, to contrast with the knottier cases in the special sciences? But if what makes a science ‘special’ is that it lacks universality of scope, almost all of physics is a special science — physicists study stars, or metals, or plasmas,

¹See, e.g., (Wallace 2003, 2010, 2012); (Seifert 2023); (Franklin 2024); (Franklin and Robertson 2021); (Ladyman and Ross 2007); (Ross 2000). (Not all these authors agree: Ladyman and Ross, in particular, interpret real patterns not as a between-levels relation from lower-level to higher-level ontology, but as a within-levels relations between ontology and phenomena.)

or lasers — and much of the work of physics is constructing one inter-theoretic relation or another.

Those relations are far harder to understand philosophically than is sometimes recognized in philosophy: the reduction of thermodynamics to statistical mechanics, in particular, is hugely controversial and comprises one of the central problems of modern philosophy of physics.² But for all that they are the simplest inter-theoretic relations we find in the sciences, and for that reason they serve as test beds for any general account of inter-theoretic relations. If we cannot understand how theories are related even within physics, we have little hope of understanding the more complex cases; if we can understand the physics case, perhaps we can learn wider lessons.

Let's consider a comparatively simple example.³ Our system will just be a sealed box, a few centimeters across, containing some fluid - chlorine gas, water, molten iron, the details will not matter. We can distinguish (at least) two levels of description of that system:

The interatomic level: The system is a collection of order 10^{24} atoms (or molecules), characterized by the masses of those atoms and by the functional form of the fairly-short-range forces by which neighboring atoms interact with one another. In some cases we might need to engage with the substructure of the atoms, but in many cases the interactions *between* atoms occur at a sufficiently different energy scale from the interactions *within* atoms that we can get away with treating them as elementary. A full description of the system requires the 6×10^{24} real numbers which represent the positions and velocities of each atom; the evolution of the system is governed by Newton's laws, supplemented by the details of the particular interatomic interactions that apply.⁴

The fluid level: The system is a continuous fluid, characterized by parameters like specific heat capacity and viscosity and by how its density, pressure and temperature vary across the box. A full description of the system requires several smooth functions, encoding those variation; the evolution of the system is governed by the Navier-Stokes equation of fluid dynamics.

For the moment, let's pretend that the interatomic-level description is the most fundamental description we have, and disregard the realms of atomic and nuclear and particle physics.

This is about as simple a case of an inter-theoretic relation as we can imagine. If we can't understand how *fluids* emerge from atoms, our prospects for

²See, e. g. , (Sklar 1993; Albert 2000; Frigg 2007; Myrvold 2022; Wallace 2023).

³For technical discussion of the relevant physics, see, e. g. , (Balescu 1997) and references therein.

⁴I assume here that the system can be understood classically, without the use of quantum mechanics except perhaps to calculate the form of the interatomic interaction. This is not always valid technically and is conceptually subtle even when it is technically permissible, for reasons connected to the quantum measurement problem (see (Wallace 2016) for discussion), but including those subtleties would if anything strengthen the points I wish to make here, so I disregard them for simplicity.

understanding how *beliefs* emerge from them look dismal. And yet already there are subtleties. Consider: what's in the box? 'Atoms', is the obvious answer: our best theory of the box quantifies over atoms, asserts that atoms exists, so (following Quine (1948)) if we accept that theory, we accept that there are atoms.

But isn't it also true that there is a fluid in the box? In the interatomic-level description there are no fluids to be found: only atoms and the void. Fluids are not quantified over according to the interatomic theory, and that (we're pretending) is our *fundamental* theory of the box, so (again following Quine) there *is* no fluid in the box. That seems absurd: the empirical evidence for water is pretty conclusive. Indeed, the empirical evidence for the *theory* of fluids is pretty solid, and that theory does quantify over fluids. So how do we reconcile this with the absence of fluids at the interatomic level? The metaphysics-of-science literature has offered three main answers.

The first is *disunity*: there is the fluid theory, valid for fluids, and the interatomic theory, valid for the interactions of molecules; both are well evidenced by the data; both apply in their particular domains; perhaps there are overlaps in their applicability but there is no hierarchy of fundamentality, no reason but prejudice to regard the 'lower-level' theory as any more of a guide to ontology than the higher-level theory. Theses of this kind have been quite popular in the philosophy of the special sciences, beginning with Fodor's disunity thesis (Fodor 1974) and continuing in the work of Cartwright (1999), Dupré (1983) and others: they are tied to the denial of reductionism, of the idea that the generalities and models of the special sciences could in any feasible way be derived from those of physics.

For myself, I find those claims overstated, and based on an overly reductive understanding of reductionism itself (cf (Dennett 1995, pp.80-83)). Among other observations, the special sciences are replete with physics ideas: biochemists discuss molecular bonds; biophysicists study the physics of animal movement; organic chemists rely on thermodynamics; doctors use CT scans and positron emission tomography. But what makes them even plausible is that reduction between physics and the special sciences is at best incomplete and in general aspirational. We may or may not have good reasons to think that cognition, say, reduces to physics in principle, or even that we will someday obtain that reduction ourselves, but today we have, at most, suggestive fragments of that reduction.

The project of reducing fluid dynamics to interatomic physics, however, is far from aspirational: huge parts of it have actually been completed. We know how to derive the Navier-Stokes equations from the microdynamical equations; we know in principle how to calculate quantities like specific heat capacity and viscosity from the microdynamics; in many cases (dilute gases most completely, but not only there) we have actually done those calculations and got answers that match the data. It's not *easy*, and the devil is sometimes in the details,⁵

⁵See (Batterman 2021) for philosophical discussion of some of the subtleties, and of the need in many cases to establish a robust 'mesoscale' understanding of a continuum system. It is also

but there is not a lot of room to deny the claim that the fluid-level physics is derivable, in the appropriate regime, from the atomic description. And then it seems inexplicable how, if we actually have a well-controlled account of the mathematical relation between the two theories, it could still be the case that their ontologies are just disconnected.

The second answer, *composition*, embraces reduction in the fullest sense: the fluid actually is present in the lower-level description, because the fluid just *is* all of the atoms together. At the formal level, this means supplementing our lower-level theory with a *mereology*, a logic of whole and part, so that as well as quantifying over the individual atoms we quantify over the various mereological sums of atoms. The fluid is then identified with one of those mereological sums — the sum of *all* the atoms — and the spatial parts of the fluid with the mereological sums of subsets of the atoms.

Composition meshes naturally with the traditional (Nagelian) approach to reduction. There (Nagel 1961), intertheoretic reduction is logical deduction of one theory from another with the aid of “bridge laws”, auxiliary axioms which coordinate the terms of one theory with those of another — for instance, “the mass of the fluid in region R is Nm iff there are N atoms each of mass m in region R ”. For that notion of reduction genuinely to *reduce* one theory to another, those bridge laws have to be more than mere contingencies: they must be something like identities. ((Lewis 1970) is perhaps the best-known working out of this idea.) And then something like mereological composition seems unavoidable: there need to be things in the ontology of the microscopic theory that we can treat as identical to the fluid and its parts, and an austere ontology of particles alone contains no candidates. The overall package is undeniably elegant, and meshes well both with the formal methods of analytic philosophy and with our intuitive understanding of composition.

Its mesh with actual physics is less impressive. We might expect to find mereology — or some rough-and-ready version of it — in the actual mathematics by which fluid dynamics was derived, but it is not at all apparent there. (In the clearest derivations I know, the fluid gets identified with some subcollection of the lowest-frequency Fourier modes of the particle distribution.) And the mereological sum of the molecules seems quite different in nature from the fluid posited by fluid dynamics — the former consists of isolated points scattered through empty space, the latter is a smooth continuum. The situation only gets worse in more esoteric regimes of physics. The statement that a proton is composed of three quarks, for instance, is somewhere between a forgivable heuristic in popular science and a story we tell children; quarks themselves are identified with certain excitations of an underlying field, *very roughly* in the way that an ordinary wave is identified as a moving disturbance of a fluid. There is no obvious logic of composition identifiable here that transcends the messy

generally the case that the higher-level theories are irreversible in the statistical-mechanics sense, whereas the lower-level theories are reversible, so that somewhere a time-asymmetric assumption has to be introduced; a full understanding of the reduction therefore also requires some engagement with the vexed question of the arrow of time; see, e. g., (Albert 2000; Price 1996; Wallace 2013).

details of the specific cases; indeed, those specific cases often mix up ontological categories, ‘identifying’ a higher-level object with something like a property of the lower-level objects.

If neither disunity nor composition offer a scientifically appropriate account of higher-order ontology, that brings us to the third answer: *instrumentalism*. Perhaps there *is* no higher-order ontology: it is just convenient to pretend that there is. The instrumentalist need not worry about establishing some precise connection between higher-level and lower-level ontology, because really there is only lower-level ontology — but it is sometimes pragmatically helpful to pretend, to act as if there is high-level ontology. That pretense can make it simpler to carry out various scientific tasks, to predict and explain the phenomena, but ultimately those phenomena are only phenomena at the high level. The instrumentalist about fluids will say that ‘fluid’ is just a convenient way of talking, but that really there are just atoms in certain configurations — ‘arranged fluid-wise’, to borrow a term from the metaphysics literature.

Applied to intentional states, this is already radical. Applied to physics, it becomes dizzying. Once we drop our pretense, we remember that the theory of classical atoms and molecules is not fundamental either. It rests on the quantum mechanics of the subatomic, which rests on the layered quantum field theories of particle physics. Even the deepest of those is not fundamental: the Standard Model of Particle Physics is an effective theory, a low-energy, large-scale approximation to something about which we can only speculate (string theory being the furthest developed of those speculations). To be an instrumentalist about ‘higher-level’ physics is to be an instrumentalist about every scientific entity ever described in an empirically-confirmed theory. Indeed, it is to be an instrumentalist about scientist themselves, about their — our — experiments and language. Even space and time are generally expected to be emergent consequences of a deeper, non-spatio-temporal physics — to the instrumentalist, they too don’t really exist.

At this point, instrumentalism about the high level starts to look incomprehensible, if not incoherent, and we seem forced to admit at least *some* higher-level ontology. But how much? Suppose two theories agree that atoms exist, but disagree about whether fluids do; what, actually, is the form of the disagreement? It cannot be stated using the vocabulary of the interatomic theory: that theory, in itself, requires no mereology. It cannot make a difference to anything stated in the language of interatomic ontology, for the same reason. What *does* it make a difference to? The question cannot be answered naturalistically: science is silent on whether fluids really exist or are just an indispensable-in-practice fiction.⁶ We seem to have been led from a quintessentially naturalistic problem — given that systems are described on various levels by various different theories, how do we think about the relation between those theories’ respective ontologies? — step by step into the deepest jungles of analytic metaphysics.

⁶If you doubt this, try asking a few fluid dynamicists whether fluids really exist or are just a useful way of talking about atoms. Then try the analogous question about atoms themselves.

3 Real patterns and Dennett's Criterion

The last section is a sort of Rorschach test for philosophers. For some, they illustrate exactly why metaphysics is indispensable, for all that neo-positivists and other philistines might claim otherwise: think clearly enough about any deep problem in philosophy of science and eventually you will have to get your metaphysics straight, and in doing so will have to engage in the practice of metaphysics. For others — like myself and, I imagine, Dennett — the fact that we seemed to be led so deeply away from anything naturalistic is an indication that something went badly wrong in our setup of the problem, and we should work out what it is.⁷

At this point 'Real Patterns' comes as a breath of fresh air. In Dennett's proposal, higher-level entities are not literally *composed* of lower-level entities, but nor are they simply instrumentalist fictions: rather, they are patterns — or, in terms more natural to a philosopher of physics, *structures* — in the distributions and dynamics of the lower-level entities. The fluid is not a mereological sum of the atoms in the box: it is a pattern in their behavior, a pattern elucidated by the mathematical derivation of fluid dynamics from interatomic physics, and it is a *real* pattern, and not just a 'mere' pattern, because it is such a powerful, explanatory, predictive pattern. There is a flexibility, a situation-dependence, to the notion of pattern not shared by something as rigid as composition: a fluid can be *truthfully* described as continuous simply because a continuum is the best way to describe the pattern that fluid dynamics discerns in the microphysics. Similarly, *pattern* need not respect ontological categories: the pattern (an object) might be realized not in the objects of the lower-level theory but in its properties and quantities.

Once you understand the idea, and once you learn a little physics, you begin to see real patterns everywhere. Light is a travelling pattern in the excitations of the electromagnetic field (just as, in 'Real Patterns', a glider is a travelling pattern in the excitations of the cells in the Game of Life). Physicists describe the vibrations in a crystal by treating it as a gas of phonons, the quanta of sound: these phonons are patterns in the movements of the atoms in the crystal. Electrons are patterns in the excitations of the quantum fields of the Standard Model.

In previous work (Wallace 2003, 2012) I called the general principle here

Dennett's Criterion: A macro-object is a pattern, and the existence of a pattern as a real thing depends on the usefulness — in particular, the explanatory power and predictive reliability — of theories which admit that pattern in their ontology.

The theory of fluids, for example, is *de facto* indispensable in studying large-scale features of certain molecular systems — various of their features would be

⁷To avoid being misunderstood: I don't want to suggest that naturalistically inclined philosophers cannot benefit from engagement with modern metaphysics. My own development of these ideas has been greatly helped by (admittedly critical) engagement with it: I'll note in particular the influence of Sider (2020, 2023).

in principle inexplicable and in practice impossible to calculate without using fluid dynamics — and so fluids are real.⁸

A real-patterns approach does *drastically* better at capturing the way physicists themselves talk about ontology than do theories based on composition or disunity (while avoiding the nihilism of a thoroughgoing instrumentalism towards physics). In physics as in philosophy of mind, real patterns offer the promise of a middle way between instrumentalism and realism, which closely fits the practice of physics. Physicists are extremely permissive about introducing new higher-level ontology, and then talking about it as robustly as they did their old ontology: they are not concerned about any particular way that the new ontology is related to the old but they do require it to be part of a theory derived or constructed from the lower-level physics, and they do require it to be *useful*.

(In my own original use of Dennett’s criterion, I used physics practice to justify the claim that ontology *in general* works this way in physics, and then applied it to one specific controversial case: the ‘branches’ or ‘worlds’ that appear in the Everett (Many-Worlds) interpretation of quantum mechanics. The idea was that the ‘worlds’ are real patterns in the evolving quantum system, so that according to the general ontological criteria which physicists tacitly use, the worlds really exist, as higher-order, non-fundamental ontology.)

That said, while I think something along the lines of the real-patterns approach is correct, by itself it still leaves some metaphysical questions unanswered. For one, while the idea of a pattern is certainly intuitive, it is harder to pin down what it means to quantify over patterns in an underlying ontology, especially when that underlying ontology is as alien and unintuitive as some of those found in physics. For another, Dennett distinguishes between *real* patterns and *mere* patterns, as (in effect) does my phrasing of Dennett’s criterion. But then what does it actually add to say that a pattern is *real*? It doesn’t mean that it *exists*, in the Quinean sense: all the patterns exist, in the sense that we’re quantifying over them when we identify some as mere and some as real. Nor will it do to get pragmatic and say that it is useful to treat some patterns *as* real: nothing in scientific practice will distinguish the scientist who does this from the scientist who simply *uses* the “real” patterns without applying the label ‘real’ to them.

Perhaps we should instead read Dennett’s criterion as some sort of abstraction principle: “Pattern X exists iff the underlying objects are arranged X-wise and positing X is . . .”. (In my own version of Dennett’s criterion, the . . . is filled in by some statement of explanatory and predictive power). If so, what is the status of that principle? It is tempting to read it pragmatically, as a statement about when it is useful to use language that way — but that gets us to quite

⁸There is a close analogy here to Dennett’s discussion of the indispensability of the intentional stance in (Dennett 1987, pp.225-27): in practice it is infeasible for us to predict what intelligent agents will do except through the intentional stance, but even if alien super-scientists did have the ability to make atom-by-atom predictions of what a human being would do, they would still be missing something explanatory if they didn’t know that the intentional stance licensed much simpler predictions.

an instrumentalist conception of higher-order ontology, one that becomes even more uncomfortable when it is not just beliefs and desires but solid matter itself that is to be understood as a real pattern. Worries along these lines lead real-pattern advocates like Ladyman and Ross (2007) to fill in the . . . with something a little more objective: informational compressibility from a physically possible perspective. But then what is the status of that principle — is it some meta-physical *a priori* truth, and if so, how do we come to know it? If two fans of the real-patterns idea fill in the blanks a little differently, how is that disagreement to be resolved? Most of all, what really distinguishes the person who holds the abstraction principle to be *true* from the person who regards it as a mere fiction? The impasse between compositional and instrumentalist approaches to higher-order ontology threatens to recur in more naturalistic guise.

There is something frustrating about this whole dialectic. Again: our starting point is as naturalistic a philosophical question as one could wish — what is the world like, given that we describe it with different physical theories at different scales — but our attempts to answer it keep ensnaring us in meta-physical tangles that naturalistic methods cannot straighten out. There must still be something wrong with how we are thinking about the problem; the real pattern idea is a major step forward but it can't be the whole of the solution.

4 Mathematics and the world

Let's step back a moment. We are trying to understand the relation between what some system, some part of the world, is like according to a higher-level theory (in our example, fluid dynamics) and what it is like according to a lower-level theory (like interatomic physics) that purports to describe the same system on different scales. But at least in the context of physics, it is odd that there can even be a substantive question — after all, *ex hypothesi* we understand how interatomic physics describes the system, *and* we understand how fluid dynamics does, *and* we understand the mathematical relation between the two systems. How can there be anything substantive left to understand?

To answer this, we need to look more closely at what these theories of physics actually are. So far I have been writing as if a scientific theory was something like a description of the world (or a part of it) in language: that is, a theory *says* what things there are in the world, and what properties they have and relations they stand in, and how all this changes over time. (It's clear that this is the notion of theory that Quine, at least, has in mind: something given in the first instance in a natural language like English, but which at least in idealization we can imagine being reformulated in the predicate calculus.) But at least in physics, scientific theories aren't like that any more, if indeed they ever were. Physics theories are highly mathematized, and the relation they bear to the systems they describe is not best thought of through linguistic notions of truth and reference and satisfaction but by some more general notion of *representation*.

Why does this matter? Because the description of the system given in the

mathematics is typically more coarse-grained than a purely linguistic description would be. Consider again the fluid: it is characterized in part by a function that represents the density at each part in the fluid, but there are two ways to think about density:

Density fundamentalism: Density is just a primitive feature of pointlike parts of a fluid; integrating it over regions of the fluid gives a derivative quantity, mass.

Mass fundamentalism: The real primitive is mass: every⁹ nonzero-volume part of the fluid has a mass; the mass of a disjoint sum of parts is the sum of their separate masses; it is density that is the derivative quantity, defined as the limit of the ratio of mass to volume as the latter approaches zero.

Try to formalize the theory of the fluid in predicate logic and you will have to make a call on questions like this: the two choices lead to different logical vocabulary and, indirectly, to different metaphysical commitments. But mathematically they are interchangeable: fluid dynamics doesn't care, and neither do fluid dynamicists. Or more precisely: they might care about which way of describing density is clearer and more useful, but not about which is true.

Examples like this multiply and become more radical as we move to parts of physics further from the manifest image. Consider a classical field theory like electromagnetism. Mathematically (and in the absence of charged matter) the theory represents the electromagnetic field by a pair of vector field, assigning an electric and a magnetic field vector to every point in space. Should we interpret this as saying that there is an extended entity — a field — whose parts occupy regions of space and have field strengths as their fundamental properties? Or is 'field' just a *façon de parler*: are 'field strengths' new properties of the points of space? The two interpretations differ radically in ontology: according to one, there are continuum many new entities, the various parts of the field; according to the other, there are no non-spatial entities at all — and yet the mathematics of electromagnetism makes no distinction between them.

Following (Wallace 2022), a *predicate precisification*¹⁰ of a mathematically-stated theory of physics is an account *in natural language* of what a system is like according to that theory. Precisifications need to commit on various questions of ontology and ideology that the theory itself is silent on, but 'precise' here is not a virtue term: a predicate precisification commits to more about the system than physics itself licenses. That is: a theory of physics underdetermines its own description in natural language.¹¹

⁹Here I elide some measure-theoretic technicalities: it is not possible to define a volume measure on *every* subset of the continuum. See, e.g., (Halmos 1950).

¹⁰I concede that this is not a beautiful phrase. I welcome suggestions for improvement, but there are false friends lurking: 'predicate description' sounds more elegant, but unavoidably suggests that what we describe in natural language is the theory itself (a mathematical structure) rather than the system that the theory represents.

¹¹There is an interesting analogy with Dennett's (1987, ch.5) distinction between *beliefs*

It is tempting to think that some predicate precisification of a successful theory must be true, so that scientific theories represent the world by virtue of constraining, though not fixing, the real description of the world in object-property language. If so, science would underdetermine reality, and there would be truths about the physical world invisible to the scientific method — perhaps permanently opaque to us, perhaps accessible through the methods of metaphysics. (The former position is one form of *epistemic structural realism* (Worrall 1989; Ladyman 1998), the view that science determines structure but not nature; the latter position characterizes the approach of metaphysicians of physics like Sider (2011, 2020), Maudlin (2018), and North (2021).) It is equally tempting to think that an interpretation of a formalized physical theory is a predicate precisification of a theory, telling us how uninterpreted mathematics is to be related to the concrete world.

Both temptations should be resisted. To make a mathematically stated theory subordinate to a natural-language description is to privilege natural language over other representational tools, and there is no non-question-begging reason to do so, and centuries of evidence from physics to the contrary (going back at least to Galileo’s famous observation that mathematics is the language of nature). Language is an exceptionally good representational tool, but it is not our only such tool: even prior to considering scientific theories themselves, think of maps or diagrams, of Kant’s thesis of the centrality of the intuition of space, and of the manifold non-linguistic forms of representation theorized in cognitive science¹²).

Nor can an ‘interpretation’ be a predicate precisification, at least if to interpret a theory is required to understand it, and if physicists understand their own theories: physics practice is not committed to any particular choice of precisification. What it is to learn a new physics theory, in practice, is not to be told how it maps to a world of objects and properties: it is to coordinate that theory with other theories, often mathematically presented, that are antecedently understood, to learn through those coordinatizations how the theory is to be used, and in due course to bootstrap one’s way to an understanding of the theory that is not purely parasitic on previous theories. Much the same is true of learning a new language, or a new word in an old language, and if the process of doing so cannot be fully described except in terms of already-understood representations, that is just the familiar Neurath/Quine observation that we must rebuild our ship while sailing in it (Quine 1969, pp.126-127).

This approach to scientific (or at least physics) theories might be called “mathematics-first realism”. (It is a form of *ontic structural realism* (Ladyman 1998), though that term has been used in wildly different ways by different authors. I understand it broadly in the sense of (McKenzie 2024).) Much more

(intentional states characterized through the intentional stance, applicable to any system which can usefully be ascribed them) and *opinions* (commitments made in language, available only to language users). Opinions are much more finely grained than beliefs, but that fineness of grain does not necessarily reflect anything factive — sometimes it is just an artifact of the sharp edges that language requires.

¹²See, e. g., (Pitt 2022) and references therein.

could be said in its explication and defense (I give my own version in much more detail in (Wallace 2022; Wallace 2024)). For now, I will take as read that it is along the right lines and consider its implications for scientific ontology.

5 Ontology as secondary

If the deep theories of physics do not represent the world in terms of objects and properties, does that mean that fundamentally there *are no* objects and properties? No, for two reasons. The first is shallowly semantic: to say that objects *do not exist* is already to concede the fundamentality of the framework for describing the world that does so in terms of objects, and then to advance a radical metaphysical thesis within that framework. (‘There are no objects’ is a claim in natural language.) The second is more substantive: object-property talk is useful, even in deep physics. Giving accounts of a physical system in language is heuristically and pedagogically helpful; it builds intuition, and helps coordinate different partial descriptions of the world with one another; it aids understanding and suggests new directions for work. A good predicate precisification of a theory is a valuable tool for thinking (Dennett 2013a). But it is not fundamental: its claims are true, insofar as they are true, in virtue not of its correspondence with the world but of it being a predicate precisification of a representationally successful mathematized theory.

This makes the pursuit of ontology into a second-order activity: it tells us not directly what the world is like but how best we can describe it with the tools of natural language. The benefit of this somewhat deflationary account of ontology is that it liberates us methodologically. If we seek *fundamental truth*, then appeals to simplicity, to intuition, to fit with the manifest image, all seem question-begging: why think that the world is simple, or that our intuitions are correct about it. But if we want that description in language of a mathematized theory that is most helpful to us, then of course these are virtues: who would choose a complicated, unintuitive predicate precisification when a simple, intuitive one is at hand? (It also permits agnosticism, even pluralism, when different, nominally-incompatible predicate precisifications are tied for first place: just use whichever one is more suitable for the task at hand.¹³)

The approach to ontology I suggest here has much in common with Dennett’s (2013b) proposal for metaphysics as “sophisticated naïve anthropology”, aimed at delivering a “metaphysics of the manifest image”, without supposing that in doing so we are “limning the ultimate structure of reality”. I am unsure quite how playfully Dennett intends his suggestion, but I intend mine entirely seriously: it is valuable to have clear predicate precisifications of theories in physics, and developing them is one (not the only) way philosophers of physics can contribute.¹⁴

¹³Compare Feynman: “every theoretical physicist who is any good knows six or seven different theoretical representations for exactly the same physics” (Feynman 1967, p.162).

¹⁴A few examples of philosophy of physics that fairly self-consciously fits that framework: Simon Saunders’ (2013, 2016) work on individuality of particles; Eleanor Knox’s (2013, 2014)

And with this account of ontology in hand, we can finally sort out the metaphysical tangles that occupied us for the first half of the paper. Consider once again our fluid, and for now assume (unrealistically) that we start off knowing only the lower-level theory, the mathematically-presented theory of interacting atoms. That theory has a (very natural and obvious) predicate precisification: the objects are the atoms; their essential properties are their masses and perhaps charges; their spatial location can be understood as an additional primitive property, or a relation between particles and spacetime points (observe that even in this paradigm case we have underdetermination, invisible to physics).

That theory also, given certain auxiliary assumptions about dynamics and initial state, mathematically instantiates (approximately) the Navier-Stokes theory of fluid dynamics: that is, some subcollection of its degrees of freedom, appropriately coordinatized, can be well approximated as having self-contained and autonomous dynamics described by that equation. There is nothing pragmatic about this claim: it is simply a piece of mathematics, a theoretical discovery about the theory which we might be led to through surprising empirical data or just work out theoretically. And because of that piece of mathematics, among the scientifically interesting facts about the system are facts about those autonomous high-level dynamics, facts better studied via the Navier-Stokes equation than via interatomic physics. In studying those facts, it will often be helpful to have a description of fluid dynamics in object-property terms, and again there is a very natural and obvious choice of predicate precisification: the objects are the parts of the fluid, the properties include densities (or perhaps masses) and temperatures, and parts of the fluid occupy parts of space.

What makes it true that, at atomic scales, there are atoms, is that we quantify over atoms in the predicate description we use for the system at those scales. What makes it true that at larger scales there are fluids and their parts is that we quantify over those parts in our large-scale description. And in each case our use of the description is legitimated by a combination of not-at-all-pragmatic facts about the mathematical representation of the system and somewhat-pragmatic facts about the best predicate precisification of that representation. Note that our ontology is *scale-relative* (Ladyman and Ross 2007, ch.4): we quantify over fluids (when talking about large-scale physics) and we quantify over atoms (when talking about smaller-scale physics) but we don't quantify over both at once.¹⁵

What is the relation between the low-level and high-level ontology? There need be no self-contained way of stating it, beyond the known relations between theory and ontology at each level and the equally-known relation between theories.¹⁶ Those relations, collectively, leave no metaphysical residue in need of explication. That said, there *might* be, in specific cases, and indeed the lack

spacetime functionalism (see also the discussion in Knox and Wallace 2023); Chris Timpson and my spacetime state realism (Wallace and Timpson 2010, Wallace 2012, ch.8).

¹⁵Not normally, anyway. Separation of scales is an important but not inviolate rule in physics: 'the subatomic particle left a trail in the vapor in the cloud chamber' is an example of a scientifically-legimate use of mixed ontology.

¹⁶Cf Guo (2023), who contrasts this 'theory-first' sense of reduction with the 'ontology-first' approach presupposed in much metaphysics of science.

of any metaphysical requirement of a relationship gives us latitude to cash out those relationships in approximate or heuristic terms. There is clearly *some* sense in which a fluid is made up of atoms — the part of a fluid in region X has properties which supervene on the properties of the atoms in region X , and some of those properties (like mass) are simply additive functions of the properties. Is that a mereological relation. . . *sorta* (cf. Dennett 2013a, pp.96-7.) It's got a lot in common with the idea that mereology formalizes, for sure. Can we sharpen things up so that the relation is really-and-truly mereological, for some particular version of mereology? Maybe. . . but it's not entirely clear what would be gained in doing so.

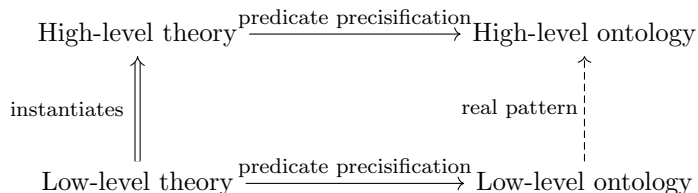
Similarly, a hydrogen atom is *sorta* made up of a proton and an electron: there's a lower-level subatomic theory whose natural ontology is electrons and nuclei, a higher-level interatomic theory whose natural ontology is atoms, and a well-understood mathematical relation between the two, giving rise to a relation between the two ontologies that looks quite a lot like composition. A proton is *sorta* made up of quarks, but now there is a lot more heuristic and metaphor in that claim. In each case something like composition is happening, but the place to look if you want sharp and exact details is the mathematics, not the ontology, and when you look, you find the composition metaphor gets filled in and corrected very differently in different cases.

And in some cases composition does not even approximately capture the intertheoretic relation. Waves, and their quantized particles, are not even *sorta* made up of the medium that is waving: the right form of words is something closer to 'moving disturbance'. But still there will be some relation between lower and higher level ontology: differences in high-level ontology reflect differences in high-level structure, which must result from differences in low-level structure, differences that are reflected in differences in low-level ontology. It would be good to have a general name for this relation, one that does not presuppose composition and reflects the structural, dynamical nature of the intertheoretic relations underlying it. I have a suggestion, one that I hope reflects the spirit of Dennett's own proposal: in this case, the higher-level ontology consists of *real patterns* in the lower-level ontology.

6 Closing the circle

Physics is atypical of the sciences: the systems it studies are more alien and un-intuitive, and yet comparatively more simple, and its theories are mathematized to a far higher degree. The uncritical use of physics as a model for science more broadly has caused much confusion in philosophy of science. So I am hesitant to draw broader lessons about scientific metaphysics from my observations here. Yet my theme has been how ideas first developed in the philosophy of mind have found purchase in philosophy of physics, so it seems a shame not to *try* returning the favor. What follows should be taken in an appropriately tentative spirit, as an invitation to those who know the territory better.

Figure 1: Theory and ontology in intertheoretic reduction



To begin with: scientific theories outside physics are for the most part¹⁷ not plausibly thought of purely as mathematical structures. Natural language runs much more deeply through these theories, far too deeply to be understood merely as a heuristic-but-dispensible predicate precisification.

I find this unsurprising, for two reasons. The first, and most obvious, is that most scientists study systems far more complicated than physicists do: no system in the institutional purview of physics has a fraction of the complexity of a bacterium, far less a human mind or an ecosystem or economy. To some degree, the mathematization of physics reflects the fact that precisely because physics is simpler, physicists have had much greater quantitative success in modelling the world.

But there is a deeper reason: language as a representational tool evolved to help us get by on human scales, dealing with phenomena on scales of centimeters to kilometers, from milliseconds to decades. On those scales, the mismatch between even a physics description of the world and the description given in natural language is not that bad, especially if we are generous in how we construe ‘natural language’ and allow it to include geometric concepts like angle or length, and certain abstracta like collections. It is only when we leave the regime of the macroscopic that linguistic descriptions show their limitations, and much of the special sciences is concerned with that macroscopic regime. So I would expect even a hypothesized ‘completed’ biology or psychology (even supposing that makes sense) to be much more language-like than physics is.

But if most scientific theories are not collections of mathematical models, nor are they simply collections of natural-language sentences. I suspect the right way to describe them is a messy mixture of interconnected representational practices: pieces of mathematics, chunks of natural language equipped with additional resources, concepts and tools and models borrowed from physics, and of course the large amount of tacit knowledge that one learns from experiment

¹⁷There are exceptions, for instance: some parts of economics, of quantum chemistry, of population genetics. Even here, though, the theory is usually to be seen as a mathematized chunk of some larger, more pluralistically-formulated, representation.

and observation and fieldwork and cannot losslessly describe. Much of how representation functions in this context will be at most partially and imperfectly described through the truth-satisfaction-reference. If so, supposing that ‘ontology’ is the right fundamental tool to describe the world according to that theory may be as unwise as in physics, and we may do better, as in physics, to treat our fully-linguistic descriptions of phenomena at least partially as a pedagogical and heuristic tool rather than the full and literal truth.

(Doing so, among other virtues, may help us to recognize that certain apparently-legitimate questions arise because of the artificially sharp distinctions that language forces us to draw. I’ll risk some tentative examples: ‘Is a fetus a human being?’, ‘are viruses alive’, and ‘are wolves and dogs the same species’ are in each case grammatically well-posed questions, yet in each case it seems that we can have a perfectly satisfactory understanding of the underlying science without having an answer to that question.)

I’ll finish by considering again the original question in the ontology of mind that provoked ‘Real Patterns’ in the first place: on Dennett’s account of intentionality, do people really have beliefs or is that just a useful fiction? Recall the context: we assume a person’s behavior is pretty well, but very complicatedly, described by some comparatively low-level theory of the brain and body: some ‘sub-personal cognitive psychology’ (Dennett 1981).¹⁸ But we also assume that same person can be quite well, and much less complicatedly, described, through the intentional stance: through attributing to them beliefs and desires that are updated according to their informational environment and according to which their actions are rational. And the intentional-stance description does not float free of the sub-personal explanation: it had better be the case that we can understand the lower-level mechanisms that make it the case that the intentional stance holds (and which in some circumstances, like brain damage or LSD, qualify or invalidate it).

Does the person has beliefs? In the account of ontology I’ve offered (and throwing caution to the wind by applying it far outside physics), the answer is a clear YES: an ontology of beliefs and desires is the most helpful ontological description of the intentional-system theory, and that theory is correct, in the sense that it does indeed give a fairly-robust, fairly-reliable account of the person’s behavior. To be is to be the value of a bound variable; our best predicate precisification of intentional-systems theory quantifies over beliefs; so, at the scales and in the circumstances where we apply that theory, beliefs exist. Maybe it will turn out that those beliefs can be identified, exactly or approximately, with some sub-personal bit of ontology, but that’s an additional scientific hypothesis not relevant to the higher-level ontology. The qualifiers on our ontological claim — ‘beliefs exist at certain scales and certain circumstances’ — do not mark beliefs off as having some second-class status: all ontology works this way, and all higher-level ontological claims are claims about real patterns in lower-level ontology, and everything posited in physics outside the most extreme regimes

¹⁸A complication is that Dennett expects intentionality to recur, in simpler but robust form, at the sub-personal level; I put this aside for simplicity.

of quantum gravity is ‘higher-level’ in this sense. If beliefs don’t *really* exist, then fluids don’t *really* exist; atoms and electrons don’t *really* exist; most likely space and time don’t *really* exist.

There appears to be room for doubt about the reality of the beliefs and desires revealed by the intentional stance because there appears to be a contrast between abstracta, mere *patterns* in physical matter, and concrete things actually *comprised* of physical matter. But from the perspective of modern physics this contrast is an illusion, caused by the fact that composition gives a reasonable — but far from perfect — account of inter-theoretic relations at scales not too far from human imagination. Step back and look at the bigger picture, and it’s real patterns all the way down.

Acknowledgements

In developing the ideas in this paper I have benefitted greatly from conversations with Neil Dewar, Alex Franklin, Bixin Guo, Eleanor Knox, James Ladyman, Kerry McKenzie, Tushar Menon, Katie Robertson, and Simon Saunders, and from discussions at the 2024 Santa Fe Institute workshop “Investigating Reality – A Philosophical, Mathematical, and Scientific Exploration”. The present volume is also an opportunity to thank Dan Dennett for the tremendous influence his work has been on my own, and for providing so compelling a model of how to do philosophy as a contribution to, and communication with, science.

References

- Albert, D. Z. (2000). *Time and Chance*. Cambridge, MA: Harvard University Press.
- Balescu, R. (1997). *Statistical Dynamics: Matter out of Equilibrium*. London: World Scientific.
- Batterman, R. W. (2021). *A Middle Way: A Non-Fundamental Approach to Many-Body Physics*. Oxford: Oxford University Press.
- Cartwright, N. (1999). *The Dappled World: A Study of the Boundaries of Science*. Cambridge: Cambridge University Press.
- Dennett, D. C. (1981). Three kinds of intentional psychology. In R. Healey (Ed.), *Reduction, Time, and Reality*. Cambridge University Press. Reprinted and expanded in Dennett (1987), pp. 43–81.
- Dennett, D. C. (1987). *The intentional stance*. Cambridge, Mass.: MIT Press.
- Dennett, D. C. (1991). Real patterns. *Journal of Philosophy* 87, 27–51. Reprinted in *Brainchildren*, D. Dennett, (London: Penguin 1998) pp. 95–120.
- Dennett, D. C. (1995). *Darwin’s Dangerous Idea: Evolution and the Meanings of Life*. Simon and Schuster.

- Dennett, D. C. (2013a). *Intuition Pumps and other Tools for Thinking*. New York: W.W.Norton.
- Dennett, D. C. (2013b). Kinds of things — towards a bestiary of the manifest image. In D. Ross, J. Ladyman, and H. Kincaid (Eds.), *Scientific Metaphysics*, pp. 96–107. Oxford: Oxford University Press.
- Dupré, J. (1983). The disunity of science. *Mind XCII*, 321–346.
- Feynman, R. P. (1967). *The Character of Physical Law*. Cambridge, Mass.: MIT Press.
- Fodor, J. (1974). Special sciences (or: The disunity of science as a working hypothesis). *Synthese 28*, 97–115.
- Franklin, A. (2024). Incoherent? no, just decoherent: How quantum many worlds emerge. *Philosophy of Science 91*, 288–309.
- Franklin, A. and K. Robertson (2021). Emerging into the rainforest: Emergence and special science ontology. Forthcoming; online at <http://philsci-archive.pitt.edu/19912/>.
- Frigg, R. (2007). A field guide to recent work on the foundations of thermodynamics and statistical mechanics. In D. Rickles (Ed.), *The Ashgate Companion to the New Philosophy of Physics*, pp. 99–196. London: Ashgate.
- Guo, B. (2023). Ontology-first vs. theory-first approach to reduction: A case study from statistical mechanics. *Philosophy of Science*, forthcoming; preprint at <https://philsci-archive.pitt.edu/22733/>.
- Halmos, P. R. (1950). *Measure Theory*. New York: Litton Educational Publishing.
- Knox, E. (2013). Effective spacetime geometry. *Studies in the History and Philosophy of Modern Physics 44*, 346–356.
- Knox, E. (2014). Newtonian spacetime structure in light of the equivalence principle. *British Journal for the Philosophy of Science 65*, 863–880.
- Knox, E. and D. Wallace (2023). Functionalism fit for physics. <https://philsci-archive.pitt.edu/22655/>.
- Ladyman, J. (1998). What is structural realism? *Studies in History and Philosophy of Science 29*, 409–424.
- Ladyman, J. and D. Ross (2007). *Every Thing Must Go: Metaphysics Naturalized*. Oxford: Oxford University Press.
- Lewis, D. (1970). How to define theoretical terms. *Journal of Philosophy 67*, 427–446. Reprinted in David Lewis, *Philosophical Papers*, Volume I (Oxford University Press, Oxford, 1983).
- Maudlin, T. (2018). Ontological clarity via canonical presentation: Electromagnetism and the aharonov–bohm effect. *Entropy 20*, 465.
- McKenzie, K. (2024). Structuralism as a stance. *Philosophy of Physics 2*, 1.

- Myrvold, W. (2022). Philosophical issues in thermal physics. *Oxford Research Encyclopedia in Physics*. doi: 10.1093/acrefore/9780190871994.013.42.
- Nagel, E. (1961). *The Structure of Science*. New York: Harcourt.
- North, J. (2021). *Physics, Structure, and Reality*. Oxford: Oxford University Press.
- Pitt, D. (2022). Mental Representation. In E. N. Zalta and U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy* (Fall 2022 ed.). Metaphysics Research Lab, Stanford University.
- Price, H. (1996). *Time's Arrow and Archimedes' Point*. Oxford: Oxford University Press.
- Quine, W. (1948). On what there is. *Review of Metaphysics* 2, 21–38.
- Quine, W. (1969). Natural kinds. In *Ontological Relativity and Other Essays*, pp. 114–138. New York: Columbia University Press.
- 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, Massachusetts: MIT Press/Bradford.
- Saunders, S. (2013). Indistinguishability. In R. Batterman (Ed.), *The Oxford Handbook of Philosophy of Physics*, pp. 340–380. Oxford University Press.
- 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.
- Seifert, V. (2023). The chemical bond is a real pattern. *Philosophy of Science* 90, 269–287.
- 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. (2023). 3D in high-D. *Journal of Philosophy*, forthcoming; preprint at https://tedsider.org/papers/3D_in_high-D.pdf.
- Sklar, L. (1993). *Physics and Chance: Philosophical Issues in the Foundations of Statistical Mechanics*. Cambridge: Cambridge University Press.
- Wallace, D. (2003). Everett and Structure. *Studies in the History and Philosophy of Modern Physics* 34, 87–105.
- Wallace, D. (2010). Decoherence and ontology: Or: How I learned to stop worrying and love FAPP. In S. Saunders, J. Barrett, A. Kent, and D. Wallace (Eds.), *Many Worlds? Everett, Quantum Theory, and Reality*, pp. 53–72. Oxford: Oxford University Press.
- Wallace, D. (2012). *The Emergent Multiverse: Quantum Theory according to the Everett Interpretation*. Oxford: Oxford University Press.

- Wallace, D. (2013). The arrow of time in physics. In H. Dyke and A. Bardon (Eds.), *A Companion to the Philosophy of Time*. Chichester: John Wiley and Sons.
- Wallace, D. (2016). Probability and irreversibility in modern statistical mechanics: Classical and quantum. To appear in D. Bedingham, O. Maroney and C. Timpson (eds.), *Quantum Foundations of Statistical Mechanics* (Oxford University Press, forthcoming). Preprint at <https://arxiv.org/abs/2104.11223>.
- Wallace, D. (2022). Stating structural realism: mathematics-first approaches to physics and metaphysics. In J. Hawthorne (Ed.), *Philosophical Perspectives Volume 36: Metaphysics*, pp. 345–378. Wiley-Blackwell.
- Wallace, D. (2023). Thermodynamics with and without irreversibility. To appear in Olimpia Lombardi and Cristian Lopez (eds.), *The Arrow of Time: From Local Systems to the Whole Universe* (Cambridge University Press, forthcoming). Preprint at <https://philsci-archive.pitt.edu/22273/>.
- Wallace, D. (2024). Learning to represent: Mathematics-first accounts of representation and their relation to natural language. Preprint: <https://philsci-archive.pitt.edu/23224/>.
- Wallace, D. and C. Timpson (2010). Quantum mechanics on spacetime I: Spacetime state realism. *British Journal for the Philosophy of Science* 61, 697–727.
- Worrall, J. (1989). Structural realism: the best of both worlds? *Dialectica* 43, 99–124.