

Same-diff?

Conceptual similarities between gauge transformations and diffeomorphisms

Part II: Challenges to sophistication

Henrique de A. Gomes*

University of Oxford
Oxford, OX2 6HT, UK

October 23, 2022

Abstract

The following questions are germane to our understanding of gauge-(in)variant quantities and physical possibility: how are gauge transformations and spacetime diffeomorphisms understood as symmetries, in which ways are they similar, and in which are they different? To what extent are we justified in endorsing different attitudes—nowadays called sophistication, haecceitism, and eliminativism—towards each? This is the second of four papers taking up this question, and it is the one that most engages with the metaphysical debates surrounding our understanding of symmetry and equivalence.

In this paper, I will provide two desiderata for the application of a treatment of symmetries known as ‘sophistication’ and show that both general relativity and Yang-Mills theory satisfy these desiderata. The first desideratum for symmetries is mathematical, and was shown to hold for general relativity and Yang-Mills in the first paper (Gomes, 2021b): (i) that they correspond to the automorphisms of the structured base sets for the models of the two theories. Here I will extend the desideratum to more general theories. The second desideratum is mathematical-physical: (ii) that the general type of structure to which Desideratum (i) refers is axiomatizable and that this axiomatization can be phrased in terms of basic physical predicates of the theory. In the third paper in the series, Gomes (2022),

*gomes.ha@gmail.com

I will provide yet a third desideratum, to deal with an issue that goes beyond the standard debate about sophistication.

Same-diff [noun]: *an oxymoron, used to describe something as being the same as something else. Often used as an excuse for being wrong.* (Urban dictionary).

Contents

1	Introduction	3
1.1	Motivation	3
1.2	Morals of this paper	4
2	Formal desiderata for sophistication: two challenges	6
2.1	Anti-haecceitism about spacetime in more detail	7
2.2	Challenges to sophistication	9
2.2.a	Extending Sophistication and invariant structure	10
2.2.b	Sophistication on the cheap?	12
3	Sophistication defended	14
3.1	Desideratum (i): symmetries induced by automorphisms of a base set	17
3.2	Desideratum (ii): axiomatization of the invariant structure	20
3.2.a	The case of spacetime metrics.	20
3.2.b	The case of connections on a principal bundle.	22
4	Summing up	25

1 Introduction

This is the second of three papers analysing the similarities and distinctions between the gauge symmetries of Yang-Mills theory and the spacetime diffeomorphisms of general relativity. The first three analyse more formal aspects while the fourth will analyse more detailed aspects of this comparison. The previous paper, (Gomes, 2021b), has given general definitions of dynamical symmetries and applied them to the specific cases of Yang-Mills and general relativity. In this paper I will try to provide more general criteria for the interpretation of symmetries, with a more metaphysical focus and taking these two theories as templates.

1.1 Motivation

Gauge theories lie at the heart of modern physics: in particular, they constitute the standard model of particle physics. Philosophers of physics generally accept as the leading idea of a gauge theory—or as the main connotation of the phrase ‘gauge theory’—that it involves a formalism that uses more variables than there are physical degrees of freedom in the system described; and thereby more variables than one strictly speaking needs to use. Hence the common soubriquets: ‘descriptive redundancy’, ‘surplus structure’, and more controversially, ‘descriptive fluff’ (e.g. Earman (2002, 2004)).

Although the main idea and connotation of descriptive redundancy has been endorsed by countless presentations in the physics literature, some celebrated philosophers, such as Healey (2007) and Earman (2002) among others, have gone beyond this connotation, and defended a stronger, *eliminativist view*. The view is that gauge symmetry must be ‘eliminated’ before determining which models of a theory represent distinct physical possibilities, on pain of radical indeterminism. For them, the connotation of ‘fluff’ is that it can have no purpose.

But radical indeterminism also threatens theories such as general relativity, embodying diffeomorphism symmetry. There the threat—which underlies the infamous ‘hole argument’ (Earman & Norton, 1987)—finds its most popular resolution in a treatment of, or rather, an attitude towards, symmetry-related models, called *sophisticated substantivalism*.

Sophisticated substantivalism is *not* eliminativist; but it is *anti-haecceitist*. This jargon can be quickly summarized: haecceitism is the doctrine that objects have an intrinsic identity (or ‘thisness’: *haecceitas*); haecceitistic possibilities involve individuals being “swapped” or “exchanged” without any qualitative difference. Similarly, *quidditistic* possibilities involve properties being “swapped” or “exchanged” without any qualitative difference. Anti-haecceitists about spacetime points thus deny that there are possible worlds that instantiate the same distribution of qualitative properties and relations over spacetime points, yet differ only over which spacetime points play which qualitative roles. Similarly, the anti-quidditist will insist that there are no two possible worlds that instantiate the same nomological structure, and yet differ only over which properties play which nomological roles.¹

¹Caspar Jacobs points out to me that there may here be some tension between the standard use of qualitative and quidditism. for a couple of general philosophy topics this term. The properties that are swapped under

In sum, according to the doctrine of sophisticated substantivalism, possible worlds can only be distinguished by qualitative properties and spacetime points have no metaphysically robust identity across possibilities—points can *only* be singled out or individuated through their complex web of properties and relations, as encoded in fields.²

A similar resolution is available for gauge symmetry, in the form of ‘anti-quidditism’; but it is much less popular in that context.³ In the case of gauge symmetry, attempting to *eliminate* the symmetry-related models is considered a more viable alternative. But is this alternative really more justified in the case of gauge symmetry? If so, why?

1.2 Morals of this paper

Given a theory with local gauge symmetry, the eliminativist seeks a second theory, empirically equivalent to the first but with possibly different ontological commitments, that has no local gauge symmetry. But not all is lost if eliminativism fails: the threat of a pernicious type of indeterminism can still be countered by *sophistication*. This is the position that, generalizing sophisticated substantivalism, accepts that there are isomorphic models within a theory’s formalism but simply denies that these models represent distinct physical possibilities. And, in both the older and the more recent philosophical literature about isomorphisms in general relativity, sophistication is, in effect if not in name, reported to be the position of the majority of theoretical physicists.⁴

But there are (at least) three worries about sophistication that so far have not been adequately addressed in the literature.

The first two indict sophistication as being only a half-way house, as too abstract to shed light on the underlying ontology of the theory. Within this theme, the first main worry—see e.g. Dewar (2017); Jacobs (2021a); Møller-Nielsen (2017); Read & Møller-Nielsen (2020)—is about

quidditism will, elsewhere, be classified as qualitative (e.g. mass values). Jacobs (2022) proposes, as a solution, that anti-haecceitism, understood as the claim that there are no distinct yet qualitatively identical possibilities, suffices. I agree, but for the sake of the dialectic between gauge and gravity, I will maintain the terminology of ‘quiddistic’, for those possibilities that differ only over which properties play which nomological roles.

²It is not a given that points *can be* singled out (equivalently: uniquely specified or individuated) by such a web of properties and relations: if this web is not sufficiently complex, the specification will fail. This type of obstruction to individuation due to homogeneity is well-known, and related to Black (1952)’s criticism of Leibniz’s Principle of the Identity of Indiscernibles. See (Pooley, 2022, Sec. 3.3) for a thorough exposition, and Wüthrich (2009) for a related debate in the context of general relativity (we will come back to this topic in Gomes (2022)).

³But recently the position has garnered support, starting with Dewar (2017) and followed by Jacobs (2021a,c); Martens & Read (2020).

⁴As reported by e.g. (Wald, 1984, p. 438), (O’Neill, 1983, p. 5), (Hawking & Ellis, 1975, p. 68). The literature about the treatment of isomorphic models and indeterminism in general relativity is vast, and, since the 1980s, tied to the ‘hole argument’ (see Brighouse (1994); Butterfield (1989); Earman & Norton (1987); Hofer (1996) for discussions). More recently, the topic has been revived by Weatherall (2018) and several responses to it: Curiel (2018); Fletcher (2020); Gomes & Butterfield (2022); Pooley & Read (2022); B. Roberts (2020); see B. W. Roberts & Weatherall (2020) for a collection of responses.

whether sophistication can be obtained all too easily, by just stipulating that symmetry-related models of any theory represent the same physical possibility, even without a clear understanding of the symmetry-invariant ontology. This worry is that this stipulation still leaves the inference from mathematical isomorphism to physical (or metaphysical) equivalence opaque or even unjustified. For instance, in the familiar case of vacuum general relativity, it questions whether we are right in assigning physical, chronogeometric significance to diffeomorphism-invariant quantities.

I will call this worry, which finds echoes in the literature on the metaphysics of spacetime (Dasgupta, 2011; Teitel, 2019), Worry (1), or ‘the obscurity of structure-type’. It besets the philosopher more than the physicist, and it will be the focus of this paper.⁵

But this first worry does not question whether or not we have perspicuous representations of each such equivalence class. For, even if we agree on the given notion of isomorphism, and agree on what kind of things are invariant under this notion, we may still not be able to characterize physical possibilities using just those things. For instance, again in the familiar case of vacuum general relativity, even if we agree to assign physical, chronogeometric significance to diffeomorphism-invariant quantities, we still may not know how to specify a particular physical possibility using just diffeomorphism-invariant predicates. In other words, a resolution of the first worry does not provide a set of quantities whose values label the isomorphism classes. The worry that, indeed, we cannot do so is the second worry: which besets the physicist more than the philosopher.

The second worry—about the obscurity of structure-tokens, or Worry (2)—is closer to the physicist’s heart, because, as described in Belot & Earman (1999, 2001), when theoretical physicists—especially those working in quantum gravity—get down to brass tacks, they need to ‘get inside’ each world; they seek more perspicuous, piece-wise characterizations of the symmetry-invariant ontology, even while endorsing sophistication at a more general level. This worry, and its resolution, will be the focus of the third paper, Gomes (2022).

The third worry should be concerning to both the philosopher and the physicist. As Belot (2018) forcefully argues, in certain sectors of general relativity and Yang-Mills theory, and contrary to sophistication, certain isomorphisms *are* taken to relate different physical possibilities. We will also sketch a reply in Gomes (2022) (see also Gomes (2021a)). A more thorough reply will be deferred to forthcoming work, since it requires a detailed analysis of subsystems in general relativity.

My hope is that, jointly, the first three of the four papers comparing gauge symmetries and diffeomorphisms will further illuminate the doctrine of sophistication for both diffeomorphisms and the gauge symmetries of Yang-Mills theory. This will be attained by providing further desiderata for sophistication in general (namely, in the list below), beyond those that have been clearly articulated in the literature—that are mostly included in Desideratum (i) in the list below.

⁵This worry is the same as Martens & Read (2020); Møller-Nielsen (2017) objection to interpretationalism and thus favoring ‘motivationalism’, as we will see in Section 2.2.b.

The desiderata can be divided into the purely mathematical (i), and the physico-mathematical (ii and iii).

- (i) That the symmetries be induced by the automorphisms of some ‘natural’ geometric structure;
- (ii) That this geometric structure is axiomatizable, and that the axiomatization can be phrased in terms of basic physical predicates assumed for theory.
- (iii) That the automorphisms of the structure correspond to changes between different choices of physical coordinate systems, or physical reference frames.

Desideratum (i) ensures the invariant structure of the models admits a natural ‘qualitative’ interpretation. Desideratum (ii) ensures the invariant structure is ‘metaphysically perspicuous’: a property that has been left rather vague within the literature (cf. Section 2.2.b).

Both general relativity and Yang-Mills theory (and also Newton-Cartan theory, which will not be treated here) satisfy all three desiderata.⁶

This second paper will discuss (i) and (ii), but not (iii), which is aimed at the second worry (about structure-tokens) and which is a formalization of Einstein’s original understanding of coordinate systems (cf. Norton (1989)). I will leave a discussion of (iii) for the third paper, Gomes (2022).

The first three papers in this series should be interesting even for those who are only concerned with the interpretation of isomorphisms in general relativity and ‘the hole argument’, and have little interest in gauge theory; it is only in Gomes (2021c) that I will focus on more detailed contrasts between the symmetries of the two theories.

Here is a brief outline about how I plan to proceed. In Section 2 we challenge the sophisticationist doctrine, with several requests for clarification. In Section 3, we respond to these requests by providing sharpened desiderata for sophistication and showing that both Yang-Mills theory and general relativity meet them. In Section 4 I conclude.

2 Formal desiderata for sophistication: two challenges

In this paper I investigate the relation between symmetry-related models and physical possibilities. The general attitude towards this relation that I will defend, introduced briefly in Section 1, is what I, and the literature, call ‘sophistication’. My main aim in this paper is to find good desiderata for sophistication, in general, and not just for the familiar cases of general

⁶And indeed, in these cases, satisfaction of (iii) follows from the satisfaction of (i). The underlying reason is the proof, contained in (Gomes, 2021b, Sec. 5), of a 1-1 correspondence between dynamical symmetries (or a certain notion thereof: see Definition 2 in (Gomes, 2021b, Sec. 2)) and merely notational changes, such as changes of bases of a vector space or changes of local coordinate systems. But I am open to the possibility that the method of proof is limited, and there may be theories for which this correspondence does not hold. Thus I distinguish the two desiderata.

relativity and Yang-Mills (as well as other theories, e.g. Newton-Cartan theory). But if we want to apply the desiderata in general, we must first verify that they apply to these particular cases that have been well-studied in the literature.

Having described the symmetry structure of both Yang-Mills theory and general relativity in (Gomes, 2021b, Secs. 3, 4), and verified that they meet the formal definitions of symmetry outlined in (Gomes, 2021b, Sec.2), and having there also developed a plausible structural interpretation of both theories, I will now develop a more general formalization of sophistication, and check whether it is also clarifying in the familiar cases.

First, in Section 2.1, expanding the initial remarks of Section 1, I describe the doctrine of sophistication and anti-haecceitism in more detail, and pose some initial concerns.

In order to better grasp the validity of sophistication in general, I will expand on the first worry about extending sophistication from general relativity to other theories: the worry about structure-type (Worry (1)). This worry is composed of two parts, which will be dealt with in this paper. Broadly:

Worry (1-i): is the isomorphism-invariant structure of the theory sufficiently clear?

Worry (1-ii): what warrant does the isomorphism-invariant part of the theory have to be the only referring part of the theory?

In the extant literature, the first worry is related to the debate about ‘internal and external sophistication’; and the second worry is often read as a request for a ‘metaphysically perspicuous’ interpretation of the symmetry-related models. I will expound these worries in Section 2.2.

Then, in Section 3 I will turn to answering the skeptic, by propounding my two Desiderata (i) and (ii), that answer (1-i) and (1-ii). Though I could identify reflections of Desideratum (i) in several places in the literature on sophistication, I believe Desideratum (ii) has not yet made an appearance, in any guise.

2.1 Anti-haecceitism about spacetime in more detail

How should we, within a given dynamical theory, interpret models that are related by a dynamical symmetry? The question has a respectable ancestry: it descends from the centuries-old, absolute-relational debate, about the nature of space, that began with the dispute between Newton and Leibniz. In the correspondence between Clarke—Newton’s spokesman—and Leibniz, the question of whether there are alternative possible universes that are qualitatively identical but have their material content translated in space was a major point of contention. According to anti-haecceitism, since these two models for the universe are qualitatively identical and, in some sense, isomorphic, they represent the same physical possibility.

Let us remind the reader of our brief description of haecceitism in Section 1.1: it is the doctrine that possibilities can match qualitatively but differ merely over *which individual* has which set of qualitative properties. Thus anti-haecceitists about spacetime points deny that there are possible worlds that instantiate the same distribution of qualitative properties and relations over spacetime points, yet differ only over which spacetime points play which qual-

itative roles. The notion of qualitative is itself contentious and used rather as a term of art in metaphysics (cf. Adams (1979)). But as a ball-park approximation, we can verbally grasp non-qualitative distinctions by the use of singular propositions: these distinctions typically require proper names, indexicals or other such devices for their expression. We will regiment the idea of ‘qualitative’ property for certain physical theories in Definition 1, in Section 3 (see also footnote 15).⁷

For now, to gather some intuition, let us briefly look at the corresponding debate about Newtonian absolute space. There, anti-haecceitism would rule as identical two universes whose material content was uniformly translated, for the theory has no preferred origin of Newtonian space. That is, space in this theory is modeled by the Euclidean affine space \mathbb{E}^3 , which contains translations as some of its isomorphisms. So, if we take one model’s matter distribution and obtain another by translating that distribution five feet to the left, these two distributions will have the same relation to the ambient, translation-invariant, Euclidean geometry. Therefore, any difference between the two models must be expressed using indexicals, or other singular terms, e.g. that fix the reference of what ‘left’ is and where the origin of \mathbb{E}^3 is in each model. But under any such comparison, a qualitative predicate for a particle, such as ‘being at the center of mass’, retains its truth value under a uniform translation of all particle positions.⁸

One consequence of the anti-haecceitist view is that spatial points have no *primitive* identity across models or worlds; such identity would necessarily be non-qualitative. The position labeled *sophisticated substantivalism*, introduced in the context of the hole argument in general relativity (Earman & Norton, 1987), is anti-haecceitist in this way. What the *sophisticated substantivalist* adds to anti-haecceitism is a commitment to a spacetime manifold, as it is usually understood. That is, the view takes the denial of primitive identity for spacetime points to be perfectly compatible with the existence of a spacetime manifold, composed of spacetime points.

In the words of Hofer (1996, p. 20): “[p]rimitive identity is metaphysically otiose, and not a necessary part of the concept of a substance”. In other words, we should not imbue every aspect of our set-theoretic constructions with referential significance. Kaplan (1979, p. 96) (as quoted in (Pooley, 2022, p. 64)) explains why it may be psychologically appealing to adopt a metaphysics in which points have primitive identity irrespective of their properties and relations, i.e. irrespective of their “clothing”:

“[T]he use of models as representatives of possible worlds has become so natural for logicians that they sometimes take seriously what are really only artifacts of the

⁷In this series of papers, I mean ‘properties’ in the more general sense: not solely as unary, or 1-place relations, but general n-place relations. This was traditionally dubbed an ‘attribute’, but I think this term is ever so slightly more abstract than ‘property’, and has fallen out of use; and so I will stick with ‘property’.

⁸There are important differences between positions and velocities: Newtonian space comes equipped with *an absolute origin for velocities*—‘being at rest’ is a qualitative fact. Therefore anti-haecceitism will identify universes related by a static shift—i.e. a time-independent spatial translation, as described above—but not by a kinematic shift, such as a Galilean boost; see Maudlin (1993) and Rynasiewicz (2022) for more about this distinction.

model. In particular, they are led almost unconsciously to adopt a bare particular metaphysics. Why? Because the model so nicely separates the bare particular from its clothing. The elements of the universe of discourse of a model have an existence which is quite independent of whatever properties the model happens to tack onto them.

I find Kaplan’s explanation, and warning, convincing. That is, *pace* the temptations from model theory, there is no clear-cut separation between the points and their clothing. Nonetheless, the sophisticated substantialist would say, the clothed points of spacetime exist!

In sum, sophisticated substantialists characterize different possibilities as differing only qualitatively, and as being about spatiotemporal properties. Thereby, sophisticated substantialists endorse Earman & Norton (1987)’s condition of *Leibniz equivalence*: that isomorphic models represent the same physical possibility. And they moreover argue that this understanding does not require a reductive, or eliminativist, picture: isomorphic models are all equally valid representations of the same spatiotemporal possibility. In that sense, all spacetime distributions of the metric and matter fields that are related by a diffeomorphism are taken as representing the very same state of affairs (see also (Pooley, 2022, Sec. 3.3) and both Pooley (in press); Pooley & Read (2022) for details).

2.2 Challenges to sophistication

As they now stand, the considerations of Section 2.1, describing sophisticated substantialism for spacetime, do not give us a general characterization, for an arbitrary theory, of what a sophisticated interpretation would be. In Section 2.2.a, I will, rather straightforwardly, extend sophistication to other theories. I will there follow Dewar (2017): a paper that kicked off a considerable recent literature.

But, even in the familiar case of diffeomorphism symmetry of general relativity, the considerations of Section 2.1 should not be assumed to be unproblematic. Indeed, many metaphysicians still find the anti-haecceitist interpretation of diffeomorphism symmetry opaque. And if we are to satisfactorily answer to these metaphysicians in the more general case, our answer should first be satisfactory for the more familiar case of diffeomorphisms. Thus, in order to set out the obstacles towards sophistication in the more general case, we must first revisit the criticisms of sophisticated substantialism.

The main idea of these criticisms is that stripping spacetime points of primitive identity—condemning them to exist solely as clothed entities, dressed by the network of properties and relations in which they take part—is suspicious, or at least not capable of eliminating an ubiquitous and pernicious form of indeterminism (see Dasgupta (2011) and Teitel (2019) for recent examples). Indeed, the label ‘sophisticated’ was introduced with a deliberately negative connotation in Belot & Earman (1999) (see also Belot & Earman (2001)), signalling this tension.⁹

⁹The label has since been accepted by many philosophers, who believe the tension is illusory. Dewar (2017)

Elaborating this suspicion, Dasgupta (2011, p. 147) issues two distinct challenges to the sophisticated substantialist: “(1) to clearly articulate what the underlying qualitative facts are like, and (2) show that they are sufficient to explain (in the metaphysical sense) individualistic facts about the manifold.’

Without further context, these short quotes require some unpacking, and a good dose of interpretation. But in brief: I see (1) and (2) as truly distinct; I take (2) to be encompassed by what I labeled the ‘worry about structure-tokens’, whose treatment I leave to the third paper, (Gomes, 2022).¹⁰ And I take (1) to be, at bottom, about the conceptual and physical obscurity of the type of structure that remains invariant under the isomorphisms. Thus (1) is encompassed by what I labeled ‘the worry about structure-type’, that I then fine-grained into Worries (1-i) and (1-ii) in the introduction to Section 2.

And, in the same gist, there are two recent debates within philosophy of physics, also corresponding to (1-i) and (1-ii). They are, respectively: (1-i): about whether sophistication is, in general, suspicious, because it can be obtained all too easily, and thus (as formulated in the introduction to Section 2), it may leave the core invariant structure shared by the symmetry-related models as only implicit, or even obscure. And (1-ii): about whether the invariant structure—to which sophistication is metaphysically committed—is sufficiently ‘metaphysically perspicuous’. Thus these suspicions, which I will describe more fully in Section 2.2.b, precisify the metaphysician’s Challenge (1).

2.2.a Extending Sophistication and invariant structure

To clarify the extension of sophistication to a wider class of theories beyond general relativity, Dewar (2017, p. 502) contrasts it with *reduction*, i.e. what I called ‘eliminativism’ in Section 1: eliminativism (or reduction) advocates a new theory that excises the structure that could distinguish the isomorphic models, so that symmetry transformations would fix, i.e. act as the identity on, models of the new, reduced theory. On the other hand, models of a sophisticated theory are acted on by, but are to be isomorphic under, symmetry transformations. And that is it: that is the extension.

When we leave the concrete examples of general relativity and Yang-Mills theory, this characterization of sophistication seems rather unsubstantial. Thus the first accusation that ‘sophistication is gotten on the cheap’ is that declaring, by fiat, that symmetry transformations are isomorphisms is too easy; Russell’s ‘theft over honest toil’ comes to mind.

extended the adjective “sophisticated” to other symmetries: where it is understood as allowing commitment to structure without reduction or elimination of the isomorphic copies.

¹⁰To see that (1) and (2) are distinct, suppose we interpret (1) such that a description of what the “underlying qualitative facts are like” is just the listing of the necessary and sufficient conditions for assessing qualitative identity of possibilities. Then, as I will argue in this Section, we can rest assured that sophistication, augmented by Desideratum (i) and (ii), answers this first challenge. But this is at best only a criterion of (meta)physical identity of entire worlds; it may very well be insufficient to ‘get inside’ each world and characterize the individual facts qualitatively; thus (2) remains a separate challenge, and it is related to my worry about structure-tokens. I will return to (2) more fully in the third paper in this series, Gomes (2022).

For, according to Dewar, the requirement that symmetry-related models be isomorphic may not be stringent at all: whenever a theory has symmetries, we could extend sophistication, even to cases where those symmetries do not correspond to any of the better known examples of mathematical structure, by altering the semantics of that theory, such that “we obtain [the new semantics] by taking the old objects, and declaring, by fiat, that the symmetry transformations are now going to “count” as isomorphisms”. Thus Dewar apparently proposes that given a space of models and a symmetry transformation acting on this space, one could just announce an invariant structure as defined implicitly by ‘whatever is left invariant’ by the action of the symmetries. In such cases, knowing the isomorphisms sheds no light on what structure they are leaving invariant.

Dewar calls this extension the ‘external method’ (Dewar, 2017, p. 502). Here is Dewar (2017, p. 502):

Is there a way to precisify what is meant? Here is one way to do so. Rather than trying to define the objects of the new semantics ‘internally’, as mathematical structures of such-and-such a kind (paradigmatically, as sets equipped with certain relations or operations), we instead define them ‘externally’: as mathematical structures of a given kind, but with certain operations stipulated to be homomorphisms (even if they’re not ‘really’ homomorphisms of the given kind). For example, one way to define vector spaces is to define them as sets equipped with operations of addition and scalar multiplication, obeying appropriate axioms. This is the internal method. The alternative is to define them as spaces of the form \mathbb{R}^k , with the further feature that linear transformations are declared to be homomorphisms—and in particular, that invertible linear transformations are isomorphisms. This is the external method.

To further explicate the external method, and contrast it with the internal, it is useful to follow Klein (1893)’s distinction, taken as a starting point for the Erlangen programme. That distinction is reflected in (Dewar, 2017, p. 502)’s jargon of ‘external vs internal’, but it is made more precise in Jacobs (2021c, Ch. 4.1), as:

The symmetry-first approach (external): finding a structure implicitly as ‘what stays constant across the symmetry-related models’;

The structure-first approach (internal): finding the symmetry-related models as those that possess the same structure.

Thus far in our treatment of symmetries (see Gomes (2021b)), we have been given some space of models of a theory and a notion of dynamical symmetry: so to use Jacobs’ jargon, we were adopting the symmetry-first approach—so that, at least at first sight, we “could only glimpse” the structure thereby defined as whatever is invariant under the symmetries. In the structure-first (or internal) approach, the aim is to give a characterisation of a structure in terms of relations defined over its domain, such that this structure is invariant under the dynamical symmetries of the theory. As Jacobs puts it:

Structure-first sophistication consists of: the stipulation of a set of relations over the theory’s sub-domains, such that the bijections which induce dynamical symmetries of the theory’s models leave these relations invariant. (Jacobs, 2021c, p. 62)

2.2.b Sophistication on the cheap?

We are here in the vicinity of two philosophical debates about symmetry, that, to my view, add further detail to the metaphysician’s challenge (1), broadly about ‘the obscurity of structure-type’, that we glimpsed in the case of general relativity in the introduction to Section 2.2.

The first debate was sketched in Section 2.2.a, and is about whether sophistication can be attained too easily, i.e. whether it thus needs further explication. For sophistication by brute force, i.e. Dewar’s ‘external’ sophistication (and also Jacob’s symmetry-first approach), seems liable to leave, even at a strictly mathematical level, invariant structure as opaque, or not sufficiently understood. In the introduction to Section 2, I related this debate to Worry (1-i).

And this lack of clarity about the common mathematical structure motivates the debate about whether symmetry-related models of a given theory should invariably be regarded as being physically equivalent, i.e. as making the same claims about all physical facts about the world, even in the absence of a clear, or ‘metaphysically perspicuous’, understanding of the common ontology of the models. In the introduction to Section 2, I related this debate to Worry (1-ii).

Møller-Nielsen (2017) introduced a helpful labelling for the two sides of this second debate on physical significance. He labels the more permissive answer—symmetry-related models are invariably physically equivalent—as the *interpretational approach* to symmetries; and he contrasts it with the *motivational approach*, which requires a characterisation of the common ontology of symmetry-related models before acknowledging physical equivalence. This is the answer he endorses.

As to the first debate, Section 2.2.a reported on an initial attempt to repel the accusation of cheapness, by distinguishing internal from external sophistication. But as we will now see, that attempt is not satisfactory.

All hands agree (e.g. Jacobs (2021c); Martens & Read (2020); Møller-Nielsen (2017)) that an implicit definition, as in Dewar’s external method, would not sufficiently explicate the common ontology of the symmetry-related models. According to Møller-Nielsen (2017, p. 1264): “[I]t is simply opaque what, according to the external approach, the world is really like.” Thus Martens & Read (2020, p. 9) endorse Dewar (2017, p. 502)’s ‘internal’ sophistication: equivalence can only be justified in cases where the symmetries coincide with isomorphisms *of some familiar structure*.

But the challenge here is to characterize just what ‘familiar structures’ are. For the role of familiarity is to stop the externally defined symmetries from being reconstrued as internal by just declaring the invariant structure to be whatever is preserved by the symmetries. This may seem tautologous, or at least a mere reconstrual, but it nonetheless threatens to collapse the distinction between internal and external. Thus far, being internal rather than external seems

to rely on whether the structure is ‘familiar’ or not: a vague and conservative criterion that does not illuminate the question.

It is helpful to discuss these issues in the supposedly clear case of general relativity. Even there, we have not yet precisely identified *the type of structure* that is invariant. The idea is that the models of the theory are structured sets, $\mathcal{D} = (D, R)$, consisting of a base (unstructured) set D and relations among the elements of this set, which we can here give just one label, R , which are supposed to represent the invariant structure. In general relativity, a model “consists of a base manifold M over which we have defined some (geometrical) structure in the form of the tensor fields” (Jacobs, 2021c, p. 60). But what exactly are the relations that stay invariant under the symmetries, i.e. which stay invariant under the (pullback of) active diffeomorphisms? Tensor fields certainly do not remain invariant. A statement such as “the metric tensor is g_{ab} ” is not symmetry-invariant; it can be true in one model and false in a symmetry-related one. The distances between the points of M seen as an unstructured base set are also not invariant, since the distance between x and y according to g_{ab} is not the same distance as that according to the dragged-along metric f^*g_{ab} . The same reasoning would of course apply to angles, Riemann curvature scalars, etc.¹¹

The conclusion is that even if we try to apply the general, straightforward definition of internal sophistication—as a property of structured sets—to the reasonably well-understood case of general relativity, the structure that remains invariant under isomorphisms is characterized as the abstract set of quantities that are invariant under pull-back: just as the symmetry-first, or external, approach—not the structure-first, or internal, approach—would suggest.¹²

In Section 3 I will overcome these shortcomings about the usual definition of internal sophistication, and thus resolve Worry (1-i). My Desiderata (i) and (ii) are the core of my response, since they give further criteria for the invariant structure, that go beyond its mere familiarity.

Moving on to the second debate, corresponding to Worry (1-ii): even supposing we have succeeded in defining some structure internally in a non-tautologous manner, there is still the matter of connecting the mathematical and the physical. Given the vast and flexible resources of mathematics, a strictly formal internal definition will hardly quell the skeptic’s suspicion that the said structure fails to support the metaphysical interpretation we would like to hang on it. As argued by Martens & Read (2020): if there is no ‘metaphysically perspicuous’ package accompanying the formulation of the symmetry-related models, sophistication is still metaphysically obscure.

¹¹And in particular, I don’t find the discussions in the standard textbooks of general relativity illuminating on this issue. Here is Wald (1984, p. 438): “If a theory describes Nature in terms of a spacetime manifold M and tensor fields, T , then if $f : M \rightarrow N$ is a diffeomorphism, the solutions (M, T) and (N, f^*T) have physically identical properties. Any physically meaningful statement about (M, T) will hold with equal validity for (N, f^*T) . On the other hand, if (M, T) and (N, T') are not related by a diffeomorphism, *and if the tensor fields T represent measurable quantities*, then (N, T') will be physically distinguishable from (M, T) .” (my emphasis).

¹²Besides the equivalence class of metrics, the smooth structure of the manifold stays invariant under arbitrary diffeomorphisms. That is because it is defined by a maximal atlas, which is already an abstract equivalence class, for all of the diffeomorphism-related charts (see (Gomes, 2021b, Sec. 5)).

Thus, with this second debate in mind, following a nomenclature suggested in (Belot, 2003, p. 220), Møller-Nielsen (2017) (see also Martens & Read (2020)) insists we still need to find a metaphysically perspicuous interpretation of the invariant structure, even in the internal case.¹³ Thus, (Martens & Read, 2020, p. 9) advise, even for internal sophistication, that more needs to be said so that we can “regard those models—interpreted as being isomorphic—as in fact representing the same physical states of affairs.” And thus there is a contradiction with (Jacobs, 2021c, p. 62), who writes:

If we agree that an interpretation of a theory provides a metaphysically perspicuous picture if it tells us plainly and clearly which entities the theory posits (ontology) and what the fundamental relations between these entities are (ideology), then structure-first sophistication is perspicuous in this sense.

I agree with Martens & Read (2020); Møller-Nielsen (2017): even assuming we have obtained a non-tautologous internal characterization of structure, we won’t necessarily have succeeded in providing a metaphysically perspicuous picture, because even the internal picture may not yet state “plainly and clearly which entities the theory posits (ontology)”.

But at this point, these requests for ‘metaphysical perspicuity’ remain vague: so far, ‘perspicuity’ seems to be perhaps best understood as ‘intelligibility’ of the mathematical structure. As with the label of ‘internal’, the label of ‘metaphysically perspicuous’ does not yet satisfy any rigorous requirements.

In Section 3 I will also respond to this second debate/Worry (1-ii). My novel proposal is to characterize perspicuity as part of my Desideratum (ii), by explicitly anchoring the mathematical structure of the theory on axioms bearing primitive physical significance.

After this review of the debates/worries about the ‘obscurity of structure-type’, i.e. (1-i) and (1-ii), that precisify Dasgupta’s Challenge (1) to sophistication, I now turn to my response.

3 Sophistication defended

In this Section, I will provide some answers to the debates reviewed in the preceding Section and to the metaphysician who is skeptical about sophistication. They are, now in a bit more detail than provided in their formulation in the introduction to Section 2:

(1-i) Whether the invariant structure is sufficiently (internally) understood, and whether it refers to some notion of ‘qualitative’ property;

(1-ii) whether the invariant structure is the only part of the theory that refers, i.e. whether it is ‘metaphysically perspicuous’.

And the metaphysician’s concerns are just a coarse-graining of these worries. To recap: he may not be satisfied, for instance, that the common ontology of two isomorphic models of general relativity is just ‘fields on the manifold, where the latter is interpreted anti-haecceitistically’.

¹³(Belot, 2003, p. 220) talks about finding a ‘perspicuous formulation’ of the symmetry-related models, Møller-Nielsen (2017) also talks about perspicuous characterization.

He is concerned that a characterization of structure as ‘qualitative’ comes all too cheaply, as a bare intuition that cannot be (meta)physically cashed out.

Before I propose my two desiderata for sophistication (now in more detail), we need some definitions (for more information about the nomenclature here, see (Gomes, 2021b, Sec. 2)). First, let Φ be the space of models of the theory whose models $\varphi \in \Phi$ are maps from a tensor bundle of valence (r, s) over a structured, finite-dimensional base set N , to a value space F , i.e.: $\varphi : T^{r,s}N \rightarrow F$. Now we can define three different notions of symmetry acting on Φ .

Let \mathcal{G} be the group of isomorphisms of Φ , so that Φ is a groupoid, i.e. a category in which every arrow is an isomorphism, with the objects of the category being the models, i.e. the elements of Φ (cf. (Gomes, 2021b, Sec. 2.3) and footnote 15 therein for more on this):

$$\begin{aligned} \mathcal{G} \times \Phi &\rightarrow \Phi \\ (g, \varphi) &\mapsto \varphi^g. \end{aligned} \tag{3.1}$$

Let $\text{Aut}(N)$ be the group of automorphisms of N , that preserve its structure (e.g. smooth, fibered, etc). Mathematically, any given notion of automorphisms of the base set can be pulled-back to define a corresponding transformation on the (values of the) fields, which can be an isomorphism in the relevant category that has the theory’s models as objects.¹⁴ Therefore $\text{Aut}(N)$ acts on $T^{r,s}N$ and thus acts through pull-back on Φ , as (cf. (Gomes, 2021b, Sec. 3))

$$\text{Aut}(N)^* : \Phi \rightarrow \Phi \tag{3.2}$$

Now, the automorphisms of N change only ‘which object plays which role’, and so we can regiment the idea of qualitative with the following definition:

Definition 1 (Qualitative) *Given, Φ and N as defined above, those properties that remain invariant under the induced action of the automorphisms of the (structured) base set, $\text{Aut}(N)^*$ will be called qualitative properties of Φ .*¹⁵

And finally, let T be a dynamical theory for $\varphi \in \Phi$, with $\overline{\mathcal{G}}$ as its group of dynamical symmetries, i.e.:

$$\begin{aligned} \overline{\mathcal{G}} \times \Phi &\rightarrow \Phi \\ (\overline{g}, \varphi) &\mapsto \varphi^{\overline{g}}. \end{aligned} \tag{3.3}$$

¹⁴As described in (Gomes, 2021b, Sec. 4): the Ehresmann connection ω is as described above, a map from the tangent bundle of the principal bundle to a Lie algebra, i.e. $\omega : TP \rightarrow \mathfrak{g}$, and the metric is a map from the symmetric tensor product of the tangent bundle to the reals $g_{ab} : TM \otimes_S TM \rightarrow \mathbb{R}$. What matters for this discussion is that an action of the automorphisms of P or M can be ‘pulled-back’ to an action on the values of these fields.

¹⁵Note that this definition is not quite the same as the ‘verbal characterization’, given in the second paragraph of Section 2.1, that referred to proper names, indexicals, etc. If we take coordinate systems to provide names for the objects of our manifolds (which is not the usual understanding of ‘names’ in this context) then Definition 1 would be analogous to the verbal characterization *provided* we impose (standard) mathematical restrictions on our ‘naming’ the objects of a manifold *and* we have an active-passive correspondence for the automorphisms of N . This was shown to hold for smooth manifolds and fibered manifolds in (Gomes, 2021b, Sec. 5).

These symmetries are here restricted to preserve the value of a salient set of quantities, and to preserve some background structure of Φ (for example, the Hamiltonian and the symplectic structure, respectively; see (Gomes, 2021b, Sec. 2.1) and Definitions 1 and 2 therein for more on this).

Now note that according to Definition 1, qualitative properties are independent of the dynamics of the theory T . Our first Desideratum will make sure that all three notions of ‘symmetry’, given in equations (3.1), (3.2), and (3.3), match.

The two Desiderata are:

- (i) Suppose theory T is given, jointly with Φ, \mathcal{G} and $\overline{\mathcal{G}}$. The isomorphisms \mathcal{G} are required to be:
 - (i_a) such that isomorphisms and symmetries coincide, $\overline{\mathcal{G}} = \mathcal{G}$ (i.e. the action of $g \in \mathcal{G}$ and of $\bar{g} \in \overline{\mathcal{G}}$ on each element of Φ correspond one-to-one) and, moreover,
 - (i_b) induced by the automorphisms of some base space:

For all $\varphi \in \Phi$, and for each $g \in \mathcal{G}$, there exists a unique $t_g \in \text{Aut}(N)$ such that

$$\varphi^g = t_g^* \varphi ; \tag{3.4}$$

- (ii) Φ and \mathcal{G} are not initially given, but jointly arise from: (ii_a) an axiomatization; that (ii_b) uses only terms that have a direct physical interpretation.

In Section 3.1 I will show that my Desideratum (i) gives at least a partial answer to Worry (1-i) by conjoining the arguments of this paper with some of the results of (Gomes, 2021b). Desideratum (i) first of all guarantees a necessary, but not sufficient condition for the interpretation of isomorphic models as physically identical: namely, that their differences are physically unobservable. The argument here, articulated in (Gomes, 2021b, Sec. 2.3) (based on an argument of Wallace (2019)), is that if observations are dynamical processes, then dynamical symmetries leave all observations invariant, and thus are unobservable. By requiring the isomorphisms of the models of the theory to be identical to the dynamical symmetries of the theory, Desideratum (i) guarantees that isomorphisms are physically unobservable.

Second, Desideratum (i) requires the isomorphisms to be induced by the automorphisms of the underlying structured base set of the theory, which I will take to mean that these automorphisms affect only ‘which objects play which role’, in line with Definition 1. Thus Desideratum (i) guarantees that the differences between isomorphic models are *not* qualitative, in a well-defined notion of qualitative.

Note that the actual relations and quantities that remain invariant under these automorphisms are still only implicitly defined, i.e. by the symmetry-first method. Nonetheless, by constraining the notion of isomorphisms, Desideratum (i) protects the definition of internal sophistication from ‘cheapness’, i.e. from being mere (internal) reconstructions of arbitrary (external) symmetries.¹⁶

¹⁶The third Desideratum, to be described in Gomes (2022), requires those automorphisms to have a passive

This remaining concern about a ‘structure-first definition’ will only be resolved in Section 3.2. There I will show that my Desideratum (ii) answers challenge (1-ii), by requiring the invariant structure to have a primitive physical significance, since it is to be defined axiomatically in terms of primitive physical terms. I will verify that this resolution applies to spacetime metrics and Ehresmann connections in principal bundles. The idea then is that other theories will be sophisticated in the appropriate way if they also admit such constructions.

3.1 Desideratum (i): symmetries induced by automorphisms of a base set

My first desideratum is essentially the content of Earman’s two ‘SP principles’ about spacetime symmetries, proposed in (Earman, 1989, p. 45-47) and often taken to be extendible to gauge theory (see e.g. (Belot, 2013, p. 9), (Jacobs, 2021c, Ch. 4, 8), Hetzroni (2021); Jacobs (2021b)). Jointly, the two principles require that the dynamical symmetries should be just those induced by automorphisms of the structured base set (cf. (Gomes, 2021b, Footnote 14) for my definition of automorphism and the previous Section, for my assumptions of a base set for the dynamical variables). In the case of general relativity, this structured base set is a smooth manifold, and, in the case of Yang-Mills theory, it is a fibered smooth manifold called a principal bundle.

To articulate the principles, Earman distinguishes internal and external parameters. External parameters are the independent, or base-set variables, i.e. the points of N in the previous Section, and thus correspond, in our two main study-cases, to spacetime points or to the points of the principal fiber bundle (as organized in orbits); internal parameters are the value spaces F —where the fields of different (field) theories take their values, taking as their argument the independent, external or base-set variables.

In more detail, Earman’s SP-principles are, in a language suitable for this paper:

SP1: Any symmetry of theory T is induced by some automorphism of spacetime $\text{Aut}(M)$, where M is spacetime, i.e.

$$\text{Given } \bar{g} \in \bar{\mathcal{G}}, \quad \varphi^{\bar{g}} = t^* \varphi \quad \text{for some } t \in \text{Aut}(M); \quad (3.5)$$

SP2: Any automorphism of spacetime yields a symmetry of theory T ,

$$\text{Given } t \in \text{Aut}(M), \quad t^* \varphi = \varphi^{\bar{g}} \quad \text{for some } \bar{g} \in \bar{\mathcal{G}}. \quad (3.6)$$

I take this to be the content of the principles (in the original context of spacetime): they ensure a theory’s spacetime has ‘just the right amount’ of structure for symmetry-related possibilities to be identified by an anti-haecceitist.

correspondence as a change of physical coordinates. Then, in tandem with an active-passive correspondence for the symmetries, like the one shown to exist for diffeomorphisms and gauge transformations in (Gomes, 2021b, Sec. 5), we can construe the differences between isomorphic models as merely notational. Jointly, these arguments regiment the idea of qualitative as ‘coordinate-independent’ as opposed to the vaguer, but similar in spirit, ‘without the use of singular terms’.

That is, if *SP1* fails then there is a symmetry of T that does not correspond to an automorphism of spacetime. For example, following the example of Section 2.1, of Newtonian mechanics, we saw that spacetime contains a standard of absolute rest: spacetime is a geometric structure, $\mathbb{E}^3 \oplus \mathbb{R}$ (cf. (Earman, 1989, p. 9-12)), whose automorphisms do not include Galilean boosts. Thus for \bar{g} being a boost, we can find no corresponding $t \in \text{Aut}(\mathbb{E}^3 \oplus \mathbb{R})$. Accordingly, ‘being at rest’ is a qualitative fact in Newtonian mechanics. On the other hand, Newtonian mechanics does not have a standard of absolute *place*, e.g. preferred origin. Thus for \bar{g} being a uniform translation of particle positions, there will always be a $t \in \text{Aut}(\mathbb{E}^3 \oplus \mathbb{R})$.

According to Definition 1, anti-haecceitism—and more specifically, the sophistication I am defending—will identify possibilities related by static shifts, i.e. spatial and temporal translations, but not by kinematic shifts, i.e. boosts; see Maudlin (1993) and Rynasiewicz (2022) for more about this distinction (see also footnote 8 in Section 2.1). That is, boosts are dynamical symmetries of Newton’s laws (which comprise the theory T), which means that the standard of absolute rest is irrelevant to the theory’s dynamics—and *SP1* fails. In short, anti-haecceitism, or the denial of non-qualitative facts, does not account for this dynamical symmetry of T .¹⁷

Accepting Møller-Nielsen (2017)’s Motivationalism, we should feel ‘motivated’ to find a new theory (e.g. Neo-Newtonian, or Newton-Cartan gravity), for which the dynamical symmetry *is* induced by an automorphism of a natural geometric structure.

Conversely, if *SP2* fails, then there is at least one automorphism $t \in \text{Aut}(M)$ of spacetime that is not a dynamical symmetry $\bar{g} \in \bar{\mathcal{G}}$ of T . So for our main case of spacetime and general relativity, and haecceitism vs anti-haecceitism about points: a failure of *SP2* would involve the theory T , i.e. general relativity, assigning a dynamical role to spacetime points and regions. That is, two models, differing only by which region supports which pattern, would not be related by a dynamical symmetry, and thus would be empirically discernible (according to the converse of the empirical unobservability thesis, discussed in (Gomes, 2021b, Sec. 2)). So theory T would provide a dynamical role for singular terms of spacetime, and thereby draw empirical distinctions between regions of spacetime which are qualitatively identical. Thus, in this case, an anti-haecceitist position about spacetime would have to rub out empirical differences.

These ideas and specific principles *SP1* and *SP2* are straightforwardly extendible to gauge theories (see e.g. (Jacobs, 2021c, Ch. 4, 8), Hetzroni (2021); Jacobs (2021b)). Now we can make the argument sketched in the introduction of this Section 3 (penultimate paragraph), more like a deductive argument:

Assumption 0: The physical differences between worlds modeled by φ_1, φ_2 are ‘in principle’ observable.¹⁸

¹⁷Note, however, that if we formulate the dynamics of non-relativistic particles through an action principle, then boosts will not preserve the value of the action: they add boundary contributions. Thus, in Newtonian mechanics, boosts are not dynamical symmetries according to my primary desideratum here.

¹⁸This assumption identifies physical differences, empirical differences, and observable differences. But here I do *not* use ‘empirical’ to denote the traditional positivist and post-positivist ‘meter-readings’ or ‘no-special-training-needed for the judgment’, or ‘the sheer look’—a very common denotation in the literature about the

Assumption I (Unobservability): Two models φ_1, φ_2 are observationally equivalent iff they are related by a dynamical symmetry $\bar{g} \in \bar{\mathcal{G}}$ of T of (3.3), i.e. iff $\varphi_2 = \varphi_1^{\bar{g}}$.¹⁹

Assumption II (Desideratum (i_a)): Dynamical symmetries coincide with the isomorphisms of the space of models of the theory, $\mathcal{G} = \bar{\mathcal{G}}$;

Assumption III (Desideratum (i_b)): the isomorphisms \mathcal{G} of the space of models Φ are induced by the automorphisms of the base set of $\varphi \in \Phi$, i.e. $\text{Aut}(N)$.

Conclusion: for T obeying Desideratum (i), the threat of physical indeterminism is repelled by a commitment to a qualitative ontology.

Proof: the threat of physical indeterminism arises from the existence of dynamical symmetries. The threat is dissolved if we are to physically identify symmetry-related possibilities. This is how far Assumptions 0 and I get us, as argued in the last paragraph of (Gomes, 2021b, Sec. 2.2), on *empirical unobservability* of dynamical symmetries. Assumption II translates dynamical symmetries to a notion of mathematical isomorphism, as described in (Gomes, 2021b, Sec. 2.2). The automorphisms of the base set $\text{Aut}(N)$ are those permutations that preserve some base set structure (smooth, fibered, etc): under these constraints, they change the ‘which point is which’, or ‘which property is which’ (according to Definition 1). And so we take these automorphisms to give only non-qualitative differences; thus by Assumption III, symmetries give only non-qualitative differences, and we obtain the conclusion.

We have already shown that the symmetries of general relativity and Yang-Mills (in the principal bundle or in the Atiyah-Lie bundle formalisms) can be understood as isomorphisms induced by the automorphisms of an underlying geometric structure, in (Gomes, 2021b, Sec. 5). So both theories fulfill Desideratum (i). What is left to say?

Assumptions II and III above take the isomorphisms of the models, the dynamics, and the automorphisms of the base space as given. Thus Desideratum (i) regiments the notion of ‘internal sophistication’ by relating these three notions, thereby severely constraining the structure that is to be left invariant. But the skeptic could still be unsatisfied: even though Desideratum (i) avoids cheap sophistication, it does not provide a truly ‘structure-first’ definition of the invariant structure. Desideratum (ii), to be treated next in Section 3.2, will fill this gap: isomorphisms will emerge from an axiomatic, or rather, synthetic, definition of the invariant structure; that is, the isomorphisms will emerge from giving the structure first.

Before we leave this Section, I want to connect it with both what has come before, in Gomes (2021b), and what is to come, in Gomes (2022). As argued in in (Gomes, 2021b, Sec. 5), fulfillment of Desideratum (i) in the case of Yang-Mills theory and general relativity also allows a deflation of symmetries by giving them a passive gloss (see also footnote 16). That is, once symmetries are construed as isomorphisms that are induced by the automorphisms of a structured base set, and we can build an active-passive correspondence for those automorphisms,

theory-observation distinction of the past fifty years. I use it to denote ‘in-principle-observable’, in a very encompassing sense of ‘in-principle’. See also (Gomes, 2021b, Sec. 2.2).

¹⁹Note that this symmetry need not respect the fibers of the vector bundle: e.g. though the metric is a section of a vector bundle over spacetime, diffeomorphisms do not preserve the fibers.

we can understand symmetry-invariance as mere notational invariance. Thus, for theories that satisfy Desideratum (i), we are motivated to construct a passive-active correspondence, which would indeed reinforce the common belief that the symmetries of those theories correspond to notational redundancy. Such a construction is the basis of our Desideratum (iii), to be expounded in [Gomes \(2022\)](#).

3.2 Desideratum (ii): axiomatization of the invariant structure

In this Section, I will put the final nail on the coffin of Worry (1) (and alongside it, of the metaphysician’s Challenge (1)) of Section 2.2, at least for the well-understood cases.

In Section 2.2 I have already argued that previous descriptions of sophistication may lack that which Desideratum (ii) would guarantee: a metaphysically perspicuous interpretation as an internally sophisticated theory. This Section will answer these two related concerns, viz. about a perspicuous metaphysics and about the isomorphisms being of some invariant structure that is understood first (or internally, but in a more robust sense than that provided by Desideratum (i)).

Here is the basic idea of the resolution. The discussions of the previous Section 3.1 take a notion of isomorphism for granted and give an account of these isomorphisms through the automorphisms of a structured base space, by satisfying Desideratum (i). Thus, I argued, invariant structure could be understood qualitatively. Now we will recover, or build-up, the space of models and their isomorphisms by first accepting certain physically meaningful patterns as the basic constituents of a theory. From a metaphysical standpoint, we would like to know what there is, first, and then deduce from that physical input some representation through models and a notion of model-isomorphism.

The axiomatic approaches to be described below—which are not new or my own—are labeled *synthetic*. Their basic axioms carry elementary *physical* meaning in the theory.²⁰ This feature is what allows us to interpret the resulting mathematical structure physically, and thus provide a metaphysically perspicuous picture of the underlying ontology.

In Section 3.2.a and in Section 3.2.b we will discuss synthetic axiomatizations of spacetime metrics and connections in principal fiber bundles, respectively.

3.2.a The case of spacetime metrics.

Here I will focus on two related approaches to axiomatization: ([Ehlers et al., 2012](#)) and ([Mundy, 1992](#)).²¹

²⁰In [Adlam et al. \(2022\)](#) such approaches are labeled ‘constructive’, but for philosophers that term strongly connotes intuitionism, *a lá* Brouwer and cousins. According to [Adlam et al. \(2022\)](#), such approaches should be distinguished from the mere *deductive approaches*, that find uniqueness results from more basic postulates that don’t necessarily carry a physical interpretation.

²¹[Mundy \(1992, p.518\)](#) admits kinship with [Ehlers et al. \(2012\)](#), and even describes the latter as “deeper”, but also condemns their argument as “sketchy and incomplete”. I do not aim to assess these matters: this Section is illustrative only. For a complete appraisal of EPS and kindred synthetic approaches, see [Adlam et](#)

We start with the seminal Ehlers et al. (2012), who set out to successively build differential topology, conformal, affine, and Lorentzian structure, from basic postulates concerning the paths of light and massive particles (see also Linnemann & Read (2021) for a pedagogical introduction to this scheme of Ehlers, Pirani and Schild (henceforth EPS)). According to (Linnemann & Read, 2021, p.2), the EPS scheme is “constructively axiomatic in the sense of Reichenbach— i.e., builds on a basis of empirically supposedly indubitable posits.” The two main initial posits define: (1) A point set $M = \{p, q, \dots\}$, called a set of events; and (2) The elements of $L = \{L, N, \dots\}$ and $P = \{P, Q, \dots\}$, where L, N, \dots, P, Q, \dots are all subsets of M , and are respectively called ‘light rays’ and ‘particles’. Of course, it is too early to call L and P ‘light rays’ and ‘particles’: they will only acquire this meaning once they are subject to further empirically-motivated posits, such as the assumption that one can bounce a light ray from the worldline of one particle to that of another and back again, etc. Indeed, (M, P, L) only acquire their standard interpretation at the end of the construction, once they have been employed to define increasing levels of structure (smooth, affine, projective, etc).

But let me focus on this Section on Mundy (1992), who also seeks an intrinsic geometric axiomatization of semi-Riemannian geometry. Such an axiomatization follows in the tradition of many others (e.g. Robb (1936), for the case of Minkowski space). Thus, Mundy (1992, p. 517) uses the intrinsic geometric axiomatization of Euclidean geometry as a template. There, we use coordinate-free primitives such as the *affine betweenness* relation $B(p, q, r)$, read as ‘ q is between p and r on a straight line-segment’, the *segment congruence* relation $C(p, q, r, s)$, read as ‘the straight line-segments between p and q and between r and s are congruent’, etc. to formulate a theory T_E in a language L_E including these primitives relations. Equivalence with the coordinate formulations is shown by a representation theorem: that each model of T_E has Cartesian coordinates, unique up to orthogonal transformations, representing B and C by the standard coordinate formulas.

Analogously, Mundy’s axiomatization of (semi-)Riemannian geometry proceeds in terms of “truly coordinate-free primitives” in which C refers to metric behavior of clocks and rods, and B refers to geodesic motion. In analogy to the Euclidean case, one replaces straight lines by geodesics, and has $B(p, q, r)$ as the relation of *geodesic betweenness* and $C(p, q, r, s)$ is a notion of *geodesic congruence*. Analogously to the Euclidean case, a third intrinsic relation, $A(p, q, r)$, of *geodesic orthogonality*, is defined infinitesimally in terms of B and C , by using infinitesimal trigonometry. The same can be done with $S(p, q, r, p'q')$, the notion of *projective separation* (which I will not recount here). The general idea is that (semi-)Riemannian geometry can, as regards its local structure, be developed in a synthetic manner, along similar lines to Euclidean geometry.

The upshot is that this language, L_R has no sentence which would be made true by a model $S := \langle M, g \rangle$, but not by $Sf := \langle M, f^*g \rangle$ (or vice versa). In the words of Mundy (1992, p. 520):

the difference between the two models [S and Sf] is not expressible in the theoretical language L_R : they are isomorphic, and therefore satisfy all of the same sentences.

al. (2022).

(Note that this is theoretical, not merely observational, equivalence.) [...] This is because the language L_R does not and cannot contain any term which “rigidly designates” the point p itself [i.e. is non-qualitative], i.e., which refers to p when L_R is interpreted over S , and also refers to p when L_R is interpreted over Sf .

Thus, while the metalanguage is able to express singular, non-qualitative features, L_R is not. Isomorphisms, expressed in the metalanguage, will map objects singled out by the same description into each other. Thus Mundy’s formal semantics under some interpretation defines a structure admitting a notion of isomorphism that coincides with isometry.

I take the synthetic approach to spacetime geometry illustrated here (twice) to provide a structure-first definition of the theory; and I also take it to answer the metaphysician’s misgivings above, about a qualitative description of spacetime structure.

To nail the point home, take Dasgupta (2011, p. 147)’s first challenge as representative of that concern.²² He proposes (p. 149, *ibid*) that a solution would need a language “PL of predicate logic with identity but no constants, in which every predicate expresses a qualitative property or relation. Every fact that can be expressed in PL is a purely qualitative fact.”

Both (Mundy, 1992) and (Ehlers et al., 2012) furnish us with this kind of language: certain relations between events, such as ‘being (non)causally connected’, or ‘having the same length’, or even ‘being connected by a massive particle’s worldline’, are primitive and qualitative, just like ‘having non-zero absolute velocity’ is a qualitative property in Newtonian mechanics (cf. footnote 8); or just like the sentence “someone loves someone” in predicate logic.

As to a metaphysically perspicuous interpretation, the constraints on these languages allow solely for the expression of isomorphism-invariant (i.e. diffeomorphism-invariant) quantities, whose interpretations are by construction attached to primitive physical posits, about point-coincidences, proper times along worldlines, etc.

In sum, jointly, we physically interpret general relativity’s fulfillment of Desiderata (i) and (ii) as taking diffeomorphism-invariant mathematical quantities to represent physically significant quantities, understood as quantities about coincidences of material point-particles, elapsed proper times along a particle worldline, etc. Though our models of these physical features are defined only up to isometry, this poses no threat to determinism: we can identify physical possibilities as given by an invariant structure which can be interpreted anti-haecceitistically.

3.2.b The case of connections on a principal bundle.

Similarly, we have two examples of axiomatization for Yang-Mills theory. These are at the very least as intricate as the previous case of general relativity. But for reasons of space, I will only briefly gloss the two constructions here.

The first is based on Barrett (1991)’s construction of a general connection ω on a principal fiber bundle, P , as discussed in (Gomes, 2021b, Sec. 4). Here, besides the spacetime structure,

²²As quoted at the start of Section 2.2, the challenges are: “(1) to clearly articulate what the underlying qualitative facts are like, and (2) show that they are sufficient to explain (in the metaphysical sense) individualistic facts about the manifold.

which could be constructed as in the previous Section, we can take the basic axioms to have physical meaning. The physical meaning of these postulates follow the prescription sketched in [Yang & Mills \(1954\)](#)'s quote at the end of [\(Gomes, 2021b, Sec. 4\)](#), and argued at length in that Section. That is, the primitive, empirically-motivated axioms, refer to the relative rotation of a charge, obtained by dragging it (i.e. parallel transporting it) along closed curves of the spacetime manifold and comparing the result with the initial charge.²³ In Yang-Mills theory, this overall rotation is called *the holonomy* (we will discuss it in more detail in [Gomes \(2021c\)](#); see [Healey \(2007\)](#) for a philosophical introduction). From these basic quantities, we can construct a principal fiber bundle and, on it, a unique Ehresmann connection-form ω (cf. [\(Gomes, 2021b, Sec. 4\)](#)), up to vertical automorphisms (the isomorphisms of the principal fiber bundle structure).

In more detail, let M be the spacetime manifold, G be a Lie group, and γ be a loop based at point x , namely $\gamma : [0, 1] \rightarrow M$, with $\gamma(0) = \gamma(1) = x$. The space of all such loops we call $\text{Loop}(M)$, and we postulate a map:

$$\text{hol} : \text{Loop}(M) \rightarrow G, \tag{3.7}$$

which we take as primitive, and obeying three postulates, which can all be empirically motivated (see [\(Barrett, 1991, Sec. 2.4\)](#)):

(H1) hol is a homomorphism of the composition law of loops and group multiplication:

$$\text{hol}(\gamma_1 \circ \gamma_2) = \text{hol}(\gamma_1)\text{hol}(\gamma_2).$$

(H2) hol takes the same value on loops which differ by a reparameterisation, or by the addition or removal of path sections which ‘double back’ on themselves. Thus, if $\gamma_1 \circ \gamma_2$ encloses zero area,

$$\text{hol}(\gamma_1 \circ \gamma_2) = \text{Id}.$$

Both these postulates have a straightforward empirical interpretation, used in the construction theorem to interpret hol as the parallel transport operation for internal vectors.

(H3) The last condition requires topological structures in the space of loops; general constraints select a unique family of topologies. Given one such specific topology and a smooth family of loops, say $\ell : \mathbb{R} \rightarrow \text{Loop}(M)$ (we could replace \mathbb{R} for any other parametrizing space, e.g. \mathbb{R}^2 , for a 2-parameter smooth family of loops), then:

$$\text{hol} \circ \ell \quad \text{is also smooth.}$$

Barrett interprets this as a condition on physical approximations: if we were to physically vary the loop corresponding to a holonomy slightly, the end result of the holonomy must also be varied only slightly.

²³In practice, this would be accomplished with both charges undergoing time-like motion, with their joint trajectories forming a closed loop.

With these postulates, Barrett proves a *construction* and a *presentation* theorem,²⁴ showing, respectively that: (1) for M a connected manifold with basepoint x , and $hol : \text{Loop}(M) \rightarrow G$ satisfying conditions H1-H3, then there exists a differentiable G -principal fibre bundle P , a point $p \in \pi^{-1}(x)$ and a connection ω on P such that hol is the holonomy mapping of this bundle and connection; (2) the presentation theorem shows that for any (P, ω) such that $\pi : P \rightarrow M$, there is a holonomy mapping on M that recovers (P, ω) up to a vertical isomorphism. And indeed, Barrett (1991) also construes his theorems as empirically motivated.²⁵

The second sort of axiomatization—which we will here only very superficially gloss—is based on the joint axiomatization of spacetime and of a Lie-algebra; and the latter proceeds in a similar fashion to the axiomatization of any finite-dimensional vector space, interpreted as a space of internal vectors (e.g. colours or charges). That is, the Yang-Mills connection is an object that mixes tensorial indices with internal indices. The natural principal bundle for the tensorial part, (see (Weatherall, 2016, Sec. 3) and (Gomes, 2021b, Sec. 4)), would be a sub-bundle of the frame bundle $L(TM)$; and the internal part to the vector bundle associated to \mathfrak{g} . These two bundles are of different types, and to make a joint object out of them, we need ‘bundle splicing’ (see e.g. (Bleecker, 1981, Ch. 7.1)). The structure that emerges from this splicing is known as the ‘bundle of connections’, or the Atiyah-Lie algebroid, and it represents sections of TP/G , which are in 1-1 correspondence with Ehresmann connections, ω . See (Kolar et al., 1993, Ch. 17.4) for the construction alluded to here, which is discussed in further detail in (Gomes, 2021b, Sec. 6) and employed in Gomes (2021c).

In sum, the moral, for the two types of theory, is the same: The isomorphisms that appear in the higher level language can be understood through a representation theorem for models of the more basic ontology of the theory. Once we have the Ehresmann connection ω , and a spacetime metric g_{ab} , built from more fundamental, empirically motivated posits, we can stop appealing to this more cumbersome axiomatic language and proceed in the standard textbook (metalanguage) terms; thereby we recover the isomorphisms of the theories (used in Assumptions II and III of the introduction to Section 3).

²⁴In fact, he calls them *reconstruction* and *representation*, but these are used in the opposite sense of Mundy (1992): for instance, his reconstruction goes from the physically transparent synthetic theory to various representations in bundles.

²⁵And the empirical motivation also applies for the reconstruction of the spacetime connection through parallel transport along loops (or holonomies):

The information in the gravitational field is involved in grouping together, in a particular way which will be described more fully below, all the possible particle paths into sets which represent the particles whose paths are “in coincidence”. These sets form the points of the space-time manifold. This is actually all that is required! What one finds is that the geometry of the particle paths is sufficient to specify the geometry of the resulting space-time. The result is a theory of gravity in which the test particles are firmly mixed up with the phenomenon of the field itself. In fact the field is “made out of” the motions of the test bodies. There is no conception of an independent gravitational field divorced from the theory of the propagation of matter. (p.22, *ibid*)

The idea of the extension of sophistication to more general theories then, is that, were they to satisfy Desiderata (i) and (ii), they would enjoy the same kind of metaphysically perspicuous interpretation of the symmetry-invariant ontology, and that anti-haecceitism/anti-quidditism would suffice to physically identify symmetry-related models.

4 Summing up

This paper has developed a view of sophistication that could be extended to more general theories and that applies to both spacetime diffeomorphisms and gauge symmetries.

The view was employed to defend the doctrine of anti-quidditism for gauge theory, but, more emphatically, the doctrine of anti-haecceitism for spacetime. For anti-haecceitism is still actively disputed today: in essence, the concern is that it could be a mere stipulation that does not bear the metaphysical interpretation foisted upon it; for the doctrine simultaneously deflates the individuation of spacetime points and reaffirms the existence of spacetime. That concern, I argued, has a close cousin in the modern literature about sophistication: the concern that the ontology that is common to all of the symmetry-related models is not sufficiently explicated by stating that symmetry-related possibilities are to be identified—*that* is also mere stipulation!

My Desideratum (i) for sophistication guaranteed that physical indeterminism could be eliminated by anti-haecceitism, or, more generally, by commitment to a purely qualitative ontology. By satisfying (i), a theory's isomorphisms emerge from the automorphisms of the base set of the theory's models. Thus isomorphic models differ only as to *which part* of the base set supports *which pattern*: isomorphic distributions are 'qualitatively' identical, in a precise sense of qualitative. By requiring the isomorphisms to be related to the dynamical symmetries of the theory and to the automorphisms of the base set, Desideratum (i) constrains, and may suffice to entirely regiment, the idea of 'internal' sophistication. And it guarantees that there are no empirical differences between isomorphic models. But this Desideratum takes the space of models Φ for granted, and with it a notion of isomorphism of the models. Desideratum (i) does not *start from* the invariant structure—it does not, strictly speaking, take the 'structure first'—and it falls short of ensuring a 'metaphysically perspicuous' construal of the structure that is common to the isomorphism classes.

Fortunately, Desideratum (ii) for sophistication picks up where Desideratum (i) stops. According to Desideratum (ii) we should understand the invariant structure first, 'internally': or rather, synthetically. This is accomplished explicitly by requiring the invariant structure to be defined axiomatically. The isomorphisms then relate different representative models of the same structure. Moreover, the axioms explicitly endow the invariant structure with a perspicuous (meta)physical interpretation, since they are expressed using basic physical posits of the given theory. In this way, we conceptually clarify the isomorphism-invariant mathematical structure while connecting that structure to the basic ontology of the theory.

Were a theory to fill Desideratum (ii) but not Desideratum (i), we would not be able to

interpret the invariant structure ‘qualitatively’ (at least in a well-regimented sense). And we would not be able to say whether an ontological commitment to the invariant structure would suffice to eliminate the threat of physical indeterminism.

I take this to close the question of whether we should endorse sophistication for spacetime diffeomorphisms, but not for gauge symmetries; that is, whether we should, in these cases, be (resp.) anti-haecceitists but not anti-quidditists. In both cases, and others that fulfill Desiderata (i) and (ii) as well, we should be ‘anti-ists’.

I will end this paper with a view towards the next in the series, Gomes (2022), by turning the gaze of our investigation ‘inward’, that is, to objects *within* each model. While sophistication as construed here also implies objects within each model can be individuated only by the qualitative properties which the models tack onto them, the doctrine is silent about what is required for uniquely specifying objects inside each physical possibility. This is related to the second worry, about structure-tokens, glimpsed in Section 1.2, and it will be the topic of the third paper Gomes (2022).

Acknowledgements

I would like to thank: first and foremost Jeremy Butterfield for many conversations on this topic, and for reading several versions of the paper; Ruward Mulder, for a careful reading and comments; and Caspar Jacobs, for reading and providing encouragement and suggestions.

References

- Adams, R. M. (1979). Primitive Thisness and Primitive Identity. *Journal of Philosophy*, 76(1), 5–26. doi: 10.2307/2025812
- Adlam, E., Linnemann, N., & Read, J. (2022). Constructive Axiomatics in Spacetime Physics Part II: Constructive Axiomatics in Context. *In preparation*.
- Barrett, J. W. (1991, Sep 01). Holonomy and path structures in general relativity and yang-mills theory. *International Journal of Theoretical Physics*, 30(9), 1171–1215. Retrieved from <https://doi.org/10.1007/BF00671007> doi: 10.1007/BF00671007
- Belot, G. (2003). Symmetry and gauge freedom. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 34(2), 189 - 225. Retrieved from <http://www.sciencedirect.com/science/article/pii/S1355219803000042> doi: [https://doi.org/10.1016/S1355-2198\(03\)00004-2](https://doi.org/10.1016/S1355-2198(03)00004-2)
- Belot, G. (2013). Symmetry and Equivalence. In *The Oxford Handbook of Philosophy of Physics*. Oxford University Press. Edited by Batterman, R.

- Belot, G. (2018). Fifty Million Elvis Fans Can't be Wrong. *Nous*, 52(4), 946-981. Retrieved from <https://onlinelibrary.wiley.com/doi/abs/10.1111/nous.12200> doi: 10.1111/nous.12200
- Belot, G., & Earman, J. (1999). From metaphysics to physics. In J. Butterfield & C. Pagonis (Eds.), *From physics to philosophy* (p. 166–186). Cambridge University Press. doi: 10.1017/CBO9780511597947.009
- Belot, G., & Earman, J. (2001). Pre-Socratic quantum gravity. In C. Callender & N. Huggett (Eds.), *Physics meets philosophy at the planck scale: Contemporary theories in quantum gravity* (p. 213-255). Cambridge University Press. doi: 10.1017/CBO9780511612909.011
- Black, M. (1952). The Identity of Indiscernibles. *Mind*, 61(242), 153–164. Retrieved 2022-07-15, from <http://www.jstor.org/stable/2252291>
- Bleecker, D. (1981). *Gauge theory and variational principles*. Dover Publications.
- Brighouse, C. (1994). Spacetime and holes. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association, 1994*, 117–125. Retrieved from <http://www.jstor.org/stable/193017>
- Butterfield, J. (1989, 03). The Hole Truth. *The British Journal for the Philosophy of Science*, 40(1), 1-28. Retrieved from <https://doi.org/10.1093/bjps/40.1.1> doi: 10.1093/bjps/40.1.1
- Curiel, E. (2018). On the Existence of Spacetime Structure. *The British Journal for the Philosophy of Science*, 69(2), 447-483. Retrieved from <https://doi.org/10.1093/bjps/axw014> doi: 10.1093/bjps/axw014
- Dasgupta, S. (2011). The bare necessities*. *Philosophical Perspectives*, 25(1), 115-160. Retrieved from <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1520-8583.2011.00210.x> doi: <https://doi.org/10.1111/j.1520-8583.2011.00210.x>
- Dewar, N. (2017, 09). Sophistication about Symmetries. *The British Journal for the Philosophy of Science*, 70(2), 485-521. Retrieved from <https://doi.org/10.1093/bjps/axx021> doi: 10.1093/bjps/axx021
- Earman, J. (1989). *World enough and spacetime*. MIT press.
- Earman, J. (2002). Gauge matters. *Philosophy of Science*, 69(S3), S209-S220. Retrieved from <https://doi.org/10.1086/341847> doi: 10.1086/341847
- Earman, J. (2004). Curie's principle and spontaneous symmetry breaking. *International Studies in the Philosophy of Science*, 18(2 & 3), 173–198. doi: 10.1080/0269859042000311299

- Earman, J., & Norton, J. (1987, 12). What Price Spacetime Substantivalism? The Hole Story. *The British Journal for the Philosophy of Science*, 38(4), 515-525. Retrieved from <https://doi.org/10.1093/bjps/38.4.515> doi: 10.1093/bjps/38.4.515
- Ehlers, J., Pirani, F., & Schild, A. (2012). Republication of: The geometry of free fall and light propagation. *General Relativity and Gravitation*, 44, 1587–1609.
- Fletcher, S. C. (2020). On Representational Capacities, with an Application to General Relativity. *Foundations of Physics*, 50(4), 228–249. doi: 10.1007/s10701-018-0208-6
- Gomes, H. (2021a). The role of representational conventions in assessing the empirical significance of symmetries. (*to appear in Studies in History and Philosophy of Modern Physics*).
- Gomes, H. (2021b). Same-diff? Conceptual similarities (and one difference) between gauge transformations and diffeomorphisms. Part I: Symmetries and isomorphisms. *Arxiv: 2110.07203. Submitted..*
- Gomes, H. (2021c). Same-diff? Part IV: A compendium of similarities between gauge transformations and diffeomorphisms. *Arxiv: 2110.07204. Submitted..*
- Gomes, H. (2022). Same Diff?
Conceptual similarities between gauge transformations and diffeomorphisms.
Part III: Representational conventions and relationism. *Unpublished.*
- Gomes, H., & Butterfield, J. (2022). Assessing the hole argument. *In preparation.*
- Hawking, S. W., & Ellis, G. F. R. (1975). *The Large Scale Structure of Space-Time (Cambridge Monographs on Mathematical Physics)*. Cambridge University Press. Retrieved from <http://www.amazon.com/Structure-Space-Time-Cambridge-Monographs-Mathematical/dp/0521099064>
- Healey, R. (2007). *Gauging What's Real: The Conceptual Foundations of Gauge Theories*. Oxford University Press.
- Hetzroni, G. (2021). Gauge and ghosts. *The British Journal for the Philosophy of Science*, 72(3), 773-796. Retrieved from <https://doi.org/10.1093/bjps/axz021> doi: 10.1093/bjps/axz021
- Hoefer, C. (1996). The Metaphysics of Space-Time Substantivalism. *The Journal of Philosophy*, 93(1), 5–27. Retrieved from <http://www.jstor.org/stable/2941016>
- Jacobs, C. (2021a). Invariance or Equivalence: A Tale of Two Principles. *Synthese*, 1–21. doi: 10.1007/s11229-021-03205-5
- Jacobs, C. (2021b). The metaphysics of fiber bundles. *Forthcoming in British Journal of Philosophy of Science*.

- Jacobs, C. (2021c). *Symmetries as a Guide to the Structure of Physical Quantities* (Unpublished doctoral dissertation). University of Oxford.
- Jacobs, C. (2022). Symmetries as a guide to the ontology of physical theories. *In preparation*.
- Kaplan, D. (1979). Transworld heir lines. In M. J. Loux (Ed.), *The possible and the actual* (pp. 88–109). Ithaca: Cornell University Press.
- Klein, F. (1893). Vergleichende Betrachtungen über neuere geometrische Forschungen. *Math. Ann.* 43.
- Kolar, I., Michor, P., & Slovák, J. (1993). *Natural Operations in Differential Geometry*. Springer.
- Linnemann, N., & Read, J. (2021). *Constructive axiomatics in spacetime physics part i: Walkthrough to the ehlers-pirani-schild axiomatisation*.
- Martens, N. C., & Read, J. (2020). Sophistry about symmetries? *Synthese*. Retrieved from <http://philsci-archive.pitt.edu/17184/>
- Maudlin, T. (1993). Buckets of water and waves of space: Why spacetime is probably a substance. *Philosophy of Science*, 60(2), 183–203. Retrieved from <http://www.jstor.org/stable/188350>
- Mundy, B. (1992). Space-Time and Isomorphism. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association, 1992* (Volume One: Contributed Papers), 515–527.
- Møller-Nielsen, T. (2017). Invariance, interpretation, and motivation. *Philosophy of Science*, 84(5), 1253-1264. Retrieved from <https://doi.org/10.1086/694087> doi: 10.1086/694087
- Norton, J. (1989). Coordinates and covariance: Einstein’s view of space-time and the modern view. *Foundations of Physics*, vol 19.
- O’Neill, B. (1983). *Semi-Riemannian Geometry With Applications to Relativity*. Academic Press.
- Pooley, O. (2022). *The Reality of Spacetime*. Unpublished manuscript.
- Pooley, O. (in press). The Hole Argument. In E. Knox & A. Wilson (Eds.), *The routledge companion to the philosophy of physics*. Routledge.
- Pooley, O., & Read, J. (2022). On the Mathematics and Metaphysics of the Hole Argument. *The British Journal for the Philosophy of Science*. doi: 10.1086/718274
- 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*, 70, 88-97. Retrieved from <https://www.sciencedirect.com/science/article/pii/S1355219819301649> doi: <https://doi.org/10.1016/j.shpsb.2020.03.002>

- Robb, A. A. (1936). *Geometry of Time and Space*. Cambridge University Press.
- Roberts, B. (2020). Regarding ‘Leibniz Equivalence’. *Foundations of Physics*, 50(4), 250–269. doi: 10.1007/s10701-018-0208-6
- Roberts, B. W., & Weatherall, J. O. (2020). New Perspectives on the Hole Argument. *Foundations of Physics*, 50(4), 217–227. doi: 10.1007/s10701-020-00334-8
- Rynasiewicz, R. (2022). Newton’s Views on Space, Time, and Motion. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Spring 2022 ed.). Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/spr2022/entries/newton-stm/>.
- Teitel, T. (2019). Holes in Spacetime: Some Neglected Essentials. *Journal of Philosophy*.
- Wald, R. M. (1984). *General Relativity*. University of Chicago Press.
- Wallace, D. (2019). *Isolated systems and their symmetries, part i: General framework and particle-mechanics examples*. Retrieved from <http://philsci-archive.pitt.edu/16623/>
- Weatherall, J. (2016). Fiber bundles, Yang–Mills theory, and general relativity. *Synthese*, 193(8), 2389–2425. (<http://philsci-archive.pitt.edu/11481/>)
- Weatherall, J. (2018). Regarding the ‘Hole Argument’. *The British Journal for the Philosophy of Science*, 69(2), 329–350. Retrieved from <https://doi.org/10.1093/bjps/axw012> doi: 10.1093/bjps/axw012
- Wüthrich, C. (2009). Challenging the Spacetime Structuralist. *Philosophy of Science*. **Volume 76**, Number 5.
- Yang, C. N., & Mills, R. L. (1954, Oct). Conservation of isotopic spin and isotopic gauge invariance. *Phys. Rev.*, 96, 191–195. Retrieved from <https://link.aps.org/doi/10.1103/PhysRev.96.191> doi: 10.1103/PhysRev.96.191