

Invariance, Interpretation, and Motivation

Thomas Møller-Nielsen

July 2016

[Forthcoming in *Philosophy of Science (2016 Proceedings)*.]

Abstract

In this paper I assess the ‘Invariance Principle’, which states that only quantities that are invariant under the symmetries of our theories are physically real. I argue, contrary to current orthodoxy, that the variance of a quantity under a theory’s symmetries is not a sufficient basis for interpreting that theory as being uncommitted to the reality of that quantity. Rather, I argue, the variance of a quantity under symmetries only ever serves as a motivation to refrain from any commitment to the quantity in question. In the process of this discussion, I address the related but importantly distinct issue of when symmetries can be said to prompt a mathematical reformulation of the relevant theory.

1 Introduction

Take the *Invariance Principle* to be the principle that only quantities that are invariant under the symmetries of our theories are physically real.¹ It is a doctrine with a distinguished pedigree: acclaimed theorists as diverse as the physicist Paul Dirac, the mathematician Hermann Weyl, and the philosopher Robert Nozick were all apparent signatories during their respective lifetimes.² *Prima facie*, however, it is something of a mystery as to how and why the principle is supposed to work. Nevertheless, there appear to be at least some uncontroversial cases where it—or something very close to it—does work.

One such example can be found in Newtonian Gravitation Theory (NGT), i.e., the theory comprising Newton’s three laws, plus his inverse square gravitational law, governing the behaviour of point particles in Newtonian spacetime. As is well known, this theory is *Galilean invariant*. This implies, among other things, that if one takes any solution to NGT and “boosts” it—that is, uniformly alters the absolute velocity of each point particle by the same amount throughout its history—one will invariably get back a solution to NGT. Boosts, in other words, are a *symmetry* of NGT: they are transformations that invariably map solutions of the theory to solutions.

¹I draw the term from Saunders (2007). Compare also Dasgupta’s (forthcoming) “symmetry-to-reality inference”.

²See, e.g., Dirac (1930, vii), Weyl (1952, 132), and Nozick (2001, 82).

Which quantity varies under this particular symmetry? The answer is obvious: absolute velocity. Thus, according to the Invariance Principle, we should conclude that absolute velocity is not a genuine physical quantity. Conversely, which quantities are invariant under this particular symmetry? Again, the answer is obvious: relative (inter-particle) distance and velocity, temporal intervals, and absolute acceleration. Thus, according to the Invariance Principle, we should conclude that NGT’s boost symmetry does not threaten these quantities’ status as genuinely physical.

As it turns out, one can successfully purge Newtonian theory of the spacetime structure required to make absolute velocity a physically meaningful quantity. More specifically, one can move to *Galilean spacetime*. (Sometimes also called “Neo-Newtonian spacetime”.)³ Here, the Newtonian posit of persisting points of absolute space—persisting points which, crucially, allow for the notion of absolute velocity to be physically meaningful—is done away with, but an *affine structure* is nevertheless preserved, which defines the “straight” or force-free (inertial) paths through spacetime. Absolute velocity is therefore not a physically meaningful quantity in Galilean spacetime, as it is in Newtonian spacetime. Nevertheless, all other Newtonian notions, including the notion of absolute acceleration, remain well-defined in Galilean spacetime. To the extent that one opts for Galilean over Newtonian spacetime, then, one has excised an ostensibly odious piece of theoretical structure from NGT.

Three important caveats are worth noting, however. First, and most obviously, none of this is to say that Newtonian theory set in Galilean spacetime is therefore the true and complete theory of the world. (It isn’t.) Second, nor is this to say that by moving to Galilean spacetime one has thereby purged Newtonian theory of all its “variant” structure. (One hasn’t. The symmetry group of Newtonian theory is actually wider than the Galilean group: it has additional symmetries.)⁴ Third, nor is this even to say that the invariant quantities one ends up with following such an application of the Invariance Principle will invariably be preserved in future theories. (For instance, there is no notion of “relative spatial distance” *simpliciter* in special relativity.) Given all of these caveats, however, one might well ask: What good is the Invariance Principle, exactly? What purpose, in particular, does it serve?

As I see it—and, I take it, as many other contemporary theorists also see it—the purpose of the Invariance Principle is essentially *comparative*. That is, it is simply supposed to lead you to a *better theory*—or a better interpretation, or characterisation, of the same theory—than the one you started with. To take the case at hand: Newtonian theory set in Galilean spacetime is a better theory than Newtonian theory set in Newtonian spacetime. It is a theory which possesses all of the theoretical virtues of its rival, but lacks any apparent ontological commitment to the unwanted variant quantity in question.

In summary, the Galilean invariance of NGT, in conjunction with the Invariance Principle, is supposed to indicate that neither absolute velocity nor

³See, e.g., Earman (1989, §2.4).

⁴See, e.g., Knox (2014). I discuss this point further in Section 4 below.

any corresponding persisting points of absolute space are genuinely real. Now to lay my cards on the table: I actually think that something *very close* to this general kind of inference—that is, from the variance of a quantity under symmetries to that quantity’s nonreality—is legitimate. The devil, however, is in the details. In particular, I don’t believe that the *mere* Galilean invariance of NGT is enough to establish absolute velocity’s nonreality. And in general, I don’t believe that the *mere* variance of a quantity under symmetries is enough to establish that quantity’s nonreality. These beliefs, as far as I can determine, put me in the minority camp in the contemporary philosophical literature on symmetries. Nevertheless, I think they are correct beliefs—and they are precisely the ones that I will attempt to argue for in the remainder of this paper.

2 Interpretational vs Motivational

In arguing for the above claims, it will prove extremely useful first to distinguish between two very different ways of thinking about symmetries.

Close cousins of the distinction that I have in mind have already been drawn in the literature. Thus, Greaves and Wallace write:

There is a widespread consensus that two states of affairs related by a symmetry transformation are really just the same state of affairs differently described. That is, if two mathematical models of a physical theory are related by a symmetry transformation, then those models represent one and the same physical state of affairs. (Greaves and Wallace 2014, 60)

They continue:

Although we agree with this consensus [...] even those who do not agree that symmetry-related states of affairs are identical at least agree that they are *empirically indistinguishable* from one another. (Greaves and Wallace 2014, 60, fn 1)

To illustrate the difference between these two ways of thinking about symmetries, consider again the example of boosts in NGT. According to the “widespread consensus” view alluded to, and endorsed by, Greaves and Wallace, boosted models of NGT are to be taken to represent the same physical state of affairs *even when the theory is putatively set in Newtonian spacetime*. In other words, according to this view, one needn’t make the move to Galilean spacetime in order not to be committed to absolute velocities; there is a way of understanding boosted models’ physical equivalence, and their associated noncommitment to the notion of absolute velocity, prior to making this move.⁵

Things are very different according to the second conception of symmetries described, and rejected, by Greaves and Wallace. According to this view, boosted models of NGT are to be regarded as physically *inequivalent*: they are not to be construed as representing the same physical state of affairs. Instead,

⁵See, e.g., Healey (2007, 114-7), for an endorsement of this view in the Newtonian context.

such models are taken to represent physically distinct scenarios, which differ in what absolute velocity they ascribe to the world’s total material content. Nevertheless, such models still represent *empirically indistinguishable* states of affairs: in a Newtonian universe, no experiment could ever help an observer determine what her absolute velocity actually is. Such boosted models therefore represent physically distinct ways for the world to be, albeit ones that are indiscernible on the basis of measurement.⁶

As previously mentioned, this distinction between different ways of thinking about symmetries is close, but not identical, to the one that I want to draw. The key reason why it is not identical is because Greaves and Wallace say nothing to the effect that the person who subscribes to the second conception of symmetries—that is, who believes that symmetry-related models invariably represent empirically indistinguishable, but not necessarily physically equivalent, states of affairs—should still be *motivated to seek* an alternative theory, or an alternative interpretation or characterisation of the same theory, according to which such models do not merely represent empirically indistinguishable scenarios, but rather represent physically equivalent states of affairs.⁷ Moreover, I claim, it is precisely this notion of *motivation* which plays a central role in correctly understanding the philosophical significance of symmetries in the general case.⁸

Here, then, is what I take to be the appropriate distinction between these two different ways of thinking about symmetries:

- **Interpretational:** Symmetries allow us to *interpret* theories as being committed solely to the existence of invariant quantities, even in the absence of a metaphysically perspicuous characterisation of the reality which is alleged to underlie symmetry-related models.
- **Motivational:** Symmetries only *motivate* us to find a metaphysically perspicuous characterisation of the reality which is alleged to underlie symmetry-related models, but they do not allow us to interpret that theory as being solely committed to the existence of invariant quantities in the absence of any such characterisation.

The central claim of this paper may now be neatly summarised: the (orthodox) interpretational view is mistaken; the (unorthodox) motivational view is correct.

Drawing the distinction in the way that I have done, however, invites the rather obvious question: What, precisely, is meant by a “metaphysically perspicuous characterisation” of reality? This is the question addressed in the next section.

⁶See, e.g., Maudlin (1993, 192), for an endorsement of this view in the Newtonian context.

⁷Compare (again) Maudlin’s (1993, 192) discussion in the Newtonian context.

⁸Note that I do not intend any of this as a criticism of Greaves and Wallace’s paper. Indeed, as Greaves and Wallace (2014, 60, fn 1) are careful to remark, the distinction they draw is orthogonal to the central topic of their paper, namely the issue of which symmetries have “direct empirical significance” (i.e., have analogues to Galileo’s ship).

3 More on Metaphysical Perspicuity

In intuitive terms, a metaphysically perspicuous characterisation of reality is one which corresponds to, or “limns”, reality’s structure in some suitably faithful way. To use another common (Platonic) metaphor, a metaphysically perspicuous characterisation of reality is one which “carves nature at its joints”. (In comparative terms: a description of reality is *more* metaphysically perspicuous than another precisely to the extent that it corresponds to, or limns, reality’s structure *more* faithfully than its rival does.)

As many readers will be aware, such a notion is frequently alluded to, and made use of, in contemporary analytic metaphysics.⁹ But metaphysical perspicuity is also, I think, a notion that is reasonably serviceable in physical (rather than “merely metaphysical”) contexts. One particularly illustrative example—albeit a slightly misleading one, for reasons that I will soon explain—drawn from physics may plausibly be found in classical electromagnetism.¹⁰ As is well known, this theory may be formulated in two different ways.¹¹ According to one such formulation, EM₁, the theory is expressed in terms of the Faraday tensor, F_{ab} , satisfying the (Maxwell) equations $\nabla_{[a}F_{bc]} = 0$ and $\nabla_a F^{ab} = J^a$, where J^a is a vector field representing the charge current density. According to the second formulation, EM₂, however, the theory is expressed in terms of the vector potential, A_a , satisfying the equation $\nabla_a \nabla^a A^b - \nabla^b \nabla_a A^a = J^b$.

These two formulations of electromagnetism are related to one another. In particular, any model $\langle M, \eta_{ab}, A_a \rangle$ of EM₂ corresponds to a unique model $\langle M, \eta_{ab}, F_{ab} \rangle$ of EM₁, via the equation $F_{ab} = \nabla_{[a}A_{b]}$. The converse, however, is not true. That is, a typical model of EM₁ does *not* typically correspond to a unique model of EM₂. More specifically, if $\langle M, \eta_{ab}, A_a \rangle$ is a model of EM₂ corresponding to a model $\langle M, \eta_{ab}, F_{ab} \rangle$ of EM₁, then so will any other model of EM₂ $\langle M, \eta_{ab}, A'_a \rangle$, where A'_a is related to A_a by a “gauge transformation” $A'_a = A_a + \nabla_a \chi$, where χ is some smooth scalar field.

It is EM₁ which, I take it, constitutes the metaphysically perspicuous characterisation of this theory. That is, it is the tensor F_{ab} which faithfully represents the fundamental ontology of the theory, namely the electromagnetic field. Not so EM₂. This second formulation may, of course, be useful for various calculational or heuristic purposes. But the key point is that the vector potential A_a *does not directly represent a genuinely real field*: rather, it is merely a mathematically convenient “shorthand” way of characterising and determining the values of the Faraday tensor, which *is* taken to represent the genuine material ontology of the theory.¹² Moreover, it is precisely by construing the vector potential in this

⁹See, e.g., O’Leary-Hawthorne and Cortens (1995, 154-7).

¹⁰Here and below, I take this theory to be set in Minkowski spacetime. Thus, the spacetime models of this theory are of the form $\langle M, \eta_{ab} \rangle$, where M is a four-dimensional differentiable manifold, and η_{ab} is the Minkowski metric.

¹¹For a recent, intriguing study of the relationship between these two different formulations of electromagnetism, see Weatherall (forthcoming). I draw heavily on his discussion over the next couple of paragraphs.

¹²Modulo, that is, certain concerns that arise as a result of the Aharonov-Bohm effect. See, e.g., Healey (2007).

way which plausibly allows us to explain and understand, in a fully transparent way, gauge-symmetry models' physical equivalence in EM_2 —namely, for the reason that they are merely notationally distinct ways of representing the same fundamental physical ontology.

As mentioned above, I think this example of metaphysical perspicuity is apt to be slightly misleading, at least when taken on its own. This is because this example might make it seem as though having a metaphysically perspicuous characterisation of the (putative) reality underlying symmetry-related models crucially relies upon one having to *mathematically reformulate* the relevant theory (or at least upon having such a mathematical reformulation already in hand), and in particular upon having to reformulate the theory so as to remove any relevant representational redundancy. However, I think this is incorrect. That is, I believe that one *can*, in fact, be in possession of a metaphysically perspicuous characterisation of the reality underlying symmetry-related models *even in the absence* of any mathematical (re-)formulation of the theory which removes the relevant representational redundancy.

Let me illustrate this point with two simple examples. First, consider the case of *shift symmetry* in NGT. This symmetry is subtly different from the case of boost symmetry, discussed above. Here, instead of uniformly altering the absolute velocity of each particle throughout its history, one enacts a global, time-independent repositioning of all matter in space. Thus, for instance, in the shifted world all of the world's material content will (*prima facie*) be located three metres to the left of where it is in the original world. The basic idea behind the “Leibniz shift” argument—the famous argument associated with this symmetry—is that the substantialist's admission of points of space as primitive objects (allegedly) has the undesirable consequence of committing her to regarding shifted worlds as physically distinct, yet nevertheless empirically indistinguishable:¹³ in intuitive terms, everything would look, feel, taste, touch and sound the same in the two (putatively distinct) shifted worlds, just as in the case of boosted worlds.

It will prove helpful to express all of this in terms of the models of the theory. Thus, take a generic model of NGT to be of the form $\mathcal{M} = \langle M, t_{ab}, h^{ab}, \sigma^a, \rho, \phi \rangle$, where M is a differentiable 4-dimensional manifold, t_{ab} is the temporal metric, h^{ab} is the spatial metric, σ^a is the timelike vector field whose integral curves represent the persisting points of absolute space, and ρ and ϕ represent the matter density and the gravitational potential field respectively.¹⁴ A shift symmetry can then be characterised as the application of the appropriate diffeomorphism (corresponding to a spatial translation) d so as to yield a new model $\mathcal{M}_{static} = \langle M, t_{ab}, h^{ab}, \sigma^a, d^* \rho, d^* \phi \rangle$. It is then alleged that \mathcal{M} and \mathcal{M}_{static} differ precisely

¹³Though see Maudlin (1993), who notes that there is an interesting (epistemological) sense in which shifted worlds in NGT are not indiscernible after all.

¹⁴Note that the canonical presentations of Newtonian spacetime (e.g., Earman 1989, §2.5) take the affine connection as ideologically primitive. I find such presentations unsatisfactory for historical rather than for philosophical reasons: in particular, it threatens to make the move to Galilean spacetime seem almost trivial, and the associated timelike vector field trivially superfluous. For more on this point, see Pooley (MS, §4.4-§4.5).

insofar as they each represent the world’s matter content as being located at distinct places in absolute space. More specifically, such Leibniz-shifted scenarios are alleged to differ precisely with regard to which particular points of space are underlying various parts of the matter fields.

For a second example, consider *diffeomorphism symmetry* in general relativity (GR). Here, similarly, the existence of this symmetry is alleged to commit the substantialist to a plurality of physically distinct possibilities that are nevertheless empirically indistinguishable. In terms of the models of the theory: taking a generic model of GR to be of the form $\mathcal{M} = \langle M, g_{ab}, T_{ab} \rangle$ and applying an arbitrary diffeomorphism d to yield a new model $\mathcal{M}_{diff} = \langle M, d^*g_{ab}, d^*T_{ab} \rangle$ (where M is again a differentiable 4-dimensional manifold, g_{ab} is the metric tensor, and T_{ab} is the stress-energy tensor which, roughly speaking, represents the model’s matter content), the two scenarios represented are alleged to differ with regard to which particular points of the spacetime manifold are underlying various parts of the metric and matter fields.¹⁵

It is my contention that neither the shift symmetry of NGT, nor the diffeomorphism symmetry of general relativity, by themselves motivate any mathematical reconstrual of the respective theories. This is because I believe there is a perfectly transparent, anti-haecceitist, “modestly structuralist”—but nevertheless fully substantialist—way of understanding such models’ representational equivalence even in the absence of any such mathematical reformulation. On this view, spacetime points are construed as genuinely real, fundamental entities. However, they are “contextually individuated”: they are not to be understood as being anything more—or less—than “nodes” in the relational, geometrical structures in which they are embedded. Shifted models in NGT and diffeomorphically-related models in GR are thus to be understood as representing the same physical state of affairs precisely because the exact same pattern of relational, geometrical structures is represented as obtaining in each case. Moreover, this view denies that there are any primitive, singular (“haecceitistic”) facts about spacetime points which would even allow for a distinction between shifted or diffeomorphically-related scenarios to be coherently drawn.¹⁶

Whence the difference, then, between the case of gauge symmetry in electromagnetism on the one hand, and shift and diffeomorphism symmetry in NGT and GR on the other? I think the answer is straightforward. In the latter cases, the models in question are *isomorphic*: they represent worlds which differ at most with regard to which particular objects are playing which qualitative roles, i.e., they represent at most haecceitistically distinct possible worlds. Hence, adopting modest structuralism (which implies anti-haecceitism) about spacetime transparently collapses the number of possibilities represented by these models to one. In the former such case, however, the relevant models are *not* isomorphic—read “literally”, gauge-related models of EM₂ assign *qualitatively distinct* arrangements of the vector field over spacetime—hence adopting a modestly structuralist ontology does not by itself collapse the number of represented

¹⁵For further details see, e.g., Earman (1989, §9).

¹⁶For further defence of this view—which is sometimes also called *sophisticated substantialism* in the literature—see, e.g., Saunders (2003), Ladyman (2007), and Pooley (2013).

possibilities to one. In order to transparently understand such models' physical equivalence, then, a mathematical reformulation of the theory is required.

To summarise the claims made thus far: according to the motivational view of symmetries, one is invariably only motivated to regard symmetry-related models as physically equivalent; moreover, one is justified in regarding such models as physically equivalent only insofar as one is in possession of a metaphysically perspicuous characterisation of the reality which is alleged to underlie them. However, it is possible to be in possession of a metaphysically perspicuous characterisation of the reality underlying symmetry-related models even in the absence of a mathematical formulation of the theory which removes the relevant representational redundancy. Such a metaphysically perspicuous characterisation is possible just in case the symmetry-related models in question are isomorphic, or are naturally understood as representing at most haecceitistically distinct possibilities. In brief: symmetry-related, isomorphic models invariably do *not* motivate a mathematical reformulation of the relevant theory (modest structuralism invariably suffices); but symmetry-related, *non*-isomorphic models invariably *do*.¹⁷

4 In Defence of the Motivational View

Let us return once more to the case of NGT. As alluded to in Section 1, the symmetry group of this theory is quite large. For not only does it include transformations corresponding to global velocity boosts of solutions' matter content, but it also includes transformations corresponding to time-dependent translational accelerations of such content (so long as the gravitational potential field is also appropriately transformed). Thus, read "literally", the symmetries of this theory include transformations that map solutions to solutions that represent physically distinct, but nevertheless empirically indistinguishable, states of affairs in which a given material system is:

1. Force-free and stationary with respect to absolute space.
2. Force-free and moving at constant absolute velocity.
3. Absolutely accelerating under a gravitational force-field.

According to the interpretational conception of symmetries, we may legitimately take all of these symmetry-related solutions to in fact represent the same physical state of affairs—despite the fact that they are naturally understood as representing radically distinct physical situations. Things are very different, however, according to the motivational conception of symmetries. On this view, we are merely *motivated to regard* all such solutions as representing the same physical state of affairs, the motivation arising from the general Occamist principle that, other things being equal, our preferred scientific theories should not allow for solutions that represent physically distinct but nevertheless empirically indistinguishable possible worlds. According to the motivational

¹⁷See also Pooley (2013, 576-7) and Weatherall (forthcoming) for recent, related arguments to this effect.

view, then (and to repeat slightly), absent a metaphysically perspicuous characterisation of the reality underlying these symmetry-related models, we have no choice but to regard them as representing physically distinct states of affairs.

For our purposes, the crucial thing to note about all of these models is that *none of them are isomorphic*—naturally understood, they do not represent at most haecceitistically distinct possible worlds. According to the criterion laid down in the previous section, then, in order to be able to transparently understand how it could be that such models may be said to represent physically equivalent scenarios, a mathematical reformulation of the theory is required.

As it turns out, such a mathematical reformulation of the theory is possible. In brief, in this reformulation one replaces the vector field σ^a with a new kind of *dynamical* inertial connection ∇^{NC} , with models of the form $\mathcal{M}_{NC} = \langle M, t_{ab}, h^{ab}, \nabla^{NC}, \rho \rangle$. Up to isomorphism, any two symmetry-related models of NGT correspond to a unique model of Newtonian gravity geometrised in this way. Thus, it is said, by moving to this “Newton-Cartan” theory one successfully removes the undesirable “gauge-redundancy” inherent in all non-geometrised versions of Newtonian gravitation theory.¹⁸

What might the defender of the interpretational view of symmetries say in defence of her view—in this context, that the move to Newton-Cartan theory is not required in order to be able to legitimately regard all symmetry-related solutions of NGT as physically equivalent?

I anticipate two likely lines of response. First, she might attempt to establish the preferability of her view over the motivational view by noting that the defender of the motivational view is committed, at least prior to the appropriate theory’s reformulation (in the context of NGT), to the existence of in principle undetectable (symmetry-variant) matters of fact. Moreover, the defender of the interpretational view might argue, this is an unpalatable consequence, one which we would do best to avoid—and one which, she might point out, the interpretational view does in fact avoid.

I agree that the admission of such in principle undetectable facts is an undesirable consequence of the motivational view. However, I do not think that this admission is sufficiently unpalatable so as to be capable of refuting the motivational view, or even of establishing the preferability of the interpretational view over the motivational view. After all, prohibitively strong versions of verificationism aside, there is nothing obviously absurd about admitting in principle undetectable facts into one’s ontology; nor is there any obvious reason why we should always be capable of discovering a theory, or a perspicuous characterisation thereof (the case of isomorphic models excepted), which succeeds in transparently explaining such solutions’ empirical equivalence by virtue of

¹⁸For further details, see, e.g., Knox (2014). Note also the important point that moving to Newton-Cartan theory is not by itself sufficient for one to be able to transparently understand as physically equivalent all symmetry-related models of Newtonian theory set in flat spacetime. This is because—as mentioned above—such symmetry-related models will typically correspond to a single model of Newton-Cartan theory *only up to isomorphism*. Thus, in order to have a *fully* transparent understanding of how it is that symmetry-related models of Newtonian theory set in flat spacetime can correspond to a single model of Newton-Cartan theory, a modestly structuralist conception of spacetime ontology is also required.

their actual physical equivalence; nor indeed is there even any obvious way of guaranteeing that there will always be such a theory or characterisation (again, isomorphic models excepted) waiting in logical space to be discovered.

Furthermore, although it is to be admitted that the Newtonian who subscribes to the merely motivational view of symmetries might indeed be committed to the possibility of there being facts beyond her epistemic grasp, it nevertheless bears emphasising that for such a Newtonian there is a perfectly good explanation as to *why* such facts are epistemically inaccessible: they are inaccessible precisely because the world is in fact accurately described by the laws of NGT, with associated models of the form $\langle M, t_{ab}, h^{ab}, \sigma^a, \rho, \phi \rangle$, and because all any Newtonian observer ultimately has empirical access to are the relative distances and velocities between material entities. For such a Newtonian, then, the empirical phenomena underdetermine the genuine physical facts; but the theory itself is able to provide a perfectly transparent explanation of the reality behind the phenomena in terms of which the underdetermination can be straightforwardly understood.

The Newtonian who adopts the interpretational construal of symmetries, however, would appear to lose this explanatory transparency. In other words, she might know *that* she may legitimately regard all symmetry-related solutions as physically equivalent; but the reality in terms of which this physical equivalence is to be understood will (absent a reformulation of the theory) remain opaque to her; she is offered no immediate explanation as to *how* such physical equivalence is to be construed, or how it could even be said to arise.

These considerations naturally suggest a second possible line of response for the defender of the interpretational view. In particular, she might claim that she *does*, in fact, have a transparent understanding of the reality underlying NGT’s symmetry-related models, and that such a transparent understanding is in fact attainable *prior* to the move to Newton-Cartan theory.¹⁹

Such a response evidently leads into deep philosophical waters very quickly. (After all, what does it mean to be in possession of a “transparent understanding” of anything?) But let me make a brief remark as to why I find this particular claim to be implausible. For note that in NGT the persisting points of absolute space are not merely “idly turning wheels” that can simply be expunged from the theory without explanatory loss: they are not “explanatorily idle” posits. This is for two main reasons. First, such points play a crucial role in the *metaphysical* explanation of what quantities like relative velocity and absolute rotation and absolute acceleration truly are: for the Newtonian, facts about particular inter-particle velocities and absolute rotations and absolute accelerations are naturally understood as being *grounded in* particular facts about (rates of change of) absolute velocities.²⁰ Second, such points provide the crucial transtemporal standard which is required in the realist’s *causal* explanation of the observable effects of noninertial motion (e.g., Newton’s famous “bucket experiment”): a standard without which Newton’s laws simply cannot be formu-

¹⁹Dewar (2015, esp. 322)—who is a recent, explicit defender of the interpretational view—is plausibly read as making this claim.

²⁰Cf. Pooley (MS, 118).

lated (at least, absent any *other* way of construing the transtemporal structure required to underwrite the distinction between inertial and noninertial motion). In short—and to the extent that the interpretational view is not supposed to reduce to a rather uninteresting form of scientific instrumentalism—it is simply not clear what causal-explanatory, *realistic* picture of the world is being propounded by the defender of the interpretational view, at least in this particular (Newtonian) context; it is simply opaque what, according to her, *the world is really like*.

Acknowledgements

For extremely helpful comments and discussion, I would like to thank Neil Dewar, James Ladyman, Niels Martens, Tushar Menon, Oliver Pooley, James Read, Simon Saunders, Alex Skinner, Teru Thomas, David Wallace, and audiences in London and Cardiff.

References

- Dasgupta, S. (forthcoming), “Symmetry as an Epistemic Notion (Twice Over).” *British Journal for the Philosophy of Science*.
- Dewar, N. (2015), “Symmetries and the Philosophy of Language.” *Studies in the History and Philosophy of Modern Science*, Vol. 52, pp. 317-327.
- Dirac, P. A. M. (1930), *The Principles of Quantum Mechanics*. Oxford University Press. (Reference is made to 1958 (4th) edition.)
- Earman, J. (1989), *World-Enough and Space-Time*. MIT Press.
- Greaves, H. and Wallace, D. (2014), “Empirical Consequences of Symmetries.” *British Journal for the Philosophy of Science*, Vol. 65, No. 1, pp. 59-89.
- Healey, R. (2007), *Gauging What’s Real*. Oxford University Press.
- Knox, E. (2014), “Newtonian Spacetime Structure In Light of the Equivalence Principle.” *British Journal for the Philosophy of Science*, Vol. 65, No. 4, pp. 863-880.
- Ladyman, J. (2007), “Scientific Structuralism: On the Identity and Diversity of Objects in a Structure.” *Aristotelian Society Supplementary Volume*, Vol. 81, No. 1, pp. 23-43.
- Maudlin, T. (1993), “Buckets of Water and Waves of Space: Why Spacetime Is Probably a Substance.” *Philosophy of Science*, Vol. 68, No. 2, pp. 183-203.
- Nozick, R. (2001), *Invariances: The Structure of the Objective World*. Harvard University Press.
- O’Leary-Hawthorne, J. and Cortens, A. (1995), “Towards Ontological Nihilism.” *Philosophical Studies*, Vol. 79, No. 2, pp. 143-165.
- Pooley, O. (2013), “Substantivalist and Relationalist Approaches to Spacetime.” In R. Batterman (ed.), *Oxford Handbook of Philosophy of Physics*. Oxford University Press.
- Pooley, O. (MS), *The Reality of Spacetime*. Book manuscript.

Saunders, S. (2003), "Physics and Leibniz's Principles." In K. Brading & E. Castellani (eds.), *Symmetries in Physics: Philosophical Reflections*. Cambridge University Press.

Saunders, S. (2007), "Mirroring as an A Priori Symmetry." *Philosophy of Science*, Vol. 74, No. 4, pp. 452-480.

Weatherall, J. (forthcoming). "Understanding Gauge." *Philosophy of Science*.

Weyl, H. (1952), *Symmetry*. Princeton University Press.