



Invariance or equivalence: a tale of two principles

Caspar Jacobs¹

Received: 8 February 2021 / Accepted: 7 May 2021
© The Author(s) 2021

Abstract

The presence of symmetries in physical theories implies a pernicious form of underdetermination. In order to avoid this theoretical vice, philosophers often espouse a principle called *Leibniz Equivalence*, which states that symmetry-related models represent the *same* state of affairs. Moreover, philosophers have claimed that the existence of non-trivial symmetries motivates us to accept the *Invariance Principle*, which states that quantities that vary under a theory's symmetries aren't physically real. Leibniz Equivalence and the Invariance Principle are often seen as part of the same package. I argue that this is a mistake: Leibniz Equivalence and the Invariance Principle are orthogonal to each other. This means that it is possible to hold that symmetry-related models represent the same state of affairs whilst having a realist attitude towards variant quantities. Various arguments have been presented in favour of the Invariance Principle: a rejection of the Invariance Principle is *inter alia* supposed to cause indeterminism, undetectability or failure of reference. I respond that these arguments at best support Leibniz Equivalence.

Keywords Symmetries · Invariance principle · Leibniz equivalence · Reduction · Sophistication · Supervaluation · Underdetermination · Indeterminism

1 Introduction

The presence of symmetries in physical theories implies a pernicious form of underdetermination. It is widely agreed that at least in some cases symmetry-related models are empirically equivalent: the possibilities they represent are observationally indiscernible.¹ For example, models related by 'shifts'—uniform translations of all the universe's matter content—are empirically equivalent because they agree on all

¹ See Roberts (2008), Wallace (2019). It is controversial whether this is true for all symmetries; for dissent, see Belot (2013), Dasgupta (2016). Middleton and Murgueitio Ramírez (2020) have argued that variant quantities *are* detectable, but see Jacobs (2020) for a response.

✉ Caspar Jacobs
caspar.jacobs@philosophy.ox.ac.uk

¹ Magdalen College, University of Oxford, High Street, Oxford OX1 4AU, USA

observable facts about distances, relative velocities and accelerations. In addition to underdetermination, *local* symmetries also seem to imply a particularly problematic form of indeterminism. Theories with local symmetries possess pairs of models that agree on all facts up until some time t , but diverge thereafter. The Hole Argument in General Relativity is a famous example of this.²

In order to avoid these theoretical vices, philosophers often espouse a principle called *Leibniz Equivalence*, which states that symmetry-related models (invariably) represent the *same* state of affairs. With Leibniz Equivalence, there simply is no underdetermination: symmetry-related models are just different ways of describing the same possibility. In the case of shifts, this means that shift-related models really represent the same state of affairs: shifts are ‘distinctions without a difference’. Moreover, philosophers have claimed that the existence of non-trivial symmetries motivates us to accept the *Invariance Principle*, which states that quantities that vary under a theory’s symmetries are physically unreal. The Invariance Principle codifies so-called ‘symmetry-to-(un)reality inferences’: inferences from the variance of some quantity to its non-reality.³ For example, (absolute) positions vary under shifts, and hence the Invariance Principle rejects them as unphysical.

Leibniz Equivalence and the Invariance Principle are often seen as part of the same package (Sect. 3). I believe that this is a mistake: Leibniz Equivalence and the Invariance Principle are orthogonal. But it is no surprise that this fact has gone unnoticed, for the principles *are* equivalent if one assumes that a quantity’s values always represent the same magnitude. I understand ‘values’ here as mathematical entities that represent physical properties (Sect. 4). Call this claim the *Value-Magnitude Link*. If we accept the Value-Magnitude Link, Leibniz Equivalence does entail the Invariance Principle. We are then naturally led to *reductionism*, which aims to formulate ‘reduced’ theories in terms of these invariant quantities. However, it is also possible to reject the Value-Magnitude Link. Indeed, a popular position in the philosophy of spacetime called *sophisticated substantivalism* does just that. If we reject this link, Leibniz Equivalence does not entail the Invariance Principle. This means that it is possible to take symmetry-related models to represent the same state of affairs whilst having a realist attitude towards variant quantities. In other words, Leibniz Equivalence and the Invariance Principle can come apart. *Sophistication*, of which sophisticated substantivalism is an instance, aims to fill this space.⁴ In Sect. 5, I explain the nature of sophistication in terms of a method of supervaluation over symmetry-related states.⁵

This means that those who accept Leibniz Equivalence as a solution to the problems of underdetermination and indeterminism face a choice: reject the Value-Magnitude Link, or accept it and hence commit to the Invariance Principle? Various arguments

² See Earman and Norton (1987), Wallace (2002).

³ For more on the symmetry-to-reality inference, see Dasgupta (2016), Read and Møller-Nielsen (2020).

⁴ For more on the reduction/sophistication distinction, see Dewar (2019), Martens and Read (2020).

⁵ This naturally raises the question: is there also a coherent position that denies Leibniz Equivalence yet affirms Invariance Principle? It has been suggested to me in conversation by Noel Swanson that Belot’s (2011) ‘modal relationism’ might fulfil this role. As a relationist, Belot accepts the Invariance Principle: only the invariant distances are physically real. But Belot’s relational facts are endowed with rich modal structures that allow for *more* possibilities than standard relationism, *contra* the collapse of possibilities entailed by Leibniz Equivalence.

have been presented in favour of the Invariance Principle: a rejection of the principle is supposed to cause indeterminism (Earman and Norton 1987), undetectability (Dasgupta 2016) and failure of unique reference (Healey 2009; Caulton 2015). As I will argue in Sect. 6, however, those arguments at best support Leibniz Equivalence. The inference to the Invariance Principle is unwarranted since this additionally requires that one assumes the Value-Magnitude Link, which proponents of sophistication reject. Therefore, these arguments cannot offer support for reductionism, contrary to the intentions of their authors.

2 Two principles

In this section, I discuss Leibniz Equivalence and the Invariance Principle in turn. Throughout the paper, I will adopt a loose definition of symmetries as transformations that in some sense ‘preserve the dynamics’. In order to explain this notion, it is useful to distinguish between two types of models. The *kinematically possible models* (KPMs) of a theory are those that are of the correct ‘form’, or contain the right sort of mathematical objects. The *dynamically possible models* (DPMs) are those KPMs that in addition satisfy a set of equations of motion. The DPMs represent ways the world could be if the theory were true. Symmetries are maps from the space of KPMs onto itself that preserve the space of DPMs. Since Belot (2013), it has been well-known that not all symmetries of this type relate empirically—let alone physically—equivalent states. For this reason I restrict discussion to just those symmetries that in addition satisfy a condition of empirical equivalence. The main symmetries of philosophical interest, from ‘shifts’ in Newtonian Mechanics and diffeomorphisms in General Relativity to the gauge symmetries of electrodynamics, satisfy this requirement.

2.1 Leibniz equivalence

Leibniz Equivalence was first suggested by Earman and Norton (1987) as a response to the Hole Argument, which states that General Relativity contains pairs of models related by diffeomorphisms that agree on all physical fact up until some time t but diverge thereafter. Because of this fact, it seems that substantialist interpretations of General Relativity are indeterministic. Earman and Norton (1987, p. 522) phrase Leibniz Equivalence as the claim that *diffeomorphic models represent the same physical situation*. With this assumption, General Relativity is *not* indeterministic, since the symmetry-related models that seem to represent diverging futures are in fact merely distinct representations of the same future. Since then, the term has been used more broadly to refer to the view that *any* symmetry (subject to the proviso of empirical equivalence) relates physically equivalent models, whether or not the symmetry in question is a diffeomorphism. I will understand Leibniz Equivalence in this broader sense:

Leibniz Equivalence: Symmetry-related models are physically equivalent.

The aforementioned shifts are a paradigmatic example: shifts preserve all distances between bodies and so relate empirically equivalent states of affairs. From Leibniz Equivalence, it follows that shift-related models are also physically equivalent: they merely represent the same state of affairs in different ways. Proponents of Leibniz Equivalence include Saunders (2003b), Baker (2010), Weatherall (2018) and Greaves and Wallace (2014). The latter, for example, 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.

This is not to say that there is no dissent; see, for instance, Roberts (2020). But it seems clear that Leibniz Equivalence is the orthodoxy with respect to symmetry-related models.

Rynasiewicz (1994) calls the negation of Leibniz Equivalence *model literalism*. According to model literalism (or literalism, for short), distinct models represent distinct states of affairs, whether or not they are symmetry-related. In other words, there is a one-to-one relation between models and possibilities. As we will see below (Sect. 6), in the presence of non-trivial symmetries literalism implies exactly those vices that Leibniz Equivalence is supposed to solve. I will argue that the denial of literalism, and hence acceptance of Leibniz Equivalence, suffices to avoid these: the further posit of the Invariance Principle is unnecessary.

2.2 The invariance principle

The *Invariance Principle* is the claim that “only quantities that are invariant under the symmetries of our theories are physically real” (Møller-Nielsen 2017, p. 1253). In other words, only those quantities that have the same value across equivalence classes of symmetry-related states represent a feature of physical reality. Put differently:

Invariance Principle. A quantity is physically real only if it is invariant under the symmetries of our theories.

Examples of variant quantities include absolute position under shifts, but also the values of the electrostatic potential in electromagnetism under gauge transformations, and (arguably) intrinsic mass under mass scalings.⁶ Examples of invariant quantities are distances and relative velocities, so-called ‘loop holonomies’ of the potential field, and mass ratios. The Invariance Principle suggests that the latter set of quantities is more fundamental than the former.

The Invariance Principle has gathered assent from a broad range of physicists and philosophers throughout the twentieth and twenty-first century. Here is a collection of representative quotes:

⁶ The latter issue is controversial: see Dasgupta (2013) for the case for comparativism, and Baker (2014), Martens (2019) for responses.

Physically real quantities are invariant under exact symmetries—this is the general lesson. (Saunders 2003b, p. 300)

[O]nce we possess the covariant representation under which the equations stay the same for all coordinate systems, the quantities in the (covariant) equations are the real and objective quantities. (Nozick 2001, p. 82)

[T]he important things in the world appear as the invariants [...] of these transformations. (Dirac 1930, p. 7)

Endorsements of the Invariance Principle can also be found in Saunders (2007), Baker (2010), Dewar (2015), Caulton (2015) and Dasgupta (2016), among others. The Invariance Principle is perhaps even more orthodox than Leibniz Equivalence.

The Invariance Principle encodes the idea of so-called *symmetry-to-reality inferences*: the inference from the variance of a quantity under the theory's symmetries to its physical unreality. Dasgupta (2016) justifies such inferences on the basis that variant quantities are undetectable. It then follows from an application of Occam's razor that theories that excise these quantities are preferable, all else being equal. But as I will argue below (Sect. 6), realism about such quantities does not entail empirical underdetermination. Therefore, Occam's razor does not justify the Invariance Principle, *pace* Dasgupta.

3 One principle?

Leibniz Equivalence and the Invariance Principle are closely related. Nevertheless the two principles *are* distinct, a fact which has often gone unnoticed.⁷ For example, Saunders (2007, p. 453) equivocates between them in the following passage:

[Relationalism] says that *only quantities invariant under exact symmetries are real*—thus relative directions, relative distances, and so on, under rotations and translations, etc. [...]. Call [this principle] the invariance principle. The distinctions among representations then correspond to nothing physically real. *Equivalently, such models, as goes their physical content, can be simply identified.* (emphasis mine)

The first italicised part of this passage is a canonical statement of the Invariance Principle. But towards the end of the quote Saunders equates this with Leibniz Equivalence, the claim that symmetry-related models are physically equivalent. Similarly, Baker (2010) writes:

[I]f changes in surplus structure are generally (as in geometry) mere descriptive changes, it follows that physical situations related by symmetries must be qualitatively identical [Leibniz Equivalence]. And if this is right, then physical quantities that change under symmetry transformations (i.e. that are not invariant) must not be fundamental quantities [Invariance Principle].

⁷ Read and Møller-Nielsen (2020, Sect. 6) are an exception: they rightly point out that these principles are not logically equivalent.

In this passage Baker, like Saunders, moves freely from Leibniz Equivalence to the Invariance Principle.

As far as I am aware, however, there is no explicit argument for the claim that Leibniz Equivalence implies the Invariance Principle to be found in the literature. It seems plausible that the following line of reasoning is implicitly assumed. Let m and m' denote symmetry-related models of a theory T , such that there is a variant quantity Q which takes on different values across these models. For instance, m and m' could denote shift-related models of Newtonian Gravitation, with Q the (variant) position of the Earth. Then:

- (1) The models m and m' represent the same physical possibility (Leibniz Equivalence);
 - (2) Distinct models represent the same physical possibility iff they agree on all physical features;
 - (3) Therefore, m and m' agree on all physical features (from (1) and (2));
 - (4) The models m and m' disagree on the value of Q ;
 - (5) Therefore, if Q is physically real then m and m' disagree on some physical feature (from (3) and (4));
- (C) Conclusion: Q is not physically real (Invariance Principle).

This argument seems valid, but in fact it is not: (5) does not follow from (3) and (4). The reason is that a quantity's values are *representative devices*, rather than physical features. Therefore, the fact that two models disagree on those values does not imply that they disagree on any physical fact. For it is possible that the values of Q represent *different* magnitudes across models. In particular, it is possible that Q 's distinct values in m and m' denote the *same* property. In that case, the two models are in agreement on all physical features after all, so the Invariance Principle does not follow.

In order to turn the above into a valid argument, an assumption that connects a quantity's values with the magnitudes they are supposed to represent is required. I will discuss such a principle in the next section.

4 The value-magnitude link

The above argument is valid *if* we assume the following principle, which I call the *Value-Magnitude Link*:

Value-Magnitude Link. The values of a quantity invariably represent the same magnitude across models.

I am not aware of any explicit statement of the Value-Magnitude Link in the literature, let alone a defence. But in light of the quotes above, I contend that the Value-Magnitude Link is a widely-shared implicit assumption in the symmetries literature.

It is important to be clear on what is meant by *values* here. In the first instance, it is tempting to think of values as *numbers*. For example, the mass value of a 5 kg object is '5'. But on this definition the Value-Magnitude Link is obviously false, since the same magnitude may be represented by different numbers in different unit systems. For example, we can also represent '5 kg' as '5000 g'. In other words,

the Value-Magnitude Link clearly fails for *passive transformations*. When I write ‘value’, then, I don’t mean the numerical value of a quantity in a certain system of units. Rather, I mean a mathematical element of the theory’s models: a particular point of a differentiable manifold, or an element of an internal ‘value space’. For example, a model of Newtonian Gravitation consists of an assignment of massive particle trajectories to a manifold M , over which some spacetime structure is defined. We can endow the manifold M with coordinates: a passive transformation then is simply a change in these coordinates. But we can also leave the coordinates fixed, and assign the particle trajectories to different points of the manifold itself. This is an *active transformation*.⁸ The values in this case are the point of the manifold M , and the Value-Magnitude Link asserts that the same point of M represents the same point of spacetime in all of the theory’s models. We can present a similar set-up for quantities such as mass. Here, quantities are functions from the theory’s domain—say, particles—into a *value space*: a mathematical structure whose elements represent determinate magnitudes of said quantity. For instance, mass is a function from particles into a mass value space, whose elements represent determinate mass magnitudes such as ‘5 kg’. As with spacetime, we can ‘coordinatise’ a value space via an assignment of numbers to its elements. This amounts to a choice of unit. But we can also leave the units fixed, and consider a different mapping from the domain into value space. For instance, an (active) mass scaling maps each particle onto a different element of mass value space. In this case, values are simply elements of a value space, and the Value-Magnitude Link asserts that these elements represent the same magnitudes across the theory’s models. Finally, I embrace realism about values: they represent real aspects of the physical world, such as positions or masses (in the case of position, this entails substantialism).⁹

With this clarification in mind, let’s consider the Value-Magnitude Link in more detail. In the previous section, I said that Leibniz Equivalence implies the Invariance Principle *if* we assume the Value-Magnitude Link. In fact, these two principles are *equivalent* on that assumption, as I will now show.

Proof Left to right. Consider a variant quantity Q . Suppose that in some model m_1 , the values of Q for x and y are $Q(x) = a$ and $Q(y) = b$. Now, consider a symmetry-related model m_2 in which $Q(x) = b$ and $Q(y) = c$.¹⁰ By Leibniz Equivalence, m_1 and m_2 are physically equivalent, so $Q(x)$ and $Q(y)$ each represent the same magnitude in both models. But since $Q(y) = b$ in m_1 while $Q(x) = b$ in m_2 , this means that the value b represents *distinct* magnitudes across those models. This contradicts the Value-Magnitude Link, so Q cannot represent a physically real quantity. Since Q is an arbitrary variant quantity, it follows that only invariant quantities are physically real.

For an example, consider again the static shift. Since static shifts are symmetries of Newtonian Gravitation, Leibniz Equivalence implies that shift-related models are

⁸ Note that on this definition an active transformation need not correspond to any *physical* transformation: if shift-related models of Newtonian Gravitation are physically equivalent, then an active transformation relates models that represent the same state of affairs.

⁹ For more on the metaphysics of value spaces, see Arntzenius and Dorr (2012), Eddon (2013), Wolff (2020) and references therein.

¹⁰ This relies on the assumption that there exists some symmetry-related model in which $Q(x) = b$; see the shift example below.

physically equivalent. In that case, the Invariance Principle implies that position, which varies under shifts, is *not* a physically real quantity. For consider the following two shift-related models, where x_A and x_B denote the locations of particles A and B respectively:

$$\mathfrak{M}_1 : \begin{cases} x_A = (1, 0, 0) \\ x_B = (2, 0, 0) \end{cases} \quad \mathfrak{M}_2 : \begin{cases} x_A = (2, 0, 0) \\ x_B = (3, 0, 0) \end{cases}$$

If these models are physically equivalent, then $(1, 0, 0)$ and $(2, 0, 0)$ both represent the same spacetime point: A 's location. Similarly, $(2, 0, 0)$ and $(3, 0, 0)$ both represent B 's location. And this means that the same value, $(2, 0, 0)$, represents A 's location in \mathfrak{m}_1 , and B 's location in \mathfrak{m}_2 . But both models represent those bodies as occupying *distinct* locations. It follows that the same values of x_A and x_B must denote different spatial locations across models. This contradicts the Value-Magnitude Link, and so variant location is not physically real.

Right to left. Consider any pair of symmetry-related models. By definition, the invariant quantities have the same values across these models. The Invariance Principle implies that only those quantities are physically real. From the Value-Magnitude Link, it follows that their values denote the same magnitudes across models. Therefore, these models agree on all physical features. Symmetry-related models are thus physically equivalent, and so Leibniz Equivalence holds. For example, consider once more two shift-related models. From Invariance Principle, it follows that only the invariant quantities—distances, relative velocities and accelerations—are physically real. By definition, those quantities have the same value in both models. From the Value-Magnitude Link, it follows that those values represent the same magnitude. Therefore, both models represent the same state of affairs. \square

On the assumption of the Value-Magnitude Link, then, Leibniz Equivalence and the Invariance Principle are equivalent. If that is the case, then “symmetry-to-reality” inferences have a clear structure: symmetry-related models are physically equivalent, and hence variant quantities are unphysical. This naturally leads to the view that we ought to formulate new theories that are solely expressed in terms of the invariant quantities, so-called *reduced theories*. On the other hand, if one were to insist that variant quantities *are* physically real, then under the assumption of the Value-Magnitude Link this entails the denial of Leibniz Equivalence, or literalism. So, on the assumption of the Value-Magnitude Link the interpretation of symmetry-related models leads to a dilemma: either variant quantities are physically real and symmetry-related models represent distinct states of affairs (literalism), or only invariant quantities are real and so symmetry-related models are physically equivalent. In the latter case, a new theoretical formalism is appropriate (reduction). There is no room for an intermediate view. But as we will see in the next section, there is no need to accept the Value-Magnitude Link. If we do not, another option becomes available.




Fig. 1 Examples of shift-related models

5 Conventional values and supervalueation

The Value-Magnitude Link seems unassailable, but is in fact denied by one of the most common views on symmetries: sophistication. In this section, I will first illustrate sophistication with a couple of examples: sophisticated spacetime substantivalism, and sophistication for the electrostatic potential. The fact that the latter example concerns an internal symmetry shows that the failure of the Value-Magnitude Link is not merely a particularity of spacetime. After the examples, I will undertake a more general discussion of the relation between sophistication, reduction and literalism.

The most well-known example of sophistication is *sophisticated substantivalism*.¹¹ According to sophisticated substantivalism, spacetime is a real substance in which bodies are located; hence, there is a sense in which (variant) positions are physically real.¹² Contrast this with relationism, which says that only distances between bodies are real. Despite the fact that sophisticated substantivalism is realist about position, it affirms the physical equivalence of symmetry-related models. It does so via an appeal to *anti-haecceitism* regarding spacetime points. Rather than possessing primitive identities, spacetime points are individuated via their qualitative relations to each other and to the universe's matter content. Since spacetime symmetries relate qualitatively identical models, symmetry-related models represent the exact same points in the same configurations. As a result, such models have identical physical content.

Therefore, sophisticated substantivalism accepts Leibniz Equivalence and rejects the Invariance Principle. It follows from the previous section that sophistication must also reject the Value-Magnitude Link. This is indeed the case. Here is a simple example. Suppose as before that models related by shifts are physically equivalent. The following three diagrams then represent the same state of affairs:

Here, the dots represent the locations of three bodies labelled 'a', 'b' and 'c' in two-dimensional position space. The location value of the dots is different in each diagram (the dots are shifted on the x -axis). On the assumption of the Value-Magnitude Link, this means that spatial location itself is not a physically real quantity, since it varies

¹¹ See Pooley (2013, Sect. 7) for an overview.

¹² The claim that sophisticated substantivalism is realist about (variant) position assumes a particular form of that view, namely what Saunders (2003a) calls *non-reductive relationism*. On the other hand, an *eliminative relationist* may understand positions as non-fundamental quantities. To confuse matters, both types of view have sometimes been called 'structuralism'; cf. Greaves (2011). I will eschew the latter term in what follows.

across the theory's shift-symmetries. But according to sophistication there *are* physically meaningful statements about spatial position. Consider the claim that particle *b* has *some* location in space (i.e. *b* is assigned *some* value in two-dimensional position space). This statement is invariant under shifts: *b* is assigned a location on all three diagrams in Fig. 1. Robert Stalnaker has summarised the point as follows: “in the relational theory of space, spelled out in this way, no intrinsic spatial properties would be real (*except the property of being located somewhere in space*)” (Stalnaker 1979, pp. 353–354, emphasis mine). This, of course, is a direct consequence of the sophisticated substantialist's commitment to space as a real substance in which particles are located.

Furthermore, sophistication is not just applicable to spacetime symmetries. Dewar (2019) has recently argued that sophistication can be seen as a general approach to the interpretation of symmetries. He mentions the example of the electrostatic potential, ϕ . The empirical consequences of ϕ are due to its gradient $E := \nabla\phi$. Therefore, uniform shifts of ϕ are symmetries of electrostatics. According to Leibniz Equivalence, then, models related by shifts in the potential are physically equivalent. It simply makes no difference if we increase ϕ everywhere by the same amount. The Invariance Principle in turn implies that the electrostatic potential is not physically real, since it is variant under the symmetries of electrostatics. Instead, the electric field E , which *is* invariant, represents the theory's fundamental field. This motivates the demand for a reduced theory formulated in terms of E alone. But as Dewar points out, such a reduced theory cannot explain the fact that E is a *conservative* force, i.e. that $\nabla \times E = 0$. On the other hand, if E is *defined* as the gradient of ϕ then this conservativeness rolls out as a mathematical theorem, since $\nabla \times (\nabla\phi) = 0$ for any ϕ . For this reason, Dewar argues that it is better to affirm the physical reality of ϕ , contrary to the Invariance Principle. This implies a rejection of the Value-Magnitude Link. For consider two symmetry-related models, one in which $\phi(x) = a$ and one in which $\phi(x) = b$. When $a \neq b$, the value of ϕ at x differs across these models. But since symmetry-related models are physically equivalent, both models must represent the same physical *magnitude* of the electrostatic potential at x . Hence, it is not the case that values invariably represent the same magnitude across models. Dewar's sophistication achieves this balance via an appeal to *anti-quidditism*, the analogue of anti-haecceitism for properties: physical properties are individuated via their position in a structure of qualitative relations. Since the distinct values a and b occupy the same structural role across symmetry-related models, they also represent the same magnitude.¹³

These examples illustrate how sophistication steers a course between equivalence and invariance. But in order to obtain a more general characterisation of the difference between sophistication and reduction, it is necessary to provide a more specific criterion for the physical content of a sophisticated theory. In other words, what we need is a statement of the *semantics* of a sophisticated theory.¹⁴ I propose the following criterion:

¹³ For more on anti-quidditism in the context of symmetries, see Martens and Read (2020).

¹⁴ This is different from the *formalism* of a sophisticated theory, which is what Dewar (2019) is chiefly concerned with. As Dewar argues, sophistication requires that a theory's symmetry-related models are *isomorphic*. On the 'internal' approach to sophistication, this usually requires a reformulation of the theory's models.

Sophistication. A theoretical statement is *physically meaningful* if and only if it is invariant across physically equivalent models.¹⁵

From left to right, this says that a statement is physically meaningful only if it is invariant across physically equivalent models. This is uncontroversial: physically equivalent models agree on all physical facts, so whatever varies across such models cannot have a consistent meaningful interpretation. But read from right to left, this claim is non-trivial: it states that *any* invariant statement is physically meaningful, *even if* such a statement involves variant quantities. In other words, sophistication locates physical meaning in the full sentence, rather than in its terms. For an example, borrowed from Van Fraassen (1989, p. 284), consider the statement “total momentum is conserved”. In Newtonian Gravitation, total momentum is a variant quantity: it varies under boosts of our reference frame. But it is true in any frame that the total momentum of all particles is conserved, and hence the statement “total momentum is conserved” is still physically meaningful under a sophisticated semantics.

We can also understand this semantic criterion in terms of the logic of supervaluation. This will provide us with another, more formal way of understanding the space between literalism and reduction. In formal semantics, the idea of supervaluation is to consider the truth-value of a sentence with respect to a class of interpretations, rather than within any particular interpretation. In the present context, the equivalent idea is to evaluate the truth-value of a sentence of the theory with respect to an equivalence class of symmetry-related models, rather than within any particular model. Thus, we call a statement *supertrue* (respectively *superfalse*) iff it is true (respectively false) in all models within a class of symmetry-related models.¹⁶ Importantly, *supervaluation does not commute with logical composition*. This means that the truth-value of a composite statement under supervaluation is not a function of the truth-values of its components under supervaluation. We can use this fact to explain the difference between reduction and sophistication: the former *first* supervaluates over ‘atomic’ theoretical statements and then composes complex ones, whereas the latter first composes complex sentences and *then* supervaluates. Let $[\phi(m)]_{\mathcal{S}}$ denote the truth value of a sentence ϕ in a model m when we supervaluate it with respect to a group of symmetries \mathcal{S} (compare the role of *interpretations* in logic). The truth conditions of $[\phi(m)]_{\mathcal{S}}$ are:

$$[\phi(m)]_{\mathcal{S}} = \begin{cases} T & \text{iff } \phi(s(m)) = T \text{ for all } s \in \mathcal{S} \\ F & \text{iff } \phi(s(m)) = F \text{ for all } s \in \mathcal{S} \\ \# & \text{otherwise} \end{cases} \quad (1)$$

where $s(m)$ denotes the model obtained from an application of s to m . In simple terms, this says that ϕ is true (false) in m iff it is true (false) in all models related to m by one of the symmetry transformations in \mathcal{S} . If neither situation obtains—if ϕ is true in some model, but false in another, symmetry-related one—then the statement is meaningless.

¹⁵ Cf. Mundy (1986). See also Barrett’s (2017) proof that any piece of symmetry-invariant structure is *implicitly definable* from the theory’s models.

¹⁶ Cf. Dewar’s (2019) supervaluationist semantics for sophistication. See also Russell’s (2018) discussion of determinacy.

For instance, the claim that a certain body is at absolute rest is true in some models of Newtonian Gravitation, but false in any model related to the first by a uniform boost. Therefore, on a supervaluationist semantics it is physically meaningless. But the claim that total momentum is conserved is true in all models of the theory, so comes out as true under the supervaluationist semantics.

For an example of a statement that is meaningful (indeed, true) for sophistication but meaningless for reduction, consider the claim that a quantity Q has *some* value in m : $\exists x(Q(m) = x)$. For example, Q can stand for the Earth's position, so the claim is that the Earth is located somewhere in space. We can write this claim as a disjunction:¹⁷

$$Q(m) = x_1 \vee \cdots \vee Q(m) = x_n \quad (2)$$

The crucial insight here is that there is a difference in truth value between the supervaluation of this sentence as a whole, and the composition of its individually supervaluated disjuncts. In other words,

$$[Q(m) = x_1 \vee \cdots \vee Q(m) = x_n]_S \not\leftrightarrow [Q(m) = x_1]_S \vee \cdots \vee [Q(m) = x_n]_S \quad (3)$$

Proof Since Q varies under S , for any x the claim $[Q(m) = x]_S$ is neither true nor false. Under standard trivalent semantics (such as Kleene), the disjunction of sentences with truth-value # itself has truth-value #. Therefore, the rhs of (3) is neither true nor false. On the other hand, $[Q(m) = x_1 \vee \cdots \vee Q(m) = x_n]_S$ *does* have a truth-value, since in any model Q has *some* value between x_1 and x_n . Therefore, the lhs of (3) is true. This completes the proof. \square

It is because of this fact that sophistication is different in terms of its physical content from both literalism and reduction. Using the example of the Earth's location, the difference is as follows. Reduction says that no statement of the form "The Earth is located at position x " is meaningful, since no such statement has the same truth-value within a class of shifted models. Therefore, the reductionist argues, *composites* of such statements are also meaningless. In particular, the sentence "The Earth has *some* location"—which we have interpreted as a disjunction of ascriptions of particular locations—has no physical relevance. But sophistication evaluates the truth-value of sentences *as a whole*. And as we have seen, the sentence "The earth has some location" *is* true across shift-related models. Therefore, the semantics of supervaluation offers a formal way of understanding the sense in which sophistication has *more* physical content than reduction.

Although it is outside of the scope of this paper, let me briefly say something about the benefits of this additional content. The rough idea is that the theory's variant quantities contribute towards its *explanatory strength*. Consider, for example, the fact that the electric force is conservative, i.e. that $\nabla \times E = 0$. On a sophisticated semantics, the definition $E := \nabla\phi$ is invariant across symmetry-related models, hence physically meaningful. Therefore, we can use it to *explain* the conservativeness of E as a result of the fact that $\nabla \times \nabla\phi = 0$ for any scalar field ϕ . On a reductionist semantics, on

¹⁷ I will gloss over any issues arising from the fact that for most quantities, this disjunction is infinitary.

the other hand, *any* statement that involves ϕ is physically meaningless as a result of the fact that ϕ itself varies under the theory's symmetries. Therefore, reductionism cannot appeal to ϕ to explain the conservativeness of E . For reductionism, it is a brute fact that E is conservative. Generally, what this means is that sophistication but not reduction allows for explanations of physical phenomena in terms of variant quantities. In fact, such explanations abound in physics. For just one example, we can explain the Twin Paradox in Special Relativity in terms of the phenomenon of time dilation.¹⁸ But this is a frame-dependent (hence symmetry-variant) phenomenon, and so reductionism accords it less reality than sophistication. Here, too, it seems that sophistication increases the explanatory strength of the theory. I'll briefly return to these ideas in the conclusion.

6 Against invariance

If one accepts the Value-Magnitude Link, one faces a stark binary choice between literalism and reduction. On the first view, symmetry-related models represent distinct possibilities: the consequence is underdetermination and possibly indeterminism. On the second view, symmetry-related models are physically equivalent. With the Value-Magnitude Link in place, this then implies the Invariance Principle: only invariant quantities are physically real. It is this principle that motivates the search for a reduced theory formulated solely in terms of these invariant quantities. Clearly, reduction is preferable over literalism. But if we reject the Value-Magnitude Link a third option becomes available: sophistication. Sophistication accepts Leibniz Equivalence, but not the Invariance Principle. As I illustrated above, this means that it remains committed to the physical reality of variant quantities. The possibility of sophistication shows that Leibniz Equivalence in itself cannot motivate the Invariance Principle.

However, there are many other arguments in the literature in favour of the Invariance Principle. Earman and Norton (1987) argue that the Invariance Principle is necessary to avoid radical indeterminism; Dasgupta (2016) claims that variant quantities are undetectable; and according to Caulton (2015), such quantities fail to refer uniquely. These arguments, if successful, *a fortiori* support reduction over sophistication. But I will argue that they fail. Specifically, I claim that all three arguments only support Leibniz Equivalence, rather than the Invariance Principle. It is only when one tacitly assumes the Value-Magnitude Link that the latter follows. But that assumption begs the question against sophistication, which explicitly rejects this link. Therefore, I conclude that these purported arguments for the Invariance Principle cannot support reduction over sophistication, contrary to their author's explicit claims.

I will now discuss the arguments in favour of the Invariance Principle one by one.

¹⁸ See, for instance, Rindler (1977, Sect. 2.14). But note that Debs and Redhead (1996) argue that such explanations are invalid exactly because they appeal to frame-dependent facts.

6.1 Determinism

The claim that symmetries lead to indeterminism of our theories is well-known. Earman and Norton (1987) argue that this happens in General Relativity due to the Hole Argument, and Belot (1998) applies a similar line of reasoning to electromagnetism. Wallace (2002) generalises the argument to any time-dependent local symmetry (for our purposes, a local symmetry is one that can act non-trivially over some finite period of time). The thrust of these arguments is that there are distinct models that agree on all physical facts before some time t , but which differ by a non-trivial symmetry transformation after t . If these models represent distinct physical possibilities, then this implies a failure of determinism: the history of the universe up to time t does not determine its state at later times.

Earman and Norton (1987) argue that we should avoid this sort of radical indeterminism not because indeterminism is in principle unacceptable, but because “[determinism] should fail for a reason of physics, not because of a commitment to substantial properties which can be eradicated without affecting the empirical consequences of the theory” (524). In a spacetime context, the ‘substantial properties’ to which Earman and Norton refer here are variant positions in spacetime. Their claim is that it is necessary to accept Leibniz Equivalence in order to avoid indeterminism, and that therefore the fundamental spatio-temporal quantities are relational rather than intrinsic. In other words, Earman and Norton advocate a form of reduction. Indeed, Earman (1989) went on to try to formulate such a reduced theory in terms of Einstein algebras, but Rynasiewicz (1992) showed that this was unsuccessful.

I agree with Earman and Norton that Leibniz Equivalence is necessary to avoid this sort of indeterminism. However, it does not follow that we have to ‘eradicate’ substantial properties, or variant quantities more generally. This is only a valid inference on the assumption of the Value-Magnitude Link. If the Value-Magnitude Link is denied, on the other hand, Leibniz Equivalence implies that we can remain realist about variant properties. In the case of spacetime, for example, sophisticated substantivalism accepts that symmetry-related models are physically equivalent, but nevertheless gives a realist interpretation of spacetime. Moreover, Dewar (2019) suggests that one can use the same strategy for non-spatio-temporal, or ‘internal’ symmetries, such as the gauge symmetry of electrodynamics. There, too, we can simply interpret gauge-related models as distinct representations of the *same* physical fields, rather than representations of distinct possible evolutions of those fields for a given initial condition. This suffices to avoid indeterminism, since the symmetry-related models that were supposed to represent distinct possible futures now considered as distinct representations of the same future. We need not accept the Invariance Principle in order to avoid radical indeterminism.

6.2 Detectability

The second argument for the Invariance Principle is epistemic. According to Dasgupta (2016), we should renounce variant quantities because they are in principle *unde-*

tectable. On his account, the symmetry-to-reality inference goes as follows (Dasgupta 2016, p. 843):

- (i) Laws L are the complete laws of motion governing our world.
- (ii) Feature X is variant under the symmetries of L .
- (iii) Therefore, X is undetectable (from (i) and (ii)).
- (C) Therefore, X is not real (from (3) and an Occamist norm that we dispense with undetectable structure).

Like Earman and Norton, Dasgupta claims that such reasoning motivates the search for an alternative theory that solely trades in the invariant quantities.

Again, I will argue that we can avoid the undetectability of variant quantities by assuming Leibniz Equivalence, which only implies the Invariance Principle if we also accept the Value-Magnitude Link. First, however, let me comment on the ‘Occamist norm’ to which Dasgupta refers. This norm states that *all else being equal* we ought to dispense with undetectable structure. The ‘all else equal’ clause is important here: Dasgupta believes that there may be good reasons to posit undetectable structure. For example, such structure may yield a simpler theory, or one with greater explanatory strength. “In that case, we would have empirical evidence of sorts that the feature is real, in the sense that our all-things-considered best empirically confirmed theory implies that it is real” (Dasgupta 2016, p. 854). Dasgupta thus believes it is possible that variant quantities *are* physically real, contrary to the Invariance Principle. Nevertheless, the spirit of Dasgupta’s argument is clearly against variant quantities. It is therefore worthwhile to show that there is a sense in which such quantities are *not* undetectable on a sophisticated view.

Here is how Dasgupta initially defines undetectability:

[S]omething is undetectable in my sense if, roughly speaking, it is physically impossible for it to have an impact on our senses. (Dasgupta 2016, p. 854)

This definition of detectability is rather weak: even absolute velocities are detectable on this definition. After all, in Newtonian mechanics relative velocities supervene on absolute velocities, and since the former impact our senses, so do the latter. For instance, suppose that we increase the absolute velocity of a ship at sea. Since this also increases the relative velocity of the ship with respect to the shore, the difference is clearly visible. Therefore, the absolute velocity of the ship has an impact on our senses. On this definition, then, the symmetry-to-reality inference is inconsequential.

However, Dasgupta quickly revises his initial definition in favour of something stronger:

It then follows [...] that the feature is undetectable [if] there is no physically possible process by which we might discover which determinate values are actually instantiated.

This is a stronger criterion, since it demands that we can measure the particular *value* of a quantity (presumably up to some degree of accuracy), rather than just its effect. Let us call this notion *measurability* (cf. Ismael and Van Fraassen’s (2003) distinction between the observable and the measurable). For instance, absolute velocity is

unmeasurable in the sense that there is no unique (numerical) value we can assign to the motion of any particular object: we can only express velocities relative to some other object. What is undetectable here are not the variant quantities themselves, but their particular values. Suppose, for example, that the current mass of the Eiffel Tower is 7 million kilograms. And suppose that there exists a physically possible world qualitatively identical to ours, but in which the masses of all bodies are doubled—hence the mass of the Eiffel tower is 14 million kilograms in that world. Since these two worlds are (by stipulation) empirically equivalent, there is no experiment that could reveal the *actual* mass of the Eiffel Tower. For such an experiment, if it were an accurate measurement of the Eiffel Tower’s mass, would have to have *different* outcomes in these worlds.¹⁹ Therefore, we cannot know, even in principle, whether the Eiffel Tower has the mass denoted by ‘7 million kilograms’ rather than the mass denoted by ‘14 million kilograms’.

If we follow Dasgupta’s Occamist advice and reject the physical reality of intrinsic mass, there simply are no mass values to measure. However, such a solution to the problem of undetectability is unnecessarily strong. For as we have seen, the problem arises when our theories imply the existence of symmetry-related worlds that are empirically equivalent yet physically distinct. This means that the issue is a failure of Leibniz Equivalence, rather than the Invariance Principle. If symmetry-related models represent the same possible worlds, then there are no differences between them, so *a fortiori* no undetectable differences. As I explained in Sect. 5, Leibniz Equivalence is consistent with the rejection of the Invariance Principle once we reject the Value-Magnitude Link. On such a view, the value ‘7 million kilograms’ does not denote the *same* mass across models. Instead, sophistication adopts an anti-quidditist picture on which physical mass magnitudes are *qualitatively identified*. Since the values ‘7 million kilograms’ in the first model and ‘14 million kilograms’ in the second instantiate the same qualitative profile (for example, the same mass ratios to external bodies), the Eiffel Tower is represented as having the same mass in both models. Sophistication implies that the Eiffel Tower has *some* mass, whereas reduction renounces intrinsic mass quantities altogether. Nevertheless, it may seem as if there is still a further question: *which* mass does the Eiffel Tower have? Now, if this is a question about the primitive identities of mass values, there simply are none on the sophisticated account—so the question is void. If, on the other hand, this is a question about the *qualitative* features of the Eiffel Tower’s mass, the answer is the same for both models: the Eiffel Tower’s mass is 100,000 times my mass, a mere fraction of the Sun’s mass, and so on. On this view, then, we *can* measure which mass the Eiffel Tower (or any other object) has, since we can measure a mass value’s qualitative features.

Therefore, sophistication avoids undetectable properties without the Invariance Principle; Leibniz Equivalence is all that’s necessary. Dasgupta’s epistemic argument does not justify the latter once we recognise the Value-Magnitude Link as a tacit but unwarranted assumption.

¹⁹ This, in a nutshell, is Roberts’s (2008) argument for the claim that variant quantities are unmeasurable.

6.3 Reference

The previous section was concerned with an epistemic argument. The final argument, found in Healey (2006) and Caulton (2015), is semantic in nature. The claim I will contest is that variant quantities cannot be given a realist interpretation because they have no unique referent.

Suppose we are trying to interpret a physical theory. Part of the aim of interpretation is to assign referents to the terms that occur in the theory in question. Some of these terms will obviously correspond to physical quantities with which we are directly acquainted, such as colour or distance. We call these *observables*. Since we already know what the observables *are*, we can simply stipulate which theoretical quantities denote them. Caulton (2015) calls this the first phase of interpretation. However, theories usually also introduce novel terms for *unobservable quantities*, such as ‘electric field’. We have no grasp on what these unobservables are apart from what the theory says about them. It seems that we cannot know what these theoretical terms refer to, unless we have *already* interpreted the theories in which they occur. This raises the worry that the interpretation of theoretical terms is an impossible mission.

Fortunately, the situation is not hopeless. Following Lewis (1970), we can construct *implicit definitions* of the theory’s theoretical terms once we have interpreted the observable terms by stipulation; this is Caulton’s (2015) second phase of interpretation. Instead of directly connecting a theoretical term to its referent, we declare that the term refers to *whatever* it is that is related to the observables in the manner described by the theory. For example, the electric field is just whatever is generated by and exerts a force on charged particles in accordance with Maxwell’s laws. For this approach to work, our theories must be *uniquely realised*. That is, there must exist in nature a unique set of unobservable quantities that the theoretical terms of our theories refer to. If this condition fails to hold, then either the theory is not realised at all (i.e. it is false), or it is *multiply realised*. In the latter case it is *indeterminate* what its theoretical terms refer to.

Now, both Healey and Caulton argue that if we interpret variant quantities as physically real, theories with non-trivial symmetries are multiply realised. Therefore, in order to avoid indeterminacy of reference we ought to embrace the Invariance Principle. But as with the two previous arguments, I believe that Leibniz Equivalence is enough to avoid a failure of reference. The further step from Leibniz Equivalence to the Invariance Principle is unwarranted. I will borrow an example from Healey (2006) to illustrate this claim. Suppose that we have a toy theory according to which subatomic particles called ‘quarks’ cluster together to form protons and neutrons. Quarks can have one of three colours: red, green and blue. Colour is a dynamically efficacious property which couples to various other quantities in the laws. However, the theory has a colour permutation symmetry, such that if one changes (for example) all red quarks to green, all green quarks to blue, and all blue quarks to red, the difference is empirically indiscernible. Finally, the dynamics of our toy theory are such that quarks are strongly confined in colour-neutral combinations of red, green and blue.

Healey argues that in this set-up the terms ‘green’, ‘red’ and ‘blue’ are referentially indeterminate. For simplicity, let’s label these properties *R*, *G* and *B*. The toy theory

has a partially interpreted model of the form $m(r, g, b)$, where r , g and b are variables that stand for the colour properties of certain quarks. Suppose that we want to implicitly define r , g and b as whatever triple of properties jointly satisfies m . Furthermore, suppose that m is satisfied if r denotes R , g denotes G , and b denotes B . In other words, $m(R, G, B)$ is a dynamically possible model. It then seems that the implicit definition of r picks out R as its unique referent, and similarly for the other terms. However, since our theory is colour permutation invariant, $m(G, B, R)$ is *also* a dynamically possible model. And this means that $m(r, g, b)$ is *also* satisfied if r refers to G , and g refers to B , and b refers to R . Therefore, an implicit definition in terms of $m(r, g, b)$ leaves it indeterminate whether r refers to R or G (or even B !), and likewise for g and b . This is a consequence of the theory's colour permutation symmetry.

Healey concludes from this that variant quantities such as quark colour are not physically real. Similarly, Caulton (2015, p. 161) writes that this “prompts appropriate reform towards a *new* formalism, in which the physical properties and relations—including the unobservable ones—are transparently represented without redundancy”. In other words, Caulton advocates reduction. However, the demand for a reduced theory is too strong: assuming Leibniz Equivalence is enough to avoid non-unique reference. Recall that the problem that causes the indeterminacy of reference of the terms r , g and b is that $m(r, g, b)$ is realised multiply by $m(R, G, B)$, $m(G, B, R)$ and $m(B, R, G)$. But if we accept Leibniz Equivalence, symmetry-related models are distinct representations of the *same* state of affairs, and hence the differences between these permuted models are merely representational. If this is the case, $m(r, g, b)$ is *not* multiply realised: it only has one physical realisation, which is redundantly represented by an equivalence class of symmetry-related models.²⁰ In other words, Leibniz Equivalence alone suffices to avoid multiple realisation, of which the indeterminacy of theoretical terms is a consequence.

On this picture, what *do* the terms r , b and g refer to? Note that these terms are *values* of the colour quantity, rather than quantities themselves. According to sophistication, then, their referent is not fixed across models: there is no unique colour property to which we can consistently refer across theoretical models with r , b or g . Rather, whether we call a colour property ‘red’, ‘green’ or ‘blue’ is a conventional choice that may vary across models. Nevertheless, it *is* possible to identify theoretical terms qualitatively. For example, we can (arbitrarily) stipulate that r refers to the colour of the quark which is located at the centre of mass of a particular proton. This is entirely analogous to the stipulation that we will use the term ‘kilogram’ to refer to the mass of the standard kilogram in Paris. On this account, there is no sense in which ‘ r ’ refers to the *same* magnitude across models. Instead, each class of symmetry-related models requires a distinct convention on how the terms ‘ r ’, ‘ g ’ and ‘ b ’ are used. But in each case, these terms have a unique referent.

²⁰ The advocate of reduction may respond that its aim is exactly to avoid such representational redundancy. I agree that that is a legitimate motivation for a reduced formalism—but note that in many cases, finding such a formalism has proved difficult if not impossible. For further discussion, see Dewar (2019, Sect. 2).

7 Conclusion

I have analysed three arguments in favour of the Invariance Principle. These arguments all start from the claim that without the Invariance Principle, several woes will befall us: indeterminism, undetectability or a failure of reference. But as I hope to have shown, all that is necessary to avoid these woes is Leibniz Equivalence. Leibniz Equivalence and the Invariance Principle are closely related, but they are not equivalent. In particular, sophistication is a view which simultaneously accepts Leibniz Equivalence and rejects the Invariance Principle. Reduction, on the other hand, accepts both Leibniz Equivalence and the Invariance Principle. But since arguments in favour of the Invariance Principle fail, the latter cannot support reduction over sophistication. Of course, this does not mean that reduction is for some reason undesirable: as far as the theoretical vices of indeterminism, undetectability and failure of reference are concerned, reduction and sophistication are exactly on a par.

What, then, could decide in favour of one or the other strategy for interpreting symmetry-related models? This is a significant question that I cannot fully answer here. But the clue lies in the ‘gap’ between reduction and sophistication. As I argued, sophistication is committed to the existence of variant quantities, even though it does not entail that such quantities have determinate values. For example, the claim “the Eiffel Tower has an intrinsic mass” is (super)true in Newtonian Gravitation, despite the fact that intrinsic mass might be a variant quantity. The benefit of sophistication is that such quantities may serve an *explanatory* purpose. For example, I mentioned earlier that the electrostatic potential ϕ can explain why the electric force E is conservative. Similarly, Martens and Read (2020) argue that the existence of intrinsic mass can explain the ‘transitivity of mass ratios’, that is, the fact that the mass ratio between two bodies a and c is equal to the product of the mass ratios between a and b and b and c , for any body b .

We can put the point differently. As we saw at the end of the previous section, theories with symmetries exhibit a certain representational redundancy: different models represent the same state of affairs.²¹ This redundancy provides one motivation for reduction, which eliminates the offending symmetries from the theory. But if the above is correct, then there is a sense in which this additional structure is *not* redundant. Instead, variant quantities such as the electrostatic potential provide the theory with explanatory resources. Far from ‘fluff’, then, such quantities are necessary for theories to discharge their explanatory duties. But in order to ‘access’ these quantities, we require the semantics of supervaluation as discussed in Sect. 5. So, sophistication offers a different way of semantically evaluating our theories which allows its ‘redundant’ structures to play an explanatory role.

For this reason, I contend that sophistication may strike the perfect balance between literalism’s excess of physical content, which leads to underdetermination, and reduction’s metaphysical sparseness, which hinders its explanatory aspirations. We ought to accept Leibniz Equivalence but reject the Invariance Principle.

²¹ Ismael and Van Fraassen (2003) argue that symmetries are “a guide to superfluous theoretical structure”; Earman (2004) calls it ‘descriptive fluff’. But Bradley and Weatherall (2020) challenge this view, arguing that theories whose symmetry-related models are isomorphic possess no surplus structure.

Acknowledgements I would like to thank Adam Caulton and James Read for valuable comments on earlier drafts of this paper. This paper was written as part of a research project funded by the British Society for the Philosophy of Science.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Arntzenius, F. & Dorr, C. (2012). Calculus as geometry. In *Space, time, and stuff*. Oxford University Press.
- Baker, D. J. (2010). Symmetry and the metaphysics of physics. *Philosophy Compass*, 5(12), 1157–1166.
- Baker, D. J. (2014). Some consequences of physics for the comparative metaphysics of quantity. <http://philsci-archival.pitt.edu/12674/>.
- Barrett, T. W. (2017). What do symmetries tell us about structure? *Philosophy of Science*, 4, 617–639.
- Belot, G. (1998). Understanding electromagnetism. *The British Journal for the Philosophy of Science*, 49(4), 531–555.
- Belot, G. (2011). *Geometric possibility*. Oxford: Oxford University Press.
- Belot, G. (2013). Symmetry and equivalence. In R. Batterman (Ed.), *The Oxford handbook of philosophy of physics* (pp. 318–339). Oxford: Oxford University Press.
- Bradley, C. & Weatherall, J. O. (2020). On representational redundancy, surplus structure, and the hole argument. [arXiv:1904.04439](https://arxiv.org/abs/1904.04439) [physics].
- Caulton, A. (2015). The role of symmetry in the interpretation of physical theories. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 52, 153–162.
- Dasgupta, S. (2013). Absolutism vs comparativism about quantity. In *Oxford studies in metaphysics: Volume 8*. Oxford University Press.
- Dasgupta, S. (2016). Symmetry as an epistemic notion (twice over). *The British Journal for the Philosophy of Science*, 67(3), 837–878.
- Debs, T. A., & Redhead, M. L. G. (1996). The twin “paradox” and the conventionality of simultaneity. *American Journal of Physics*, 64(4), 384–392.
- Dewar, N. (2015). Symmetries and the philosophy of language. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 52(Part B), 317–327.
- Dewar, N. (2019). Sophistication about symmetries. *The British Journal for the Philosophy of Science*, 70(2), 485–521.
- Dirac, P. A. M. (1930). *The principles of quantum mechanics*. Oxford: Clarendon Press.
- Earman, J. (1989). *World enough and spacetime*. Cambridge: MIT Press.
- Earman, J. (2004). Laws, symmetry, and symmetry breaking: Invariance, conservation principles, and objectivity. *Philosophy of Science*, 71(5), 1227–1241.
- Earman, J., & Norton, J. (1987). What price spacetime substantivalism? The hole story. *The British Journal for the Philosophy of Science*, 38(4), 515–525.
- Eddon, M. (2013). Quantitative properties. *Philosophy Compass*, 8(7), 633–645.
- Greaves, H. (2011). In search of (spacetime) structuralism. *Philosophical Perspectives*, 25, 189–204.
- Greaves, H., & Wallace, D. (2014). Empirical consequences of symmetries. *The British Journal for the Philosophy of Science*, 65(1), 59–89.
- Healey, R. (2006). Symmetry and the scope of scientific realism. In W. Demopoulos & I. Pitowsky (Eds.), *Physical theory and its interpretation: Essays in Honor of Jeffrey Bub, the western Ontario series in philosophy of science* (pp. 143–160). Dordrecht: Springer.
- Healey, R. (2009). Perfect symmetries. *British Journal for the Philosophy of Science*, 60(4), 697–720.

- Ismael, J., & Van Fraassen, B. C. (2003). Symmetry as a guide to superfluous theoretical structure. In K. Brading & E. Castellani (Eds.), *Symmetries in physics: philosophical reflections* (pp. 371–392). Cambridge: Cambridge University Press.
- Jacobs, C. (2020). Absolute velocities are unmeasurable: Response to middleton and murgueitio ramirez. *Australasian Journal of Philosophy*.
- Lewis, D. (1970). How to define theoretical terms. *Journal of Philosophy*, 67(13), 427–446.
- Martens, N. C. M. (2019). Machian comparativism about mass. *The British Journal for the Philosophy of Science*.
- Martens, N. C. M., & Read, J. (2020). Sophistry about symmetries? *Synthese*.
- Middleton, B., & Murgueitio Ramírez, S. M. (2020). Measuring absolute velocity. *Australasian Journal of Philosophy*, 1–11.
- Møller-Nielsen, T. (2017). Invariance, interpretation, and motivation. *Philosophy of Science*, 84(5), 1253–1264.
- Mundy, B. (1986). On the general theory of meaningful representation. *Synthese*, 67(3), 391–437.
- Nozick, R. (2001). *Invariances: The structure of the objective world*. Cambridge, MA: Harvard University Press.
- Pooley, O. (2013). Substantivalist and relationalist approaches to spacetime. In R. Batterman (Ed.), *The oxford handbook and of philosophy of physics*. Oxford: Oxford University Press.
- Read, J., & Møller-Nielsen, T. (2020). Redundant epistemic symmetries. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*.
- Rindler, W. (1977). *Essential relativity: Special, general, and cosmological*. Berlin: Springer.
- Roberts, B. W. (2020). Regarding ‘Leibniz equivalence’. *Foundations of Physics*, 50(4), 250–269.
- Roberts, J. T. (2008). A puzzle about laws, symmetries and measurability. *The British Journal for the Philosophy of Science*, 59(2), 143–168.
- Russell, J. S. (2018). Quality and quantifiers. *Australasian Journal of Philosophy*, 96(3), 562–577.
- Rynasiewicz, R. (1992). Rings, holes and substantivalism: On the program of Leibniz algebras. *Philosophy of Science*, 59(4), 572–589.
- Rynasiewicz, R. (1994). The lessons of the hole argument. *British Journal for the Philosophy of Science*, 45(2), 407–436.
- Saunders, S. (2003a). Indiscernibles, general covariance, and other symmetries: The case for non-reductive relationalism. In *Revisiting the foundations of relativistic physics*, ed. A. Ashtekar (pp. 151–173).
- Saunders, S. (2003b). Physics and Leibniz’s principles. In K. Brading & E. Castellani (Eds.), *Symmetries in physics: philosophical reflections* (pp. 289–307). Cambridge: Cambridge University Press.
- Saunders, S. (2007). Mirroring as an a priori symmetry. *Philosophy of Science*, 74(4), 452–480.
- Stalnaker, R. (1979). Anti-essentialism. *Midwest Studies in Philosophy*, 4(1), 343–355.
- Van Fraassen, B. C. (1989). *Laws and Symmetry*. Oxford: Oxford University Press.
- Wallace, D. (2002). Time-dependent symmetries: The link between gauge symmetries and indeterminism. In K. Brading & E. Castellani (Eds.), *Symmetries in physics: Philosophical reflections* (pp. 163–173). Cambridge: Cambridge University Press.
- Wallace, D. (2019). Observability, redundancy and modality for dynamical symmetry transformations. <http://philsci-archive.pitt.edu/16622/>
- Weatherall, J. O. (2018). Regarding the ‘hole argument’. *The British Journal for the Philosophy of Science*, 69(2), 329–350.
- Wolff, J. E. (2020). *The metaphysics of quantities*. Oxford: Oxford University Press.