Symmetries and the philosophy of language

Neil Dewar *

August 18, 2014

Faculty of Philosophy, University of Oxford

*email: neil.dewar@philosophy.ox.ac.uk; telephone: +447799762727

Abstract

In this paper, I consider the role of exact symmetries in theories of physics, working throughout with the example of gravitation set in Newtonian spacetime. First, I spend some time setting up a means of thinking about symmetries in this context; second, I consider arguments from the seeming undetectability of absolute velocities to an anti-realism about velocities; and finally, I claim that the structure of the theory licences (and perhaps requires) us to interpret models which differ only with regards to the absolute velocities of objects as depicting the same physical state of affairs. In defending this last claim, I consider how ideas and resources from the philosophy of language may usefully be brought to bear on this topic.

1 Introduction

The question of how the symmetries of a theory bear upon its representational content has been a subject of much recent discussion.[1] It is the contention of this paper that models related by a symmetry transformation are merely different ways of representing the same physical state of affairs (at least, with respect to qualitative properties); and that utilising resources from the philosophy of language provides an insightful way of defending this claim. This is both because it can provide us with new arguments to motivate such an interpretational stance, and because it can illuminate what structural features of such pairs of models make this interpretational stance permissible.

The structure of the paper is as follows. In section 2, I outline the theory that will be our worked example throughout the paper: namely, that of Newtonian gravity set in (full) Newtonian spacetime,[2] or "Newtonian gravitation" (NG) for short. In section 3, I outline some apparatus for approaching the symmetries of this theory (an apparatus which should generalise to other similar theories); and in section 4, I discuss how models related by different kinds of symmetry relate to one another. With this much setting-up done, I turn in section 5 to consider why models related by boosts should be taken to represent observationally identical states of affairs, and why this licences the dismissal of absolute velocities from our ontology. In section 6, I go on to argue—against the received wisdom—that we can implement this dismissal without altering our theory, i.e., merely by making acceptable interpretational stipulations regarding the theory. In section 7, I discuss the situations in which such an interpretational strategy would be advantageous. Finally, in section 8, I consider whether such an interpretational stance may in fact be not merely acceptable, but positively required—at least if an unpleasant indeterminacy of reference is to be avoided. It is in sections 6 and 8, in particular, that we will see how ideas from the philosophy of language (specifically, ideas regarding synonymy and translation) may be usefully borrowed for the purposes of philosophy of physics.

Two final remarks before we begin. I will approach NG via its models: that is, by specifying what kinds of mathematical structures will count as *kinematically possible models*, and then picking out a subset of those as *dynamically possible models*. The kinematically possible models are, roughly speaking, objects of the right mathematical type to represent a physically possible world; one can think of them as representing the metaphysically possible worlds. The dynamically possible models are then those which do, in fact, represent physically possible worlds. In this paper, I will assume that all dynamically possible models represent physically possible worlds. (This is not a truism: some views impose metaphysical constraints which mean that some dynamically possible models represent metaphysically impossible, and *a fortiori* physically impossible, worlds.)[3]

---

[1] See, for example, [Saunders, 2003b], [Dasgupta, 2009], or [Belot, 2013].

[2] That is, Newtonian rather than neo-Newtonian (aka Galilean) spacetime.

[3] Such as Maudlin's metrical essentialism [Maudlin, 1988].

2 Newtonian gravitation

Consider the theory of Newtonian gravitation, set in (full) Newtonian spacetime. The kinematically possible models of this theory are of the form $\langle \mathcal{A}, \mathcal{T}, \phi, \{\langle \mathbf{x}_i, m_i \rangle\}\rangle$, where

- $\mathcal{A}$ is a three-dimensional (Euclidean) metric affine space (that is, a set of points X equipped with a faithful, transitive action of some (Euclidean) normed vector space V on X);

- $\mathcal{T}$ is a simply connected, one-dimensional manifold equipped with a metric and orientation;

- ϕ is a scalar field on $\mathcal{A} \times \mathcal{T}$; and

- $\{\langle \mathbf{x}_i, m_i \rangle\}_{i \in I}$ is a set of ordered pairs, each consisting of a smooth function $\mathbf{x}_i$: $\mathcal{T} \to \mathcal{A}$ and a scalar $m_i \in \mathbb{R}$ (with I just being an index set)

In such a model, $\mathcal{A}$ represents absolute space and $\mathcal{T}$ represents absolute time. I will refer to the product space $\mathcal{A} \times \mathcal{T}$ as the (Newtonian) *spacetime structure*, and abbreviate it as N. The terminology here is a little unhappy: the structure presented is referred to by Penrose[4] as *Aristotelian spacetime*, which—as we shall come to later—is used in the philosophy of physics literature to mean something else. So instead, I follow Saunders[5] in calling it Newtonian spacetime. This also accords with the influential terminology of Friedman[6] and Earman[7]—at least, insofar as what they and I call "Newtonian spacetime" match in their structural essentials. However, there are some differences in the manner of construction, which raise some issues: these are discussed further in section 7.

ϕ and $\{\langle \mathbf{x}_i, m_i \rangle\}_{i \in I}$ represent, respectively, the gravitational potential and the gravitating particles (with the ith particle having mass m_i and trajectory $\mathbf{x}_i$); I will refer to these structures as the *dynamical structure*, and abbreviate them as P. Given that both $\mathcal{A}$ and $\mathcal{T}$ are equipped with metrics, it is straightforward to define the velocity $\dot{\mathbf{x}}_i$ and acceleration $\ddot{\mathbf{x}}_i$ of a particle, and the gradient $\nabla \phi$ and Laplacian $\nabla^2 \phi$ of the potential. In order for a kinematically possible model $\mathcal{M} = \langle N, P \rangle$ to be a dynamically possible model, the dynamical structure P must satisfy the following equations for any $\mathbf{x} \in \mathcal{A}$ and $t \in \mathcal{T}$:

$$\ddot{\mathbf{x}}_i(t) = -\nabla \varphi(\mathbf{x}_i, t) \tag{1a}$$

$$\nabla^2 \varphi(\mathbf{x}, t) = 4\pi G \sum_i m_i \delta(\mathbf{x} - \mathbf{x}_i) \tag{1b}$$

---

[4][Penrose, 2004, chap. 17]
[5][Saunders, 2013]
[6][Friedman, 1983]
[7][Earman, 1989]

3 The symmetries of Newtonian gravitation

The symmetries of the theory sketched above are, of course, well-known; as is the fact that they come in two important classes. The *spacetime symmetries* of a given Newtonian spacetime N are the automorphisms of N. Since N is a product space $\mathcal{A} \times \mathcal{T}$, an automorphism f of N will map $(\mathbf{x}, t) \mapsto (f_{\mathcal{A}}(\mathbf{x}), f_{\mathcal{T}}(t))$, where $f_{\mathcal{A}}$ and $f_{\mathcal{T}}$ are automorphisms (i.e., isometries) of $\mathcal{A}$ and $\mathcal{T}$ respectively. Such isometries consist of translations, reflections and (in the case of $\mathcal{A}$) rotations.[8]

In this essay, I will be concerned with only the continuous symmetries.[9] Thus, the spacetime symmetries we are interested in consist of the translations (both temporal and spatial) and spatial rotations. The set of such symmetries for one Newtonian spacetime are referred to as the *Newton group* for that spacetime. The Newton group of any Newtonian spacetime is isomorphic to that for any other; consequently, we can think of them as faithful representations of a single abstract Newton group.

Towards introducing the second kind of symmetry, note that if we apply a member of the Newton group (or one of the discrete spacetime symmetries) to the *dynamical* structure of a model, we always obtain an isomorphic model. That is, given a kinematically possible model $\langle N, P \rangle$, define the image of $P = \langle \phi, \{\langle m_i, \mathbf{x}_i \rangle\} \rangle$ under a map $d : N \to N$ as $d[P] = \langle \phi', \{\langle m_i, \mathbf{x}'_i \rangle\} \rangle$, where

$$\phi'(\mathbf{x}, t) = \phi(d_{\mathcal{A}}^{-1}(\mathbf{x}), d_{\mathcal{T}}^{-1}(t)) \tag{2a}$$

$$\mathbf{x}'_i(t) = d_{\mathcal{A}}(\mathbf{x}_i(d_{\mathcal{T}}^{-1}(t))) \tag{2b}$$

If d is a member of the Newton group, then $\langle N, d[P] \rangle = \langle d[N], d[P] \rangle$, and so is isomorphic to $\langle N, P \rangle$. Because the conditions picking out dynamically possible models of NG are purely structural, they apply to any given model $\mathcal{M}$ if and only if they also apply to any model $\mathcal{M}'$ which is isomorphic to $\mathcal{M}$.

Thus, for any kinematically possible model $\langle N, P \rangle$ and any member d of the Newton group for N, $\langle N, P \rangle$ is dynamically possible if and only if $\langle N, d[P] \rangle$ is. Let us introduce a little terminology, and say that two models $\mathcal{M}$ and $\mathcal{M}'$ are *co-dynamical* if either both or neither are dynamically possible. We can then say that applying any member of the Newton group to a model yields a co-dynamical model. However, the Newton group is not the only set of continuous transformations with this kind of feature: the dynamical (im)possibility of any kinematically possible model is preserved under *boosts*.

A boost in N is specified by any vector v from the vector space V underlying $\mathcal{A}$ and any time $t_0 \in \mathcal{T}$; the associated boost is then $b : (\mathbf{x}, t) \in N \mapsto (\mathbf{x} + \tau\mathbf{v}, t) \in N$, where τ is the oriented[10] distance between t_0 and t. If we add such boosts to the

---

[8]Strictly, $\mathcal{T}$ is also invariant under rotations; but because it is one-dimensional, the only "rotation" is the identity map.

[9]See [Pooley, 2003] and [Huggett, 2003] for discussion of some of the issues raised by spatial reflection symmetries (and by violations of those symmetries).

[10]That is, if t_0 precedes t with respect to the orientation of $\mathcal{T}$, then $\tau > 0$, and if t precedes t_0, then $\tau < 0$.

Newton group for N, we obtain what is known as the *Galilei group* for N; as before, we will also refer to an abstract Galilei group, of which the Galilei group for any given spacetime is a faithful representation. Boosts are *not* spacetime symmetries: a boost cannot be decomposed into a (single) automorphism of $\mathcal{A}$ and an automorphism of $\mathcal{T}$. Nevertheless, it is easy enough to check that $\langle N, P \rangle$ satisfies the equations (2) if and only if $\langle N, b[P] \rangle$ does so, for any boost b. As such, a boost—like a member of the Newton group—is a *dynamical symmetry*, a mapping whose application to (just) the dynamical structure of a kinematically possible model yields a co-dynamical model. It follows that each member of the Galilei group for N is a dynamical symmetry.

The Galilei group is not the full set of dynamical symmetries for NG: in addition to the discrete symmetries, one can apply time-dependent accelerations which are accompanied by appropriate alterations to the gravitational potential.[11] However, all the members of the Galilei group are dynamical symmetries, even though some of them (*viz.*, the boosts) are not spacetime symmetries—which is all that is required for our purposes. In particular, this has the consequence that two models related by a boost will not, in general, be isomorphic to one another—and so the fact that they are co-dynamical cannot be explained by the mere fact that the equations of NG are sensitive to the details of a model only up to isomorphism.

Note, however, that a dynamical symmetry can only relate two models which are constructed on (numerically) identical spacetimes. This is a consequence of the fact that we have chosen not to restrict our attention to a single such spacetime.[12] The reason for this is that if all models share a spacetime, then between any two models there is an immediately privileged map between their spacetimes: namely, the identity. However, it is not generally the case that this privilege is of physical significance—so keeping models with different spacetimes in play is a good way of ensuring we do not accidentally imbue it with such significance.

This means that we can say a little more about the relationship between two models $\mathcal{M} = \langle N, P \rangle$ and $\mathcal{M}' = \langle N', P' \rangle$ (where N and N', though isomorphic, have not been identified with one another) that are, in some sense, related to one another in the same way that $\langle N, P \rangle$ and $\langle N, d[P] \rangle$ are. It obviously can't be simply that applying a member of the Galilei group to $\mathcal{M}$'s dynamical structure yields $\mathcal{M}'$, since $N \neq N'$. Nor can it be that $\mathcal{M}$ and $\mathcal{M}'$ are isomorphic, since $\langle N, P \rangle$ and $\langle N, d[P] \rangle$ are not (in general) isomorphic. Rather, the relevant relationship is that $\mathcal{M}$ and $\mathcal{M}'$ should be isomorphic *to within a Galilean transformation applied to (only) the dynamical structure*: that is, that there should be both a mapping $f : N \rightarrow N'$ such that $N' = f[N]$, and a member d of the Galilei group for $\langle N' \rangle$ such that $P' = (d \circ f)[P]$.[13] In such a case, we will say that the two models are *pseudomorphic*. (Note that if d is also a member of the Newton group, then the two models will be isomorphic as well.)

Note that one cannot pick out some unique member of the Galilei group which

---

[11] For further discussion of this symmetry, see [Saunders, 2003a] or [Knox, 2013].

[12] Unlike, for example, [Earman, 1989].

[13] Hence, this includes the relationship between $\langle N, P \rangle$ and $\langle N, d[P] \rangle$ as the special case where the isomorphism in question is just the identity.

bears, as it were, sole responsibility for the two models' being pseudomorphic to one another: if the two models are isomorphic under f to within a Galilean transformation d, then for any Newtonian transformation h, they will be isomorphic under $h \circ f$ to within the transformation $d \circ h^{-1}$. Of course, if we single out some particular mapping $f : N \to N'$, then we can say that *relative to that mapping* the two models are related by this or that member of the Galilei group; but this is all we can say. Note that this is not at all obvious if we focus our attention on models whose background spacetimes have been identified with one another. For in that case, it is tempting to say that a given pair of models are related by a member of the Galilei group in some more absolute sense, not realising that they are, in fact, related by that member of the Galilei group only relative to the choice of the identity map as the appropriate mapping.

4 Qualitative and haecceitistic differences

We turn now to the question of what kinds of differences there could be between the worlds depicted by two pseudomorphic models. As mentioned above, pairs of pseudomorphic models come in two importantly different varieties: isomorphic and non-isomorphic. We consider these in turn.

So first, suppose that we have two isomorphic models of NG: and in fact, suppose that the models are $\mathcal{M} = \langle N, P \rangle$ and $\mathcal{M}' = \langle N, d[P] \rangle$, where d is a spatial translation. It seems that for any given possible world, each model, considered by itself, is just as well- or ill-suited to represent that world as the other. However, when considered together, we might be tempted to offer the following gloss on how they represent compared to one another: whatever world $\mathcal{M}$ represents, $\mathcal{M}'$ represents a world in which all the material contents of the world represented by $\mathcal{M}$ have been translated through space. The reason, of course, is that the representation of space-time N is (literally) identical in the two models, which naturally encourages us to think that each point of N represents the same space-time point in both models: that is, that the mathematical identity of manifold points represents the transworld identity of spacetime points.

Whether this interpretational move is compulsory, forbidden or neither has been the subject of much discussion on the topic of the static shift; more generally, the issues regarding the interpretation of isomorphic models forms the core of the literature on the Hole Argument. For now, however, I want to emphasise that even supposing that the worlds represented by $\mathcal{M}$ and $\mathcal{M}'$ do differ in this way, they do not differ in any *qualitative* respect, but only over which space-time points play which qualitative roles.[14] For example, suppose that both models contain just one inertial trajectory.

---

[14]I will be assuming that we have a clear enough grip on the distinction between qualitative and non-qualitative properties (notwithstanding the difficulties of giving an uncontroversial characterisation of the distinction). The basic idea is that a property is non-qualitative just in case it depends in ineliminable fashion on some particular individual: for example, the properties of *being David Beckham*, *being five metres from David Beckham*, or *being the same weight as David Beckham* are

Then the only difference between the models is that in $\mathcal{M}$, it is *this* point of $\mathcal{A}$ which is the range of the sole trajectory in the model, whilst in $\mathcal{M}'$ it is *that* point of $\mathcal{A}$ which is the range of the sole trajectory in the model; and correspondingly, taking mathematical identity to represent transworld identity, $\mathcal{M}$ depicts a particle which occupies *this* point of absolute space, whilst $\mathcal{M}'$ depicts a particle occupying *that* point of absolute space. There is a term for the view that some possible worlds differ only over which objects play which role: *haecceitism*.[15] We will follow this usage, and (correspondingly) say that worlds which differ only over which objects in them play which role differ *merely haecceitistically*.

By way of contrast, now consider two models of NG related by a boost: say, by letting $\mathcal{M}$ be as before, and considering $\mathcal{M}'' = \langle N, f[P] \rangle$ where f is a boost. This time, the natural or obvious interpretation is that the two models depict *qualitatively* different worlds, with $\mathcal{M}$ representing a particle at absolute rest and $\mathcal{M}''$ depicting a particle moving with some non-zero absolute velocity. And note that this interpretative stance is independent of haecceitistic considerations about whether $\mathcal{M}$ and $\mathcal{M}''$ represent the same space-time points, and if so whether they differ over which point has which characteristics. Indeed, whether $\mathcal{M}$ and $\mathcal{M}''$ agree or disagree over which objects play which roles is not really well-defined, given that the qualitative differences between them mean that they disagree over what roles there are to be played.

Note that I have been assuming, for ease of exposition, that if two worlds differ at most haecceitistically, then they are qualitatively identical; and, conversely, that if two worlds are qualitatively identical, then they differ at most haecceitistically. For the sake of simplicity, I will continue to make this (reasonable) assumption. In general, then, we will suppose that isomorphic models are *qualitatively synonymous*: the worlds they represent differ at most over which objects play which roles. Non-isomorphic models, however, represent qualitatively different worlds—or at least, at first pass they do so.

In this paper, I focus my attention on models that are pseudomorphic but not isomorphic. In the next section, I review an argument to the effect that the differences between such models will be empirically undetectable (even in principle) and consider the ramifications for believing that there are such things as absolute velocities. Then, in section 6, I argue that rather than being compelled to change our theory as a result, we can instead interpret these models as representing qualitatively identical worlds after all. In section 7, I consider the positive case for doing so in more detail. And finally, in section 8, I argue (somewhat tentatively) that a failure to interpret these models in this way will lead us into indeterminacy of reference.

A small point about presentation before we begin: I will speak sometimes of the observational equivalence of boosted worlds, and sometimes of the undetectability of absolute velocity. I don't believe anything of significance hinges upon this, and it

---

all non-qualitative properties. Qualitative properties, by contrast, are those which do not depend on a particular individual in this fashion: for example, *being green* or *being the same weight as some footballer*.

[15]Note that this is only one of several uses of the term "haecceitism"; this usage of the term matches that of [Melia, 1999] and [Pooley, 2002].

will aid comprehensibility to have both vocabularies available.

5 Undetectability

First, then, we are seeking to argue that two models related merely by a boost will be observationally indistinguishable. The literature contains a general argument to this effect: I will outline the argument only briefly, and refrain from discussing the subtleties that arise in a more thorough treatment.[16] So suppose that NG were true, and that absolute velocity could be detected; for the purposes of simplicity, suppose that it could be detected with absolute reliability. That means that in any physically possible world in which the detection procedure is performed on some object O, the output of the procedure is some display of O's absolute velocity. So let W be a world in which such a procedure is performed, with the corresponding result. But now consider a world, W', which is identical to W save that a uniform boost is applied to all of W's material contents. By the boost-symmetry of NG, W' is physically possible. So if the detection-procedure is performed in W', then the output of the procedure will be different to its output in W. But all the things in which performing a detection-procedure might plausibly consist—putting O next to some apparatus, pressing buttons, making sure the apparatus's parts are all present and correct—will be invariant under the boost. So the detection-procedure *is* performed, and its output *is* different to its output in W. Yet all the things in which that output might plausibly consist—ink on paper, semaphoric flags, the emission of a particular tone—are similarly invariant under a boost. So in fact, the output of the procedure is the *same* in W' as it is in W. By contradiction, therefore, no such detection procedure can exist.

The above is, as promised, a sketch rather than a full argument; in particular, our "plausible" assumptions that the performance and output of the detection-procedure are boost-invariant raise interesting subtleties when examined in more depth. But as the argument has been discussed extensively elsewhere, we will merely suppose that these can be ironed out satisfactorily. What would follow from the undetectability of absolute velocity?

Well, one venerable conclusion is as follows: we should not believe that there is any such thing as absolute velocity. The basic idea is fairly straightforward: if the facts about absolute velocity can have no impact on anything observable by us, then claiming that there are such facts goes well beyond our scientific-realist warrant for believing that there are facts of a certain sort; consequently, we should refrain from believing that there are such facts.

There are a couple of things worth noting about this argument. The first is to stress its disanalogy with traditional positivist or instrumentalist arguments against some domain or other of unobservable facts. In part, this is simply because our conception of "observable" can be significantly more relaxed than the conception typically

---

[16]The argument given in this section is due to [Roberts, 2008]; besides the treatment there, it is also discussed by [Russell, 2011, §3.4], and in several papers by Dasgupta (see e.g. [Dasgupta, 2011], [Dasgupta, 2013], and [Dasgupta, 2014]).

found there: not only can we not perceive absolute velocity directly, but (partly as a consequence) we cannot build any kind of device which might succeed in reliably correlating absolute velocities with perceptible outputs. More importantly, though, it is because rather than failing to meet some externally imposed standard of observability, absolute velocity is deemed unobservable by the *very same theory of which it is a part*: namely, NG (interpreted literally). The only warrant that we might have for believing in absolute velocity flows from believing in NG; but believing NG (together with certain plausible beliefs about our mental and perceptual faculties) would entail believing that absolute velocity was in principle empirically inaccessible, and so beyond the scope of our realist warrant.[17]

The other point I wish to make is to contrast this case with the case of isomorphic models. As discussed in section 4, the differences between the worlds represented by isomorphic models are at most haecceitistic, whereas non-isomorphic models, interpreted literally, postulate qualitative differences. As a consequence, there is a difference in the epistemic cost of the situation that results. This difference has received a fair amount of attention recently, so it will be worth taking a moment to explain it. The starting point is an observation from Maudlin that there is an important disanalogy between the static and kinematic shift arguments:

> Consider first the kinematic shift. [...] The universe as a whole may be at rest, or travelling uniformly five meters per second due north, or 888 meters per second in the direction between Earth and Betelgeuse, and so on. According to Newtonian dynamics no possible observation can reveal its actual state of motion. So Newton must postulate that the universe has a physically real but empirically inaccessible property. In this sense, states which differ only in their net absolute velocity are observationally indistinguishable.
>
> What if we try to make the same argument using the static shift? Various positional states of the universe as a whole are possible: It could be created so my desk is here, or three meters north of here, or 888 meters from here in the direction from Earth to Betelgeuse, and so on. Which is the actual state of the world? Now the answer is easy: In its actual state, my desk is here, not three meters north or anywhere else. So in the kinematic case, unlike the static case, sensible physical questions can be asked but cannot be answered by observations. To even formulate the appropriate question in the static case one must indexically pick out a spatiotemporal location, and it is then no great trick to observe what material body that location actually contains.[18]

Before we turn to the argument itself, we should emphasise that what is being contrasted here is not the static and kinematic shifts *per se*, but rather the status of those

---

[17]This point is made by [Friedman, 1983, p. 219] and [Healey, 1999].
[18][Maudlin, 1993, p. 190]

shifts *in NG* (or any theory set against Newtonian spacetime as a backdrop);[19] in our terminology, the contrast is between isomorphisms, and mere pseudomorphisms (i.e. pseudomorphisms which are not isomorphisms). Because merely pseudomorphic models (interpreted literally) describe worlds that differ qualitatively, the character of each can be described in a language that avoids making reference to particulars; as such, we can then ask questions about whether (for example) the actual world is best described by this model or that model. By contrast, the haecceitist about isomorphic models won't think that we can indicate that the points of a model are supposed to represent the points of the actual world, save by a stipulation that this is so; this is why being in a position to ask whether a model correctly represents the haecceitistic facts about the actual world entails knowing the answer.

This is perhaps clearest if we try extending Maudlin's observation to Aristotelian or Galilean space-time. In Aristotelian space-time, we add further structure: we pick out a privileged point C of $\mathcal{A}$, corresponding (intuitively) to the centre of the universe.[20] As such, static shifts are no longer isomorphisms: two Aristotelian worlds related by a static shift *would* differ qualitatively, since there would, in each case, be some fact about (for example) how far the universe's centre of mass is from its absolute centre.[21]

Galilean space-time, on the other hand, has substantially *less* structure than Newtonian space-time. In our terms, Galilean space-time is best thought of as a fibre bundle E where $\mathcal{A}$ is the prototypical fibre and $\mathcal{T}$ the base manifold; there is a privileged affine connection D on E (i.e. a connection on the tangent bundle TE), but no bundle connection on E itself (since a bundle connection provides a privileged mapping between the fibres at two points, which is more structure than is desired).[22] As a result, *neither* the static nor kinematic shifts introduce qualitative differences: the only difference between two kinematically shifted Galilean worlds is with regards to, for example, which spacetime points the universe's centre of mass successively occupies. Hence, we are unable to ask for a number which specifies the speed with which the universe as a whole (or any part thereof) is travelling; the best we can do is to ask, for example, whether my desk is travelling as fast as my desk or not.[23]

Hence, if Maudlin is correct, then there is an epistemic problem associated with our absolute velocity (if we interpret pseudomorphic models literally), but not one

---

[19]This point is noted by Maudlin in the paper (see [Maudlin, 1993, p. 192]), although it comes out a little more clearly in his discussion in [Maudlin, 2012, chap. 2–3].

[20]Note that is *not* what Penrose calls "Aristotelian space-time" (as per the remarks in section 2); rather, it corresponds to the usage of [Earman, 1989, pp. 34–35].

[21]If you're tempted to say that any plausible physics in Aristotelian spacetime should locate the centre of the universe at the centre of mass of the universe, then you should be equally tempted to say that any plausible physics in Newtonian spacetime should judge the absolute velocity of the universe's centre of mass to be zero.

[22]See [Penrose, 2004, chap. 17]; for alternative ways of getting at essentially the same structure, see [Earman, 1989] or [Saunders, 2013] (although see Section 7 for discussion of what might be at stake in such choices of presentation). Here, the terminology of Penrose, Earman and Saunders agrees (although Saunders also uses—and prefers—the label "neo-Newtonian" to "Galilean").

[23]The point that in Galilean spacetime, the kinematic shift reduces to a merely haecceitistic difference is stressed by [Pooley, 2002, p. 59].

associated with our absolute position (even if we interpret isomorphic models haecceitistically). To coin some terminology, let us say that the kind of ignorance Maudlin is concerned with is *qualitative ignorance*: it is the kind of ignorance I have when I do not know the qualitative character of the world (as when I interpret NG literally). However, Maudlin's conclusion from this that "the static shift does not result in an indistinguishable state of affairs, nor does it imply that there are any real but empirically undeterminable spatiotemporal facts about the world"[24] has recently been challenged: both Dasgupta and Russell articulate accounts of detection or observation according to which our position in absolute space *is* undetectable or unobservable after all (although, in light of the above, we are unable to express what it is that we do not know).[25] In effect, these analyses articulate a notion of what we might call *haecceitistic ignorance*: it is the kind of ignorance I have when I do not know which objects are playing which roles (even if I know the qualitative character of the world).

Now, there's certainly an interesting question about whether haecceitistic ignorance really is a substantial kind of ignorance. The point I wish to make here, though, is a simpler one: that even if it is, it is still worse to be qualitatively ignorant than to be merely haecceitistically ignorant; and as such, the epistemic problems arising from interpreting pseudomorphic models literally are still worse than the problems arising from interpreting isomorphic models haecceitistically. The reasoning is pretty straightforward. If I am qualitatively ignorant, then I am ignorant of the qualitative character of the world; hence, I am ignorant of what qualitative roles are being played; I therefore *cannot* be in a position to know the haecceitistic character of the world, i.e., to know which individuals are playing which qualitative roles. Note that as a result, the epistemic case against reading pseudomorphic models literally can be accepted even by someone who rejects anti-haecceitism about isomorphic models.

6 In defence of qualified realism

So suppose that we agree, having taken on board considerations like those above, that if we like NG, we ought not to like absolute velocity. What's our next move?

One common response urges that our next move has to be to look for a theory which lacks commitment to absolute velocity, which we can put forward in place of NG. There's an obvious alternative: *Galilean* Newtonian gravity, i.e., Newtonian gravity set in Galilean spacetime (models of which, recall, lack any standard of absolute rest and therefore contain boosts as a spacetime symmetry). So, for example, Pooley writes that

> The sophisticated substantivalist's way to evade both kinematic shift arguments is to [. . .] move from Newtonian to Galilean spacetime. The difference between the possibilities represented by [a pair of boosted models set on the same Galilean spacetime] (on the stipulation that each point of

---

[24][Maudlin, 1993, p. 191]
[25]See [Dasgupta, 2009, §2.4] and [Russell, 2011, §3.5].

[the background manifold] represents the 'same' spacetime point in both models) *is* merely haecceitistic. He can therefore reinterpret [one of the models] and see [the pair of models] as representing the same possible world.[26]

Similarly, Dasgupta argues that absent such an alternative theory, one is rationally compelled to accept absolute velocities:

Note, then, that if this how symmetry-to-reality inferences work then we can only draw the conclusion of the inference when we have the alternative theory in hand and have shown that all else is equal. This explains why it was rational for Newton to believe in absolute velocity even though he knew that it was variant in NG and undetectable. The reason this was rational for him was that he had no good alternative theory to hand. He had good reason (his "bucket argument") to think that relationalism was not empirically adequate. And relationalism was the only alternative view he knew of (he was not aware of Galilean space-time structures in which there is a well defined feature of absolute acceleration (as required by his bucket argument) but no absolute velocity). So for Newton, all else was not equal and he was rational to believe in absolute velocity.[27]

However, it's not clear to me that this is really necessary. In this section, I want to look at whether there's an alternative option for the realist: to refrain from believing that all the structures in the models of NG correspond to structures in physical reality. That is, can we put forward NG as our theory, whilst interpreting it in such a way that we avoid commitment to absolute velocity? Let us call such an interpretation a *qualified realism* about NG.[28]

The first task is to articulate what such a position would look like. Here's a first-pass suggestion: we should claim that the cross-time identification of points in $\mathcal{A} \times \mathcal{T}$ which share a first member expresses nothing physically significant. Unfortunately, however, such a suggestion is inadequate, since this would mean that we no longer affirmed the structure needed to underwrite the differentiation between the dynamically possible and dynamically impossible models. It is straightforward enough to verify that a mere fibre bundle of $\mathcal{A}$ over $\mathcal{T}$[29] is unable to determine a standard of rotation, and therefore cannot be used to capture Newtonian dynamics (in which rotation is dynamically significant). So clearly absolute space is of *some* physical significance, which we risk eradicating if we are too zealous from the off.

A second-pass suggestion would be that we ought, then, to reformulate the theory so that we are in a position to make explicit which of its structures we will take to represent genuine physical structures: that is, when we have Galilean Newtonian gravity

---

[26][Pooley, 2002, p. 59]; emphasis in original.

[27][Dasgupta, 2014, p. 17]

[28][Saunders, 2013] can be seen as engaged in a similar project to that discussed here, although with a stronger focus on Newtonian gravity as presented in the *Principia*; moreover, the main focus of his discussion is the gravitational gauge symmetry rather than the boost symmetry.

[29]Which would correspond to so-called *Leibnizian spacetime* ([Earman, 1989, pp. 30–31]).

to hand, and can then affirm that the only structures of NG about which we are re-
alists are those which have a counterpart in Galilean Newtonian gravity. Although
this is now a coherent position, it's still problematic. This approach means that we
can only take a qualified-realist stance towards a theory when we have an alternative
theory which makes manifest which parts of the first theory are surplus structure. But
if we have such an alternative theory in hand, then qualified realism doesn't seem like
a particularly interesting position. The extra structure is made manifestly redundant
and avowedly meaningless; putting forward models containing such a structure is
a little like insisting on shouting "yabba-dabba-doo!" at the end of each sentence of
the theory, whilst reassuring everyone that this is merely a nonsensical ejaculation.
Behaving in this fashion isn't wrong, *per se*, so much as just—at best—pointless and
rather tiresome. And in at least some cases, we just won't *have* such a formulation: in
the case of NG, for example, the mathematics necessary to exhibit its surplus structure
clearly was not developed until centuries after the theory was proposed.

A more interesting thought, then, would be to ask whether there is some way in
which we could be anti-realist about part of a model without being required to ex-
plicitly single out the parts of the model one is anti-realist about. I think the answer
is yes. The trick is to stipulate which models are synonymous, rather than specifying
which bits of a model one reads literally or not: we express our qualified-realist atti-
tude by affirming certain non-isomorphic models as synonymous, which commits us
to denying that the respects in which such models disagree correspond to any physi-
cally significant difference. And of course, we know which models we'll be affirming
as synonymous: the models which are pseudomorphic to one another. Hence, we
conjoin our mathematical theory NG to the semantic stipulation that pseudomorphic
models are synonymous, with the consequence that only whatever structure such
models have in common is to be interpreted realistically. In our specific case, this
means that the physical ontology posited by the theory (so interpreted) is the same
as that posited by Galilean Newtonian gravity (straightforwardly or literally inter-
preted). The easiest way of elaborating this position is to consider, in turn, the objec-
tions that might be raised against it.

First objection: to do this is to renounce scientific realism. Response: scientific re-
alism surely isn't the position that given an empirically adequate theory, one is duty-
bound to believe in every piece of structure postulated by that theory, regardless of
what the theory should say about that structure's explanatory or theoretical role. If
anything, the kind of interpretative stance under discussion is *more* consonant with
a realist attitude than a blind literalism. For realism is a matter of taking seriously
what the theory has to say; and if the theory says that certain quantities are in prin-
ciple unobservable, then to take the theory seriously may come apart from taking it
literally.

Second objection: even if we grant that one may, in general, not be realist about
all aspects of a theory's structure, it is not acceptable to do so when those aspects
are indispensable; but the "interesting" cases above are precisely those in which the
surplus structure is indispensable. Response: indispensable for what? If they were
indispensable for explanation, then it does seem that they ought to be retained. But

if they're only indispensable for the purposes of characterising or representing the structures that are (explanatorily, hence physically) significant, then it seems that we have much greater leeway to regard them as mere mathematical scaffolding, without being honour-bound to accord them physical significance. If I say that "that rock looks like a dragon's head", then I have certainly made use of the notion of dragons to characterise or represent how the rock looks; perhaps even (if there are sufficiently many ways the rock could have looked, or if our language for expressing shape-predicates is sufficiently anaemic) an indispensable use. But the acceptability of such locutions hardly commits us to a belief in dragons. Similarly, the ubiquity of potentials as tools for calculation is not generally taken to necessitate a commitment to their existence.[30]

The relevant question, then, is when we have good grounds for thinking that some piece of mathematical apparatus in our theory is of merely representational or calculational significance. Of course, explicitly identifying the apparatus, and showing that a dynamically equivalent theory can be put forward that does without it, is one way of doing so. But the argument of section 5 did not rest upon doing so: all it required was that our theory include models which were pseudomorphic but not isomorphic.

Third objection: to put forward NG as your theory, and then deny the existence of absolute velocity, is just straightforwardly contradictory. Response: why? Yes, sometimes we think that someone is being inconsistent if they make use of the existence of something, and then deny that that thing exists: notably, if the thing figures in a proffered explanation. But in plenty of other cases, such "weaselling away" of apparent ontological commitments is perfectly acceptable: for example, no-one would think me inconsistent for saying that "that rock looks like a dragon—but of course, there are no dragons". Such locutions are an acceptable aspect of explaining how one's theory is to be understood (i.e., of specifying an interpretation), by indicating which parts are to be understood as mere aids to representation and (so) not interpreted literally.[31]

Fourth (and final) objection: we cannot just declare models to be synonymous as and when we please; there must be some constraints, or else we will end up collapsing all observationally equivalent models together. Response: I've already mentioned one important constraint—the requirement that our theory posit sufficient physical ontology to explain the observable phenomena that we see. The argument of section 5 gives us good reason to think that we will not cut into that ontology by dis-

---

[30]Ignoring concerns raised by the Aharonov-Bohm effect.

[31]Both this point and the point above draw heavily on the recent debate in the philosophy of mathematics literature over the so-called "indispensability argument". For further explanation of the kind of line taken here (that representational indispensability does not entail ontological commitment) see [Melia, 1995], [Melia, 2000], or [Yablo, 2005]; for critical responses, see e.g. [Azzouni, 2009] or [Colyvan, 2010]. It is, incidentally, worth noting that much of the recent discussion has accepted that ontological commitment requires *explanatory* indispensability, and focused instead on whether mathematics is explanatorily indispensable in this sense (see e.g. [Baker, 2009]—from whom the dragon example is taken—or [Leng, 2010, ch. 9]).

avowing commitment to structure which is variant under symmetry; declarations of synonymy that go beyond this, on the other hand, do risk doing so. However, we can do better than this. Suppose that we conceive of the task of interpreting a theory as being that of providing, for it, a Davidsonian theory of meaning.[32] In the context of a natural language L, this amounts to providing a theory that will entail all theorems of the form

$$s \text{ is true-in-}L \text{ if and only if } p, \tag{3}$$

where s is a placeholder for a structural description of a sentence and p a placeholder for a synonymous sentence. In our case, this will mean providing a specification of the co-representationality relations holding amongst the kinematically possible models.

So imagine, as it were, that we have been handed a great stock of all the sentences in question (respectively, all the kinematically possible models), and are seeking to organise them into theorems of the form (3) (respectively, to lay down putative relations of co-representationality). Without presupposing knowledge of the meanings of the sentences (models), are there constraints that we can impose on how such an organisation is to be done which could ensure that we get only true theorems (respectively, only correct relations of co-representationality)? Davidson suggests there are; and it turns out that these constraints (or analogues thereof) carry over to our case.

The first constraint is simply that in each instance of (3), the left- and right-hand-sides coincide in truth-value.[33] Even in the natural-language case, this is a very weak condition: it still permits manifestly deviant theorems such as

$$\text{'snow is white' is true-in-English if and only if grass is green.} \tag{4}$$

In our case, however, the condition is trivially satisfied: given that all models of NG are false (since NG is a false theory), anytime we connect two models we know that they will coincide in truth-value. Given that it is satisfied (albeit trivially), we could simply allow that this first constraint doesn't translate very well from the natural-language setting; however, it is tempting to ask if there is an alternative condition that will serve as an appropriate analogue in our case.

I suggest that a natural analogue to this first constraint is that two models are co-representational only if they are co-dynamical.[34] After all, although the presentation of a theory does not include information about which models are true and which false, it does at least include a bifurcation of the models into those which are dynamically possible and those which are dynamically impossible. Moreover, if we consider *why* Davidson's theory includes this first constraint (i.e., what role that constraint

---

[32]See e.g. [Davidson, 1965] and [Davidson, 1967].

[33]That said, there are some subtleties in the extent to which the truth-values, for Davidson, may be taken as given: at various places, Davidson avers that what is really being taken as basic is not facts about which sentences are true, but facts about which sentences are "held true" by which members of the linguistic community (the latter facts being, allegedly, accessible from purely behavioural evidence). See [Davidson, 2001, pp. 25–27], and [Glock, 2003, §8.1].

[34]I thank an anonymous referee for requesting a fuller explanation of this analogy.

plays), then at least one reading is that it is there to encode the contribution to meaning of the linguistic community's background knowledge.[35] The background knowledge required to correctly interpret a physical theory is knowledge of the laws of the theory (in the minimal sense of knowing which models are dynamically possible); hence, imposing a constraint that co-representational models are co-dynamical serves to encode the salient background knowledge. The satisfaction of this constraint by pseudomorphic models is, of course, entailed by the definition of pseudomorphism.

The second important constraint is that a specification of the meanings of sentences must also constitute a specification of the meanings of words, in such a way that if synonymous words are put together in the same way, we obtain a synonymous sentence. The analogue for us is that our specification of co-representationality must enable a specification of the relations of co-representationality on the *structures*, that says under what circumstances two structures represent things the same way, and under what circumstances two collections of structures are to count as having been put together the same way. The force of the constraint comes out as follows: an incorrect specification of the synonymy relations amongst words will lead to the synonymy relation amongst sentences pairing some truths with falsehoods, in violation of the first constraint. Thus, although (4) abides by the first constraint, the sentences obtained by rendering "snow" synonymous with "grass", and "white" with "green" (say, that 'snow is white' is true if and only if snow is green) will not.

The point to appreciate is that given this, limiting the co-representationality relation to pseudomorphic models can be seen as a means of abiding by this constraint. For only this will enable us to count the relation that holds between the absolute rest field and a parallel worldline as synonymous with the relation holding between the absolute rest field and a worldline at some angle to it (in different models), without thereby counting as co-representational any models which are not co-dynamical. If, on the other hand, we tried to extend the co-representationality claim beyond pseudomorphism, then we would not expect to find any consistent way of relating the *internal* structures of each model to one another in a manner coordinated with the co-representationality (at least, provided that co-representational models are required to be co-dynamical).

For example, consider a pair of models $\langle N, P \rangle$ and $\langle N, g[P] \rangle$ where g is just some diffeomorphism on N that is not in the Galilei group. By picking an appropriate such g and P, one can set things up so that $\langle N, P \rangle$ and $\langle N, g[P] \rangle$ are co-dynamical: for example, suppose that the former describes one particle in circular orbit around another, the latter describes one particle in elliptical such orbit around another (whose respective masses are the same), and g simply distorts the circle into an ellipse. Although declaring $\langle N, P \rangle$ and $\langle N, g[P] \rangle$ to be co-representational would abide by the first constraint, it would violate the second: an arbitrary dynamically possible model $\langle N, P' \rangle$ will *not* generally be such that $\langle N, g[P'] \rangle$ is dynamically possible.

In most cases, of course, an arbitrary pair of co-dynamical models will not even be

---

[35]This would provide an alternative justification for focusing on data regarding what the community holds true, rather than data regarding what *is* true (see footnote 33).

such as to allow *any* mapping of the structures of one onto those of the other: for example, if the two models have different numbers of particles, or particles of differing masses. Hence, the second constraint also prevents us identifying the representational content of these models; this is analogous to the fact that the second constraint, in the natural-language case, will prohibit taking sentences with radically different logico-grammatical forms to be synonymous.

7 In favour of qualified realism

So, we can defend qualified realism from objections. That, of course, hardly amounts to a positive case for it; in this section and the next, I undertake to provide such a case. Earlier, I mentioned that we may be in a position to know that a theory contains surplus structure, but not to clearly exhibit which aspects, precisely, are surplus. This section explores this issue further, by looking at different ways of presenting Newtonian spacetime. In particular, I compare and contrast the approach to defining Newtonian spacetime used in this paper (which I will refer to as the *Penrose* approach[36]) with the perhaps more familiar approach of, for example, Friedman[37] or Earman[38] (which I will refer to as the *standard* approach).[39]

Recall that the Penrose approach defines Newtonian spacetime as the product of a three-dimensional metric affine space $\mathcal{A}$ (representing absolute space) with a one-dimensional metric space $\mathcal{T}$ (representing absolute time). On the standard approach, by contrast, Newtonian spacetime is defined as a four-dimensional manifold $\mathcal{M}$ equipped with[40]

- A smooth, vanishing, curl-free one-form t_a on $\mathcal{M}$,

- A contravariant, symmetric tensor field h^{ab} on $\mathcal{M}$, of type $(2,0)$ and signature $(0,1,1,1)$,

- An affine connection D, and

- A vector field A^a on $\mathcal{M}$,

---

[36]Since this approach is expounded in [Penrose, 2004, chap. 17].

[37][Friedman, 1983, §III.1]

[38][Earman, 1989, chap. 2]

[39]I thank one anonymous referee for pushing me to make more clear the positive case for qualified realism, and another for suggesting that it might be enlightening to compare the significance of different formalisms.

[40]In the below, Latin indices *a*, *b*, etc., are abstract indices (see e.g. [Wald, 1984, §2.4]).

such that the following equations are satisfied:

$$h^{ab}t_b = 0 \tag{5a}$$
$$D_a t_b = 0 \tag{5b}$$
$$D_a h^{bc} = 0 \tag{5c}$$
$$K^a{}_{bcd} = 0 \tag{5d}$$
$$t_a A^a = 1 \tag{5e}$$
$$D_a A^b = 0 \tag{5f}$$

where $K^a{}_{bcd}$ is the curvature tensor of D. The integral surfaces of t_a represent spatial hypersurfaces; t_a also encodes the temporal metric and orientation, whilst h^{ab} represents the spatial metrical structure (by inducing a three-metric on each integral surface of t_a). Equation (5a) stipulates that the spatial and temporal metrics must be orthogonal. D represents the affine structure of the spacetime; it must be compatible with the metrical structure (equations (5b) and (5c)), and flat (equation (5d)). Finally, A^a encodes the structure of absolute space: the worldlines of the points of absolute space are the integral curves of A^a. Equation (5e) tells us that the worldline of any point of absolute space is (future-pointing) timelike along its length and parameterised by absolute time, whilst equation (5f) tells us that those worldlines remain a constant spatial distance from one another.

Thus, the Penrose approach begins by defining (the mathematical representatives of) space and time, and then defines Newtonian spacetime as a certain construction (namely, the Cartesian product) from these; in particular, this means that two spacetime points which are temporal stages of the same spatial point will be represented by two pairs with identical spatial factors. The standard presentation, by contrast, constructs Newtonian spacetime "all at once", by postulating structures of a certain sort upon a four-dimensional manifold; although, of course, those structures enable one to recover first a decomposition into space and time, and a representation of trans-temporal identity of spacetime points (namely, the integral curves of A^a).

Whether this amounts to a substantial difference is a rather delicate matter, touching on some subtle issues in the philosophy of mathematics. One might feel that on the Penrose approach, the cross-time persistence of points of absolute space is more tightly "keyed in" to the formalism than on the standard view, since (in a slogan) it represents identity by identity. However, we have to be somewhat careful; this claim becomes problematic if we are more structuralist about mathematics (or more weakly, if we think that only the structural aspects of mathematics are relevant for their purpose as representational devices).[41] We often describe a product $X \times Y$ as consisting of (all and only the) pairs of points $\langle x, y \rangle$ such that $x \in X$ and $y \in Y$; structurally, however, the role of $X \times Y$ is exhausted by the fact that it comes equipped with canonical projections $(X \times Y) \to X$ and $(X \times Y) \to Y$.[42] And the standard

---

[41] For further discussion of this latter idea, see [Weatherall, 2014].
[42] [Awodey, 2010, §2.4]

19

definition of Newtonian spacetime enables us to naturally define privileged maps $\pi_{\mathcal{S}} : \mathcal{M} \to \mathcal{S}$ and $\pi_{\mathcal{L}} : \mathcal{M} \to \mathcal{L}$, where $\mathcal{S}$ is an arbitrary integral surface of t_a and $\mathcal{L}$ an arbitrary integral curve of A^a: for any $m \in \mathcal{M}$, $\pi_{\mathcal{S}}$ takes m to the point of $\mathcal{S}$ lying on the same integral curve of A^a as m, and $\pi_{\mathcal{L}}$ takes it to the point of $\mathcal{L}$ lying on the same integral surface of t_a as m. In terms of these maps, Newtonian space-time also represents identity by identity: $m, n \in \mathcal{M}$ represent successive stages of the same spatial point just in case $\pi_{\mathcal{S}}(m) = \pi_{\mathcal{L}}(n)$. More generally, given that the structures $\mathcal{A} \times \mathcal{T}$ and $\langle \mathcal{M}, t_a, h^{ab}, D, A^a \rangle$ are definitionally equivalent (in that all the structure present in one may be uniquely defined in terms of the structure present in the other),[43] it seems most natural to say that the choice is that of how to present a given structure, rather than what structure to use.

However, the choice of presentation can make a difference. There are three differences that the use of the Penrose approach makes; two of these are fairly minor, but the third bears directly on the issues in this paper. First, the Penrose presentation is somewhat simpler, at least for our purposes. The standard presentation is more easily generalised to presenting relativistic spacetimes; however, essentially for this reason, it invokes a good deal more differential-geometric apparatus than is needed if one is working only with classical spacetimes. Second, the Penrose presentation is more faithful to the historical conception of space and time: insofar as Newton had an (implicit) grasp of the concept of spacetime, it was surely as something definitionally posterior to the notions of space and time (rather than the other way around).

The third difference is more interesting. In the above, I said that the advantage of qualified realism is that it helps us deal with cases in which we know that a theory contains surplus structure, but are not in a position to expunge that structure from the theory. If Newtonian gravitation has been presented to us in the Penrose manner, then (if we are unfamiliar with fibre bundles) we will typically be in such a position. If it is presented in the standard manner, however, then we will not: for in that case, it rapidly becomes obvious that we can still pose the dynamics of the theory using only t_a, h^{ab}, and D (and so A^a is identified as surplus to requirements).

It is worth pointing out that this is somewhat independent of the issue about whether space and time are defined prior to spacetime, or vice versa. For one can define Newtonian spacetime in a manner broadly analogous to that used by the standard approach, yet without making so manifest the redundancy of absolute space.[44] That is: rather than defining a Newtonian spacetime as a tuple $\langle \mathcal{M}, t_a, h^{ab}, D, A^a \rangle$ subject to equations (5), one defines it as a tuple $\langle \mathcal{M}, t_a, h^{ab}, A^a \rangle$ subject to the equations

$$h^{ab} t_b = 0 \tag{6a}$$

$$t_a A^a = 1 \tag{6b}$$

$$\pounds_A h^{ab} = 0 \tag{6c}$$

$$\mathcal{R}^a{}_{bcd} = 0 \tag{6d}$$

---

[43]See [Hodges, 1997, chap. 2].

[44]This construction is due to Oliver Pooley; see [Pooley, 2013, chap. 4].

where $\pounds_A$ is the Lie derivative with respect to the vector field A^a, and $\mathcal{R}^a_{bcd}$ is a Riemann curvature tensor encoding the curvature of the integral surfaces of t_a. The intuitive explanations of these equations are the same as those for (5), modulo the fact that equation (6c) is doing the work of equation (5f) and equation (6d) is doing the work of equation (5d).

These structures are sufficient to formulate the dynamical laws of Newtonian gravitation.[45] Rather than go into the details of how this is done, I merely observe that although defining (the analogue of) equation (1b) draws only upon t_a and h^{ab}, defining (the analogue of) equation (1a) makes crucial use of A^a: in effect, one uses A^a to define a notion of absolute velocity, with (absolute) accelerations then being defined as changes in absolute velocity. Hence, in this presentation—as in the Penrose presentation—the redundancy of absolute space is not made manifest, but must be teased out by non-trivial mathematical work.

Appreciating this helps us to see that the manifest redundancy of absolute space in the standard presentation is, in fact, quite a peculiar feature of that formalism (although its peculiarity is apt to be disguised by its familiarity). The point is simply that once one is given t_a, h^{ab} and A^a, there is a uniquely natural affine connection D definable from those ingredients (with the definition simply being satisfaction of equations (5b), (5c), and (5f)); hence, introducing D as an additional primitive is not well-motivated within the project of defining Newtonian spacetime.

Speaking generally, then, I suggest that the situation in which qualified realism offers a distinctive prescription (*viz.*, that of knowing that there is surplus structure around, but being unsure of where) is the generic case. As a result, we cannot simply cross out all the formulae featuring the surplus structure—that would throw the baby out with the bathwater. If we cannot tell the baby from the bathwater, however, then something like qualified realism is surely the only option.

8 Indeterminacy

Hence, an interpretation of the models in this fashion is consistent with a thoroughgoing realist attitude; and in many situations, such an interpretation may be the most appropriate realist attitude. In this final section, I wish to explore a more radical suggestion: that in fact, it may be the *only* appropriate attitude, if we restrict ourselves to intra-theoretical considerations for determining the representational content of a theory.

We noted above that boosted models are related to one another in a fashion reminiscent of synonymous sets of sentences, and suggested that that helped explain why interpreting them as synonymous (i.e., co-representational) made sense. By pursuing this thought a little further, we arrive at a caution against not doing so: namely, that we will be exposed to an indeterminacy of reference. The argument is an adaptation of Quine's more general argument, and runs as follows.

---

[45]See [Pooley, 2013, chap. 4].

First, consider some particular dynamically possible model of NG $\mathcal{M} = \langle N, P \rangle$, and consider its spacetime N; select some particular time $t \in \mathcal{T}$. Now consider the family of boosts $\{b_{\mathbf{v}}\}_{\mathbf{v} \in V}$, where $b_{\mathbf{v}}$ is the boost specified by t and $\mathbf{v}$ (in the sense defined above, in section 3), and define $\mathcal{M}_{\mathbf{v}} := \langle N, d_{\mathbf{v}}[P] \rangle$ (so that $\mathcal{M} = \mathcal{M}_{\mathbf{0}}$); hence, $\mathcal{M}_{\mathbf{v}}$ differs from $\mathcal{M}$ by a boost of $\mathbf{v}$. If we interpret NG literally, then we conclude that $\mathcal{M}_{\mathbf{v}}$ and $\mathcal{M}_{\mathbf{w}}$ represent distinct possible worlds $W_{\mathbf{v}}$ and $W_{\mathbf{w}}$, whose material contents have different absolute velocities from one another.

In other words, what we have is a putative correspondence between the mathematical models and the physically possible worlds; let's call this correspondence I_0. However, there is now a problem analogous to that faced by Quine's notorious field linguist: we can put forward a systematic reinterpretation of the models which will preserve the requirement that dynamically possible models represent physically possible worlds. For consider the correspondence $I_{\mathbf{a}}$, for some $\mathbf{a} \in V$, which puts $\mathcal{M}_{\mathbf{v}}$ into correspondence with $W_{\mathbf{v}+\mathbf{a}}$. What makes this interpretation scheme any worse than I_0? You might say: because we *stipulated* that $\mathcal{M}_{\mathbf{v}}$ was to represent the possible world $W_{\mathbf{v}}$. But all that does is push the problem back to the label "$W_{\mathbf{v}}$", whose connection to any given possible world is just as questionable as $\mathcal{M}_{\mathbf{v}}$'s. In effect, we are exploiting the fact that boosts are a symmetry to demonstrate that interpreting "absolute rest" as "absolute velocity of $\mathbf{a}$", etc., is an acceptable reinterpretation (acceptable, that is, in the sense of being both systematic and dynamical-possibility-preserving).[46]

What is the relationship to Quine's original argument? In my view, the above provides a substantially more concrete and plausible example for the indeterminacy thesis than the examples adduced by Quine. For example, after introducing the infamous field linguist trying his best to translate "gavagai" as either "rabbit" or "rabbit part", Quine notes that as he offers more and more hypotheses relating terms of the native language to those of English, and as those hypotheses continue to match sentences to sentences of similar usage, he will feel more and more confident that those hypotheses are correct. However, Quine continues,

> [. . .] it seems that this method, though laudable in practice and the best we can hope for, does not in principle settle the indeterminacy between "rabbit", "undetached rabbit part", and "rabbit stage". For if one workable system of analytical hypotheses provides for translating a given native expression into "is the same as," perhaps another equally workable but systematically different system would translate that native expression into something like "belongs with". Then when in the native language we try to ask "Is this *gavagai* the same as that?" we could as well be asking "Does this *gavagai* belong with that?" Insofar, the native's assent is no objective evidence for translating "gavagai" as "rabbit" rather than "undetached rabbit part" or "rabbit stage".[47]

---

[46][Healey, 2006] gives an argument similar to this, although based on ideas from Lewis (specifically, [Lewis, 1970]) rather than Quine, and more naturally located within a syntactic approach to theories than the broadly semantic approach taken here.

[47][Quine, 1969, p. 33]

If we examine the above carefully, we see that a remarkable amount of work is being done by that "perhaps". The evidence for the claim that "a grand and ingenious permutation of these denotations, along with compensatory adjustments in the interpretations of the auxiliary particles, might still accommodate all existing speech dispositions"[48] is typically taken to be the possibility of proxy functions:[49] functions which map each object into the domain of discourse into some other such object.[50] Given a proxy function f, we reinterpret every term t as referring to $f([[t]])$ ($[[t]]$ being the standard referent of t), and every predicate P as applying to some object o if and only if $[[P]]$ applies to $f^{-1}(o)$ ($[[P]]$ being the standard interpretation of P). However, this does not suffice to show that the alternative system of meaning-hypotheses thereby obtained will be "equally workable but systematically different". On the contrary, an arbitrary proxy-function will yield predicates that apply in an arbitrary and gerrymandered fashion to objects possessing no real unity. Whether that can be made to matter is controversial, of course; but nevertheless, it does seem a little premature of Quine to conclude that these observations have "made nonsense of reference."[51]

In the context of a theory containing symmetries, by contrast, we *are* in a position to check that there are indeed such permutations of meaning that will preserve all speech dispositions, yet will also yield extensions for predicates that are as unified as their original extensions were. Regarding the first, the relevant dispositions are the circumstances under which a model might be put forward as representing the actual world; and given the observational equivalence of pseudomorphic models, this means that the conditions under which one model could be put forward are equally good reasons to put forward any of its pseudomorphic companions. Regarding the second, the systematic nature of the alteration means that the new extension of any particular absolute-velocity predicate is the old extension of some other such predicate: so the latter must be as unified as the former.

The argument is not bulletproof, though. In particular, one might be concerned that a given model is naturally most suited to represent some particular world: namely, one in which two temporal stages of spatial points are represented by some $(\mathbf{x}, t)$ and $(\mathbf{x}, t')$ if and only if they are temporal stages of the very same spatial point. Or, if it is disputed that such considerations of "naturalness" ought to be applied to meaning-determinations, then one could simply try ostension: just *stipulate* that the relationship of "being stages of the same persisting object"—the kind of relationship with which one is acquainted from ordinary cases of persisting physical objects—is to be represented by the mathematical relationship of "sharing a first member." In other words, do not the (putative) facts regarding the persistence of the points of absolute space provide the means by which a model could be yoked to one world in particu-

---

[48][Quine, 1969, p. 48]

[49]I thank Simon Saunders for pushing me to consider the role of proxy functions in Quine's thinking.

[50]See [Quine, 1969, p. 55 ff.], or [Quine, 1992, §12]. I pass over Quine's other example—that of the Japanese classifiers—for the simple reason of not knowing enough about Japanese to be able to make any kind of informed comment.

[51][Quine, 1969, p. 48]

lar?[52]

A response has to accept that these show how one could make the representation relation determinate, after all. However, it bears emphasising that they do so only by importing resources external to the theory itself. For all the theory cares, one can use I_0, or $I_{\mathbf{a}}$, or whatever correspondence one likes; and that observation stands, regardless of whether extra-theoretical considerations of (metaphysical) naturalness, or stipulations grounded in ostension, serve to privilege one correspondence over the others. Thus, insofar as one is interested in accounting for the meaning imbued upon the parts of a theory by their role within that theory, the reference of these models remains indeterminate—unless, that is, they are all taken to refer identically, to one and the same possible world.

9 Concluding remarks

These, then, are my reasons for thinking that such an attitude towards the interpretation of pseudomorphic models is defensible; note that although I have confined my attention in this paper to the specific case of boost symmetries in NG, the arguments should readily admit of generalisation to other symmetries in other theories.[53] The final comment I wish to make is that such a process is the natural extension of a way of specifying the content of a theory that is—or at least, should be—wholly uncontroversial. For giving a theory by offering an overly rich class of models, and then asserting that models of a certain kind are co-representational (so that one is committed only to whatever structures those models have in common) is *precisely* what is involved in giving a theory in coordinate form! That is, suppose that we specify our spacetime structure by giving a specification of the nature of the adapted coordinate systems, before then going on to specify the laws of the theory as expressed in those coordinate systems. The laws serve to pick out a class of models, all set on $\mathbb{R}^4$; to avoid committing ourselves to the presence of an absolute centre of the universe, an absolute time function, and absolute orientations, we specify that certain classes of models (namely, those related by a coordinate transformation taking adapted coordinates to adapted coordinates) are physically equivalent, so that the differences between them wash out.

If this kind of method is acceptable as a means of disavowing commitment to spatiotemporal structure, then there is simply no reason why its extension defended here should not be. If it is not, then we are left with the conclusion that—contrary to what

---

[52]I owe this objection to Oliver Pooley.

[53]Moreover, by advocating that one's theoretical commitments be expressed by a particular attitude towards the relationships between different models (rather than an attitude towards the components of those models), the lessons of this paper could naturally be generalised to link up with recent work on theoretical equivalence: in particular, [Halvorson, 2012] and [Weatherall, 2011] both stress the importance of relationships between models (not just within models) when making judgments of theoretical equivalence. I hope to follow up this connection in future work. Thanks to an anonymous referee for suggesting this connection.

one might think—Einstein did not, in fact, succeed in putting forward a theory without absolute simultaneity in 1905, since every solution of his theory he considered was framed in some coordinate system or other.[54] Indeed, arguably we do not have a theory meeting this demand even now: after all, even "intrinsic" formulations of differential geometry start off by defining the topological and differential structure of a manifold by specification of the coordinate systems that get that structure right![55]

Of course, this isn't to say that there is *no* value to reformulating a theory's formalism in such a way that the surplus structure is made manifest, so that we can move to a formalism in which it is expunged entirely. Such a presentation lets us see what it is we are committed to by our (qualified) realism about the theory; if we want to know the answers to specific questions about the nature of a theory's ontology and ideology, then this is invaluable. Moreover, intrinsic formulations may help theoretical progress, by *ensuring* that the surplus structure is not taken seriously: by his own account, the main reason it took Einstein so long to formulate General Relativity was "the fact that it is not so easy to free oneself from the idea that the coordinates must have an immediate metrical meaning."[56] But if we lack the mathematical tools to do so, then I maintain that there is nothing wrong with recognising that one's realism will only extend to structures that are invariant under the symmetries—whatever those may turn out to be.

Acknowledgements

References

[Awodey, 2010] Awodey, S. (2010). *Category theory*. Oxford University Press, Oxford; New York.

[Azzouni, 2009] Azzouni, J. (2009). Evading truth commitments: The problem reanalyzed. *Logique Et Analyse*, 52(206):139–176.

---

[54]Incidentally, the same remark cannot be made about Newton, given that his theory was presented geometrically (as [Maudlin, 2012, chap. 2]) observes.

[55]I thank David Wallace for this observation.

[56][Einstein, 1982, p. 67]; quoted in [Maudlin, 2012, p. 26].

[Baker, 2009] Baker, A. (2009). Mathematical explanation in science. *The British Journal for the Philosophy of Science*, 60(3):611–633.

[Belot, 2013] Belot, G. (2013). Symmetry and equivalence. In Batterman, R. W., editor, *The Oxford Handbook of Philosophy of Physics*. Oxford University Press, New York.

[Colyvan, 2010] Colyvan, M. (2010). There is no easy road to nominalism. *Mind*, 119(474):285–306.

[Dasgupta, 2009] Dasgupta, S. (2009). *Symmetry as a guide to reality*. Ph.d. thesis, New York University.

[Dasgupta, 2011] Dasgupta, S. (2011). The bare necessities. *Philosophical Perspectives*, 25(1):115–160.

[Dasgupta, 2013] Dasgupta, S. (2013). Absolutism vs. comparativism about quantity. In *Oxford Studies in Metaphysics, vol. 8*. Oxford Univerity Press, Oxford.

[Dasgupta, 2014] Dasgupta, S. (2014). Symmetry as an epistemic notion (twice over). *The British Journal for the Philosophy of Science*, Forthcoming.

[Davidson, 1965] Davidson, D. (1965). Theories of meaning and learnable languages. *Logic, Methodology and Philosophy of Science*, 2:383–394.

[Davidson, 1967] Davidson, D. (1967). Truth and meaning. *Synthese*, 17(1):304–323.

[Davidson, 2001] Davidson, D. (2001). *Inquiries into truth and interpretation*. Oxford University Press, Oxford; New York.

[Earman, 1989] Earman, J. (1989). *World enough and space-time: absolute versus relational theories of space and time*. MIT Press, Cambridge, Mass.

[Einstein, 1982] Einstein, A. (1982). Autobiographical notes. In *Albert Einstein: Philosopher-Scientist*, volume 1. Open Court, La Salle.

[Friedman, 1983] Friedman, M. (1983). *Foundations of space-time theories: relativistic physics and philosophy of science*. Princeton University Press, Princeton, N.J.

[Glock, 2003] Glock, H.-J. (2003). *Quine and Davidson on language, thought, and reality*. Cambridge University Press, Cambridge, UK; New York, NY, USA.

[Halvorson, 2012] Halvorson, H. (2012). What scientific theories could not be. *Philosophy of Science*, 79(2):183–206.

[Healey, 1999] Healey, R. (1999). Quantum analogies: A reply to Maudlin. *Philosophy of Science*, 66(3):440–447.

[Healey, 2006] Healey, R. (2006). Symmetry and the scope of scientific realism. In Demopoulos, W. and Pitowsky, I., editors, *Physical Theory and its Interpretation*, number 72 in The Western Ontario Series in Philosophy of Science, pages 143–160. Springer Netherlands.

[Hodges, 1997] Hodges, W. (1997). *A shorter model theory*. Cambridge University Press, Cambridge; New York.

[Huggett, 2003] Huggett, N. (2003). Mirror symmetry: what is it for a relational space to be orientable? In Brading, K. and Castellani, E., editors, *Symmetries in Physics: Philosophical Reflections*, pages 281–288. Cambridge University Press, Cambridge.

[Knox, 2013] Knox, E. (2013). Newtonian spacetime structure in light of the equivalence principle. *The British Journal for the Philosophy of Science*, Forthcoming.

[Leng, 2010] Leng, M. (2010). *Mathematics and reality*. Oxford University Press, Oxford.

[Lewis, 1970] Lewis, D. (1970). How to define theoretical terms. *The Journal of Philosophy*, 67(13):427–446.

[Maudlin, 1988] Maudlin, T. (1988). The essence of space-time. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 1988:82–91.

[Maudlin, 1993] Maudlin, T. (1993). Buckets of water and waves of space: Why spacetime is probably a substance. *Philosophy of Science*, 60(2):183–203.

[Maudlin, 2012] Maudlin, T. (2012). *Philosophy of physics: space and time*. Princeton University Press, Princeton.

[Melia, 1995] Melia, J. (1995). On what there's not. *Analysis*, 55(4):223–229.

[Melia, 1999] Melia, J. (1999). Holes, haecceitism and two conceptions of determinism. *The British Journal for the Philosophy of Science*, 50(4):639–664.

[Melia, 2000] Melia, J. (2000). Weaseling away the indispensability argument. *Mind, New Series*, 109(435):455–479.

[Penrose, 2004] Penrose, R. (2004). *The road to reality*. Jonathan Cape, London.

[Pooley, 2002] Pooley, O. (2002). *The reality of spacetime*. DPhil thesis, University of Oxford.

[Pooley, 2003] Pooley, O. (2003). Handedness, parity violation, and the reality of space. In Brading, K. and Castellani, E., editors, *Symmetries in Physics: Philosophical Reflections*, pages 250–280. Cambridge University Press, Cambridge.

[Pooley, 2013] Pooley, O. (2013). *The Reality of Spacetime*. Unpublished draft.

[Quine, 1992] Quine, W. V. (1992). *Pursuit of truth*. Harvard University Press, Cambridge, Mass.; London, revised edition.

[Quine, 1969] Quine, W. V. O. (1969). Ontological relativity. In *Ontological Relativity and Other Essays*, pages 26–68. Columbia University Press, New York.

[Roberts, 2008] Roberts, J. T. (2008). A puzzle about laws, symmetries and measurability. *The British Journal for the Philosophy of Science*, 59(2):143–168.

[Russell, 2011] Russell, J. S. (2011). *Possible worlds and the objective world*. Phd thesis, New York University.

[Saunders, 2003a] Saunders, S. (2003a). Indiscernibles, general covariance, and other symmetries: the case for non-eliminativist relationalism. In Ashtekar, A., Howard, D., Renn, J., Sarkar, S., and Shimony, A., editors, *Revisiting the Foundations of Relativistic Physics: Festschrift in Honour of John Stachel*. Kluwer, Dordrecht.

[Saunders, 2003b] Saunders, S. (2003b). Physics and Leibniz's principles. In Brading, K. and Castellani, E., editors, *Symmetries in Physics: Philosophical Reflections*, pages 289–308. Cambridge University Press, Cambridge.

[Saunders, 2013] Saunders, S. (2013). Rethinking Newton's principia. *Philosophy of Science*, 80(1):22–48.

[Wald, 1984] Wald, R. M. (1984). *General relativity*. University of Chicago Press, Chicago.

[Weatherall, 2011] Weatherall, J. O. (2011). Are Newtonian gravitation and geometrized Newtonian gravitation theoretically equivalent? Unpublished draft.

[Weatherall, 2014] Weatherall, J. O. (2014). Regarding the "Hole Argument". Unpublished draft.

[Yablo, 2005] Yablo, S. (2005). The myth of the seven. In Kalderon, M. E., editor, *Fictionalism in Metaphysics*, pages 88–115. Oxford University Press, Oxford.