Representational schemes for theories with symmetry

Henrique de A. Gomes

University of Oxford
Oriel College, Oxford, OX2 1TQ, UK

April 22, 2025

Abstract

In the philosophical literature, symmetries of physical theories are most often interpreted according to the general doctrine called 'traditional sophistication' (TS). But even this doctrine leaves two important gaps in our understanding of such theories: (A) it allows the individuation of isomorphism-classes to remain intractable and thus of limited use, which is why practising physicists frequently invoke 'relational, symmetry-invariant observables'; and (B) it leaves us with no formal framework for expressing interesting counterfactual statements about different physical possibilities. I will call these *Limitations* of TS. Here I will show that a new Desideratum to be satisfied by theories with symmetries allows us to overcome these Limitations. The new Desideratum is that the theory admits what I will call *representational schemes* for its isomorphism-classes. Each such scheme gives an equally valid reduced formalism for a theory.

Contents

1 Introduction

This paper is about how theories of physics deal with symmetries. I will focus on theories whose equations of motion can be defined variationally from a Lagrangian, and for which we can define dynamical symmetries as 'a group of transformations whose action on the models leaves the Lagrangian invariant up to boundary terms' (cf. (Gomes, 2021d, p. 5)).[1]

[1]The group of transformations is described independently of the models, but its action can be model-dependent, i.e. it may act differently on different models.

Traditional Sophistication (henceforth TS) is a doctrine about how to view symmetry-related models (cf. (Martens & Read, 2020, p. 323)). At a very broad level, the doctrine says that, if dynamical symmetries can be understood as mathematical isomorphisms of the models of the theory, then any two models that are related by a dynamical symmetry can be taken to represent the same physical possibility.

However, as widely recognised (cf. (Jacobs, 2022; Martens & Read, 2020) for recent appraisals), 'isomorphism' is a very flexible notion, and may in effect not constrain the set of dynamical symmetries that can be interpreted via TS. Nonetheless, TS *is* compelling for—and usually understood to apply only to—theories that exhibit a particular mathematical feature. To exhibit this feature (called Criterion (ii) in (Gomes, 2021c)),[2] a theory must first of all have its models expressed by values of dynamical variables that are functions of a set of independent variables. This set constitutes a fixed, 'absolute', or non-dynamical background which is endowed with some structure of its own. For spacetime theories such as Newtonian mechanics, special or general relativity, it is a spacetime manifold endowed with a geometric structure, such as: Newtonian absolute space and time, a Minkowski metric, or just a smooth structure, respectively in the three cases (see (Earman, 1989, Ch. 2) for a longer and more detailed list). But we could also have models of theories that incorporate other fields, such as those of the Standard Model of particle physics, for which the non-dynamical background structure doesn't pertain solely to spacetime. The fixed background structure of these theories includes 'internal' spaces over each spacetime point (they are called 'a vector bundles', cf. (Gomes, 2024)).

Finally, we can state Criterion (ii): dynamical symmetries are uniquely represented by automorphisms of the models' common background structure—e.g. the diffeomorphisms of a smooth manifold; the fiber-preserving linear automorphisms of a vector bundle, etc. Satisfaction of Criterion (ii) makes TS compelling since then, even if isomorphic models distribute the values of the fields differently over a base set, that base set is *structurally indifferent* to these differences. In other words, its satisfaction ensures isomorphic models can only differ with respect to which objects we arbitrarily choose to represent which values.[3]

[2]Criterion (i) in (Gomes, 2021c) requires dynamical symmetries to be in a 1-1 relation with the isomorphisms of some well-defined mathematical structure common to all of the models of the theory. As I said above, this Criterion is very easily satisfied, due to the vast resources of mathematics and flexibility of what constitutes a well-defined mathematical structure. Criterion (ii) is a very straightforward generalization of (Earman, 1989, p. 45)'s famous 'SP principles' for spacetime symmetries, as also described in (Jacobs, 2022; Martens & Read, 2020).

[3]Suppose isomorphisms of the models *were* induced by bijections of the background structure, but assume one such bijection is *not* an automorphism of the background structure. One could argue that two models related by the induced transformation would still only differ insofar as they distributed properties differently over the background. But in such a case, the background structure is *sensitive*, not indifferent, to the differences between the transformed models. This is what occurs, for example, for boosts in a system of Newtonian particles: boosts can be represented by bijections of a Newtonian spacetime and their action on the configuration yields dynamical symmetries, but two boost-related models have different states of motion with respect to the background absolute frame. In contrast, a static shift, shifting the origin of space by a certain amount in any direction, would preserve the background absolute frame and thus would satisfy Criterion (ii) cf. (Maudlin, 1990).

Thus, one natural way of interpreting TS is through *anti-haecceitism*: a metaphysical doctrine that denies that objects have a primitive 'thisness', and so cannot differ solely as to which objects represent which qualitative properties. Under anti-haeccetism, objects can only be individuated by their qualitative properties, such as their place in the network of relations to other objects.

Here, I am mostly concerned with two types of theory that satisfy Criterion (ii), and so for which TS is compelling: (a) general relativistic theories and (b) gauge, i.e. Yang-Mills, theories. In (a) dynamical symmetries of the Einstein-Hilbert Lagrangian are isometries of Lorentzian metrics, which are generated by smooth diffeomorphisms of the underlying smooth manifold where the metric is defined (and so satisfy Criterion (ii)). Similarly, in (b) dynamical symmetries of the Standard Model Lagrangian give a kind of isomorphism of models, called a *gauge transformation*, that can be construed as spacetime-dependent fiber-preserving linear isomorphisms for (internal) vector spaces attached to each spacetime point (and so similarly satisfy (ii)).

I take TS to significantly illuminate our understanding of such theories. However, there are certain questions about symmetry and physical equivalence that are not touched on by TS but are nonetheless important. For lack of a better name, I call these questions 'limitations' of TS. Here I will argue that we can answer these questions, or overcome these limitations, by a 'relationist' approach; one that depends on TS but also goes beyond it.

There are three main such Limitations, of which I here aim to overcome only two (which are aspects of what Gomes (2021c) called Worry (2)), which I will label Limitations (A) and (B):[4]

(A) TS gives us a sufficient condition of identification of what two models represent about the world (or contrapositively: a necessary condition of non-identification). That condition is the existence of a suitable isomorphism between the models. Thus TS commits the theory to a symmetry-invariant ontology, but it does not require the theory to have explicit, symmetry-invariant models of that ontology. Without such models, at a practical level, TS's sufficient condition of identification falls short. For given two different models, the sufficient condition is not, *prima facie*, tractable: one cannot try out every diffeomorphism—or, more generally, automorphism of the models' common background structure—to see if one will bring the two models into coincidence. This Limitation motivates many theoretical physicists to seek relational characterizations of the symmetry-invariant ontology; even while generally endorsing TS (as described in (Belot & Earman, 1999, 2001)).

[4]I will set aside what Gomes (2021c) called Worry (3): as argued persuasively by Belot (2018), there is a consensus among physicists that, for certain sectors of general relativity, TS is wrong, because some isomorphisms of models relate *different* physical possibilities. Namely, in the context of spatially asymptotically flat spacetimes, some diffeomorphisms—those that preserve the asymptotic conditions but don't asymptote to the identity map—are interpreted as relating different physical possibilities. The same occurs in Yang-Mills theory (Giulini, 1995). But the reasons given for TS (see (Gomes, 2021c)) were general; they applied at the level of the whole theory. So TS, as a doctrine about the whole theory, i.e. all its sectors, seems refuted.

(B) By effectively denying that objects have primitive identities, TS says nothing about any correspondence between spacetime points that belong to non-isomorphic spacetimes. Similarly, for gauge theory, TS relinquishes any notion of correspondence between internal directions of vector bundles belonging to non-isomorphic models. But such correspondences are crucial for counterfactuals to be expressed and interpreted. And beyond counterfactuals, these correspondences are crucial for making rigorous sense of the idea of superpositions of classical spacetimes; a topic that is central to various efforts to combine quantum mechanics with gravity (for recent discussion of why such correspondences are needed in order to thus combine the fileds, cf. (Kabel et al., 2024) and references therein).

In this paper it will turn out that answering Limitation (A) in a certain way will also answer Limitation (B). So I begin in Section 2 by elaborating (A). I will there explain why these Limitations are not merely metaphysical: they are very relevant for current issues in theoretical physics. Then, in Section 3, I introduce the idea of *representational schemes*.[5] Conceptually, such schemes are choices of (idealised) physical relations with which we can completely describe each physical possibility without redundancy. So each scheme gives an equally valid reduced formalism for a given theory, and (a sub-class of) such schemes will suffice to answer both (A) and (B).

In Section 4 I discuss how choices of representational schemes are possible in the absence of an empirical breaking of the symmetry of the theory. In Section 5 I will develop the theory of counterparthood based on representational schemes and show that overcoming Limitation (A) via (a sub-class of) representational schemes also overcomes Limitation (B). That is, the resolution of Limitation (A) via certain types of representational schemes offers not only tractable conditions of identity for the entire universe, they also provide local correspondence relations between spacetime points (and, respectively, for internal vector spaces, in the case of gauge theory) in non-isomorphic models. And in Section 6 I conclude by describing a limitation of representational schemes and a possible route to overcome it. In Appendix A I provide several examples of representational schemes; focusing on the case where these schemes are choices of gauge-fixing, I will illustrate some of their important properties, such as non-locality.

2 The challenge of individuating structural content

Here is how I will organise this Section. In Section 2.1 I will expound Limitations (A) and (B). In Section 2.2 I will relate Limitation (A) to issues in the metaphysics of spacetime, but argue that the Limitation is not solely metaphysical: it has echoes on current theoretical physics. In Section 2.3 I will try to narrow down what could count as a resolution of Limitation (A); this is where I first introduce *representational schemes*, which will be developed in the rest of the paper.

[5]These were previously called 'representational conventions' (cf. (Gomes, 2022)). But a 'convention' overemphasises arbitrariness, which is a misleading connotation; whereas 'schemes', while still signalling that a choice has to be made, connote that this choice must satisfy certain constraints.

2.1 Limitations A and B in more detail

2.1.a Limitation (A): the individuation of isomorphism-classes

TS provides a sufficient condition for two models to represent the same physical possibility. As argued in recent discussions (cf. (Fletcher, 2020; Pooley & Read, 2022)), to assess this condition in any specific case one needs to first fix a context in which the two models are being applied. Within that context, the sufficient condition is then very simple: it is satified iff there is an isomorphism between the representing models.

But here is Limitation A: even once we have fixed a notion of isomorphism and a context of application for the theory, we may not have a tractable way to individuate the *isomorphism-classes*. That is, we lack an explicit isomorphism-invariant description of the models, so how should we, in practice, assess whether two given models are isomorphic? For sometimes it is exceedingly hard, via direct inspection of all candidate automorphisms of the background structure, to tractably verify whether any two given models are isomorphic or not. That is why mathematicians often seek a condition equivalent to the existence of isomorphisms that is easier to verify.

Indeed, the classification problem for a type of mathematical object is a very familiar topic within mathematics. It is sometimes trivial—e.g. finite-dimensional vector spaces are isomorphic iff they have the same dimensions—and it is sometimes hard, e.g. seeking all the topological invariants of topological spaces. (Note here that 'all' means that two spaces having the same set of invariants implies that the two spaces are homeomorphic.)[6]

An analogy to purely philosophical discussions may be helpful here. Thus we can take Leibniz's venerable Principle of the Identity of Indiscernibles (PII) to provide a sufficient condition of identity of objects: viz. if objects x and y have exactly the same qualitative properties, then $x = y$.[7] But this sufficient condition is almost always intractable, since we cannot check it for *all* the properties of an object (even just the qualitative ones). Hence the philosophical endeavour to formulate identity conditions for a specific sort of object, i.e. tractable sufficient conditions for identity (e.g. the claim that, as we have learned from crime dramas, for humans, the identity of fingerprints, or dental records, is sufficient for personal identity). So the analogy is that: (i) Leibniz's PII is like the way that the definition of a mathematical structure fixes, *ipso facto*, a sense of isomorphism, but gives no guidance about how to ascertain whether two structures are isomorphic; so that (ii) the search for tractable sufficient conditions for identity is like the mathematician's seach for a solution to the a classification problem.

As an example, in general relativity, we may postulate that models are Lorentzian manifolds

[6]For instance, mathematicians place great importance in theorem's like Perelman's, which shows that all three-dimensional compact manifolds with a certain homotopy group are homeomorphic to the three-dimensional sphere.

[7]There are many ways to interpret every term in this abbreviated description: the scope of 'objects', the meaning of 'qualitative', etc. In my usage so far, qualitative properties correspond to those represented by symmetry-invariant properties of the models. Thus TS commits us to a qualitative ontology, but it gives us no explicit, qualitative description of that ontology. Having said that, for this short example, I need no such details.

whose geometric structure is invariant under isometry, and that a sufficient condition for two models representing the same physical possibility is 'being isometric'. But given two Lorentzian metrics, how are we to ascertain *when* they are isometric?

To show that two given models are *not* isomorphic, all that TS implies is that if we scan the entire infinite-dimensional space of gauge transformations or diffeomorphisms, and do not find a gauge transformation (respectively, diffeomorphism) between the models, then they are not isomorphic. To be more explicit about this problem: given two coordinate-based descriptions of metrics $g^1_{\mu\nu}, g^2_{\mu\nu}$, we would either have to stumble upon a coordinate change that relates them (and so infer that they represent the same physical possibility), or else stumble upon some coordinate-invariant function that can be defined on all spacetimes and that takes different values on the two metrics (and so infer that they represent different physical possibilities). And thus if we fail on both counts, we could not decisively conclude that the two metrics correspond to same geometry, nor that they correspond to different geometries.

In the following, I will focus on giving tractable and sufficient conditions of identity, or a tractable notion of individuation for isomorphism-classes, by providing invariant descriptions of the isomorphism-classes. But I do not aim to give unique such descriptions. Quite the opposite: I claim there are many equally valid choices, differing perhaps with respect to theoretical virtues and depending on the context of application. Each such choice provides an explicit, reduced formalism for the theory, that can be broadly classified as 'relational' in spirit.

In our search for such invariant descriptions, we face two main obstacles. The first is the obstacle of *completeness*: the description may only be partial; the conditions that it provides are insufficient to ascertain whether any two isomorphism-classes are identical. One could try to overcome this obstacle by taking the union of many such partial descriptions, but then one faces the second obstacle, of *consistency*. Partial invariant descriptions may be inconsistent, and assessing consistency may again be an intractable endeavour. I will now illustrate these obstacles in the case of Riemannian geometries.

As our first attempt to find tractable and sufficient conditions of individuation, take a familiar list of isomorphism-invariant (in this case, isometry-invariant) properties. For instance: (1) "the manifold contains a two-dimensional surface, of area X, bounded only by two geodesics, of length, L"; or (2) "the manifold is geodesically convex"; or (3) "for all points x, there exists a unique point y whose geodesic distance πR from x is greater than all other points", etc. All of these properties are invariant under isometries. Property (3) is instantiated say, in the two-sphere, where it describes anti-podal points, and it is not instantiated on a plane. Indeed, it is only property (3) that comes close to fully characterising the geometry: it would characterise a metric sphere of some dimension. And that is only because the metric sphere is a highly unusual, highly symmetric model. In contrast, (1) and (2) are far from uniquely characterising a physical possibility, or even points or regions of the world in which these properties hold. To do so in general we would need an infinite list of properties like (1) and (2).

These examples illustrate the first obstacle that gets in the way of a description via a denumerable list of geometric invariants: *completeness*. A *complete* set of invariants (also called in the literature on gauge theory a complete set of 'observables', cf. (Henneaux & Teitelboim, 1992, Ch. 1.2), (Bergmann, 1961)) is a set whose elements, taken together, distinguish two isomorphism classes of models; i.e. any two models in two such classes disagree about the value of some or other observable in the set.

But clearly, any *finite*, consistent list of geometric invariants is, generically, *incomplete*. That is because a Riemannian geometry has infinitely many degrees of freedom. Indeed, even restricting the space of Lorentz metrics to those that satisfy the Einstein equations still leaves us with continuously many degrees of freedom at our disposal to fix the model, i.e. two physical degrees of freedom per space point. And the space of Yang-Mills connections also has continuously many degrees of freedom (cf. (Henneaux & Teitelboim, 1992, p. 29)). This means that, generically, no finite or even denumerable list will completely fix a particular isomorphism-class of a field theory.[8]

This leads us to the second obstacle, of consistency: it is not clear from inspection when two items in a list such as the one above are inconsistent. For instance, in the simple case where (3) is part of the list, the L of item (1) would have to be smaller than $4\pi R^2$. We arrive at this explicit inequality because in this simple case (3) fixes the geometry to a large degree, and we can then use familiar theorems of spherical geometry. But in general cases the consistency of a long list would pose an intractable problem: I call this the problem of *consistency* of generic lists of symmetry-invariants; or, more generally, the problem of *consistent descriptions*. We will encounter it again in Section 3.4, when we discuss non-locality. In sum, generically, any such description by a list of invariants can be subject to the problems of incompleteness and inconsistency.

So much by way of stating Limitation (A). To sum up: though TS says only symmetry-invariant properties of the models can represent physical properties, it neither provides nor 'sees the need' to provide a symmetry-invariant description of those models. Thus, in practice, TS does not help us ascertain when two models of our theory are isomorphic; it gives no guidance about how to find a consistent and complete set of observables, i.e. a set of observables whose compossible values label, one-to-one, the isomorphism classes.

2.1.b Limitation (B): the individuation of objects

I now turn to Limitation (B). Let us start by seeing it as an analogue, for objects or other features, such as regions, *within* a model, of Limitation (A) which applied to *models as a whole.* That is: if we are confronted with two models (whether isomorphic or not), and consider some intra-model object or feature in one model—such as a spacetime point or region in general relativity, or a frame for a vector space in gauge theory—then the general question arises: which (if any) object or feature of the other model corresponds to it?

[8] And then there is the related worry, about a non-denumerable list of invariants being vastly overcomplete! In this case, the elements of the list could be invariant but, since vastly overcomplete, still have to satisfy constraints that are in practice intractable. So this problem is one of consistency, as above, rather than one of completeness.

Agreed: stated so generally, this question is vague: it of course depends on what the scientific or philosophical meaning or role of 'correspond' is meant to be. In narrowing down this meaning, we must be careful not to make 'correspondence' too constraining: it could require too much similarity to be of any use. For instance, in the case of general relativity, suppose we take the notion of 'correspondence' to require the matching of all observables attached to a region. In such a case, two regions would correspond iff all their physical properties matched. In this case, no interesting comparisons can be made.

Again, TS gives no guidance about what properties or relations two spacetime points (or regions) in two models need to have in common in order to 'correspond' in this light-handed, flexible way. In particular, it gives no guidance as to which observables the points (regions) must match in their values, in order that they 'correspond'.

This is Limitation (B), and this paper will describe how we can overcome it. As with overcoming Limitation A, it will become apparent that there are many different, independent sets of relations that can 'uniquely characterise, or specify points as places in a structure' across possibilities. As I will argue in Section 4, these sets differ only with respect to theoretical virtues specific to the context of application.[9]

2.2 An upshot for theoretical physics

In this Section I will present evidence that the implications of these Limitations are not strictly metaphysical: they are important for current issues in theoretical physics.

I will start with Limitation (A), and with a connection to the hole argument (cf. (Gomes & Butterfield, 2023b; Pooley, 2023) for recent appraisals). The argument articulates the threat of a pernicious form of indeterminism that could arise in general relativity due to the the existence of isomorphic models. By stipulating that isomorphic models represent the same physical situation, TS disarms the threat.

Practicing physicists, almost unanimously (implicitly or explicitly) endorse TS in the case of general relativity. And yet, as most extensively catalogued by Belot & Earman (1999, 2001), they still consider certain issues related to the hole argument as relevant for understanding the theory:

> Far from dismissing the hole argument as a simple-minded mistake which is irrelevant to understanding general relativity, many physicists see it as providing crucial insight into the physical content of general relativity. (Belot & Earman, 1999, p. 169)

[9] Incidentally, this plurality of sets of relations that are sufficient for individuation is essentially absent from the allusions to individuation in the literature on symmetry and equivalence. On the contrary, this literature much more often (if not always) considers only the totality of properties and relations in which points stand with respect to each other. But, as in the analogy to Leibniz's PII, such totality gives us no tractable handle on these questions. Here I will propose a more tractable notion of relationism, that further picks out model-dependent subsets of relations.

But in truth, even these physicists think the hole argument is relevant to the understanding of general relativity only insofar as it is related to the question of observables, which in Section 2.1.a I tied to Limitation A.[10] As described by Isham:

> The diffeomorphism group moves points around. Invariance under such an active group of transformations robs the individual points of M of any fundamental ontological significance. [...] This is one aspect of the Einstein 'hole' argument that has featured in several recent expositions. [...] It is closely related to the question of what constitutes an observable in general relativity—a surprisingly contentious issue that has generated much debate over the years. (Isham, 1992, p. 170)

As to Limitation (B), overcoming it may be necessary in order to discuss local spacetime counterfactuals in full generality, as I briefly described in Section 2.1.b. Modal intuitions—'such a region of spacetime *could have been* more curved'—invoke *some* relation between spacetime regions across physical possibilities. Since we took TS to deny that spacetime points have primitive identity across physical possibilities, we need an alternative way to establish *local* inter-world comparisons. This obstacle is widely recognised by physicists; as remarked by (Penrose, 1996, p. 591):

> The basic principles of general relativity—as encompassed in the term 'the principle of general covariance' (and also 'principle of equivalence') —tell us that there is no natural way to identify the points of one space-time with corresponding spacetime points of another.

In the context of spacetime, Limitation (B) has led to some puzzlement about how general relativity meshes with standard conceptions of scientific laws and scientific understanding. For instance, according to (Hall, 2015, p. 270), 'the ability to provide sharp and determinate truth conditions for a wide range of counterfactuals is precisely what lends a good physical theory its explanatory power'. And here is (Curiel, 2015, p.1):[11]

> General relativity poses serious problems for counterfactual propositions peculiar to it as a physical theory. [...] Given the role of counterfactuals in the characterization of, inter alia, many accounts of scientific laws, theory confirmation and causation, general relativity once again presents us with idiosyncratic puzzles any attempt to analyze and understand the nature of scientific knowledge must face.

[10] As I described in Section 2.1.a, observables are quantities that are invariant under the symmetries of the theory, that we could use to describe physical content without redundancy). See e.g. (Donnelly & Giddings, 2016; Rovelli, 2007; Thiemann, 2003) for general arguments surrounding this difficult issue of 'gravitational observables' and as evidence of physicists' concern with the topic. There are also many examples from the string theory literature as well: e.g. see (Harlow & Wu, 2021) for an explicit, and rather complicated basis of observables in the simplified context of a two-dimensional gravitational theory called Jackiw-Teitelboim gravity; and (Witten, 2023).

[11] See (Jaramillo & Lam, 2021), for a more recent echo of this puzzlement, in the context of the initial value formulation of general relativity.

I do not aim here to judge the validity of these arguments about counterfactuals in general relativity. And although I will focus on counterpart relation for objects, I believe there are viable ways to talk about counterfactuals for spacetimes that don't require any such specific notion. With some effort, some counterfactuals can be articulated globally, using only observables. *Pace* this caveat, I maintain that finding counterpart relations for objects is the most direct and complete manner of articulating counterfactuals.

Indeed, the topic of spacetime counterparts has become even more timely than that of finding invariant observables, since it is germane to current research on the superpositions of gravitational fields (cf. (Kabel et al., 2024) for a recent appraisal). Agreed, 'superpositions of spacetime' is a quantum theme, which I will steer away from here. But, as I will explain in Section 3, Limitations (A) and (B) only take center stage in theoretical physics in the context of quantization or in the treatment of subsystems. Strictly speaking, in the classical domain in the context of the whole Universe there is no need for unique representation: a single representation is good enough, and TS suffices.[12]

2.3 What are we looking for?

In Section 2.1.b, I reformulated Limitation (2) more precisely as stating two shortcomings of TS. To recap: Limitation (A) is that we have not specified how individual physical possibilities are to be tractably individuated. Limitation (B) is that we have not provided invariant descriptions of objects within each physical possibility, in a way that allows some notion of correspondence—or rather, 'counterpart relations'; cf. (Gomes & Butterfield, 2023a)—between objects across possibilities.

Focussing on Limitation (A) in Section 2.1, I argued that denumerable lists of geometric invariants at best incompletely characterise an isomorphism-class, or at worst characterise nothing at all. So what *would* succeed in overcoming Limitation (A), by providing consistent representations of isomorphism-classes with tractable conditions of identity? I propose that it is a:

Definition 1 (Representational scheme) *A representational scheme for a given type of (isomorphism-invariant) structure is a complete, isomorphism invariant description of the isomorphism-classes of the theory, such that two such descriptions are identical iff the two isomorphism-classes are identical. Thus, mathematically, such a description is an explicit map σ from the space of models of the theory, call it Φ, to some fixed value space V (the space of tuples, maybe infinite tuples, of values of a complete set of observables) such that, for two models $\phi, \phi' \in \Phi$,*

$$\sigma(\phi) = \sigma(\phi'), \qquad \textit{iff} \qquad \phi \sim \phi', \tag{2.1}$$

where the equivalence relation is given by isomorphism.

[12] This is why I partly agree with Fletcher (2020); Weatherall (2018)'s deflation of the hole argument. It is important to note that, though these arguments and counterarguments have been given within the context of general relativity, there are closely analogous ones for Yang-Mills theory.

Two remarks are in order about this definition. First, I said the map is 'explicit': I did this to avoid implicit definitions that are in practice often intractable.[13] Second, it should be clear that I intend the definition to encompass many different choices of equally valid reduced formalisms. Although there is an element of choice, they are mathematically and physically constrained: this is why I called them *schemes* (cf. footnote 5).

But what are good examples of representational schemes for field theories with non-trivial isomorphisms, the cases that this paper focuses on?

In field theories, representational schemes are most conveniently provided through choices of gauge fixing. In that case, we require the output of a representational scheme to be itself a model of the original theory. This is why, in Definition 1, I denoted representational schemes with the familiar notation for sections of a fibre bundle, viz σ; in later Sections this will be a central case. So we define:

Definition 2 (Representational scheme via gauge-fixing) *A representational scheme via gauge-fixing is a 1-1 map between the equivalence classes given by isomorphisms of the theory and a subset $\mathcal{F}_\sigma$ of models of the theory, $\sigma : \Phi \to \mathcal{F}_\sigma \subset \Phi$, such that the two models in the subset are identical iff their arguments are isomorphic; i.e. they obey* (2.1).

Although it can be applied more widely, this type of representational scheme has many advantages, particularly in the case of field theories.

The first is that, for descriptions given through representational schemes satisfying Definition 2 the problem of *consistency* of descriptions, as described in Section 2.1, is also tractable. This advantage of Definition 2 over Definition 1 will be explained in Section 3.4. The second advantage is tightly related to the first: it is only for representational schemes satisfying Definition 2 that we can straightforwardly overcome Limitation (B).

Lastly, I should mention that it is not only gauge-fixings that bear a relationship to my notion of representational schemes as originally stated in Definition 1. In the current literature, they also encompass 'relational' (or 'material', or even 'quantum') reference frames, as well as the notion of *dressing*. And while a construal via gauge-fixing or dressed quantities is most apt to overcome Limitation (A), a construal via relational reference frames is most apt to overcome Limitation (B). We will visit the relationship between these three construals in what follows (see especially Section 3.3).

In sum: I will show Limitation (A) is overcome when we have provided a representational scheme for our models that satisfies Definition 1. As to Limitation (B), I will show it it is partially overcome for representational schemes via gauge-fixing, i.e. satisfying Definition 2.

[13] The condition excludes, for example, very abstract maps such as the projection onto the equivalence classes: $\pi : \Phi \to \Phi/\sim$. It is not always clear that the definition makes mathematical sense, in the first place; but, even if it did, it would not be an explicit map and thus would fail to satisfy Definition 1.

3 Representational schemes

Following up on Section 2.2, in Section 3.1 I will briefly describe two contexts in which representational schemes take center stage in theoretical physics; in which, that is, physicists cannot avoid facing Limitations (A) and (B): quantization and the treatment of subsystems. In Section 3.2, I will introduce the mathematics of representational schemes via gauge-fixing. In Section 3.3 I will make the relationship between representational schemes via gauge-fixings, dressings, and relational reference frames explicit. In Section 3.4 I will discuss the nonlocality of representational schemes and confront the question of consistency of partial descriptions.

3.1 Representational schemes in physics: a brief history

Section 3.1.a will explain why representaional schemes are important in quantization, and introduce the idea that gauge-fixings, dressings, and relational reference frames may be different sides of the same coin. Section 3.1.b describes the second context in which representational schemes are important: the treatment of subsystems.

3.1.a Gauge-fixing and relational observables in quantization

In Section 2.2, I argued that theoretical physicists are also preocuppied with matters of symmetry and equivalence. More specifically, with overcoming Limitation (A) by finding complete sets of 'observables' that fully describe the theory and tractably individuate isomorphism-classes. But in the classical domain, there is less need for unique representation, or for conditions of individuation in general: any model representing a given physical situation is as good as any other. What does it matter if there are many ways to represent the same physical situation? One is enough.

But passing over to the quantum domain, it matters: one is enough, more than one is too many. In the context of quantization, we need to eliminate redundancy of representation, even for technical reasons. So physicists mostly talk about complete observables in the context of quantization, where, at least in certain approaches it is ineluctable. For instance, in the path integral, or sum-over-histories approach, the tools of perturbative field theory fail if we do not treat the symmetry-related histories as being one and the same, or physically identical.[14]

There are in general two preferred ways to treat symmetry-related histories as one and the same. One approach is to reconstrue the theory, e.g. the path integral, in terms of symmetry-invariant, relational variables: this is the 'purist's' approach; common among experts in general relativity and those working on quantum gravity. In contrast, particle physicists working on gauge theory will not spend much time discussing relational invariant observables. Within that community the usual approach to redundancy is to just "fix the gauge". To fix the gauge is

[14]In slightly more detail: the propagator involves the inverse of the Hessian around a classical solution; since the Hessian of the action vanishes along the generators of the symmetry, it diverges. Thus one would obtain divergences at every order of the perturbation series: these divergences cannot be neatly separated into energy scales and so they are immune to the machinery of effective field theories.

to satisfy Definition 2 by implementing auxiliary conditions that are to be satisfied by a single element in each class of symmetry-related models. Gauge fixing is very useful in practice and widely employed in the actual calculations of gauge theory and (classical and perturbative) gravity. But it is also often (erronously) assumed to do violence to the original symmetry of the theory.

The contrast is clear: in the less applied, more conceptually-driven domains, gauge-fixing is often overlooked in lieu of a choice of 'relational observables' and, more recently, in lieu of 'dressed fields', which are a particular type of relational observables.[15]

But I maintain that this assumption, this contrast sketched above, between gauge-fixing, relational observables (and reference frames) is in one way misleading. For as we will see in Section 3.3.b, gauge fixings can be understood in terms of symmetry-invariant composites of the original fields, called *dressed fields*, and therefore do no violence to the original symmetry of the theory. Dressed fileds give rise to a complete set of observables that can be understood relationally, i.e. as providing a description of the isomorphism class with respect to a physical or *relational reference frame*. On the other hand, some complete sets of 'observables' admit no construal in terms of gauge-fixing; which is why the contrast is not *completely* misleading. In Section 3.3, I will give a conceptually unified account of dressed quantities, material reference frames, gauge-fixings, and relational observables.

3.1.b A brief history of relational reference frames and subsystems

Apart from the more familiar context of quantisation, discussed in the previous Section, representational schemes in the guise of dressed quantities—which will be discussed at length in Section 3.3.b—sometimes understood as descriptions of a physical state relative to a relational reference frame, have recently been invoked in the treatment of subsystems in gauge theory and general relativity. Indeed, our treatment of representational schemes here is inspired by an application to subsystems, which came chronologically first, in (Gomes & Riello, 2017) (where they were described as 'abstract material reference frames', or 'relational connection-forms'). This initial general treatment was illustrated with explicit examples in (Gomes et al., 2019; Gomes & Riello, 2018, 2021), where a concept of 'abstract reference frames' was compared to dressings in gauge theory and gravity; in (Gomes, 2019, 2021a) I gave a conceptual treatment of these ideas (I'll add some detail in Section 6.2).

The reason these notions of dressing or relational reference frames were first introduced in the context of subsystems was that there are subtleties in construing the isomorphisms of subsystem states as dynamical symmetries; subtleties which can be treated by appealing to properties of representational schemes. In particular, there are subtleties about the symmetry-invariance of a bounded subsystem's dynamical structures, such as its intrinsic Hamiltonian, symplectic structure, and variational principles in general.[16]

[15] Bagan et al. (2000); Lavelle & McMullan (1995, 1997) were, to my knowledge, the first to explore the relationship between dressings and gauge-fixings, in the context of QCD.

[16] The point being that dynamical structures on bounded subsystems are *not* invariant under arbitrary gauge

The point, summarised in footnote 16, is that by describing a model relative to some subset of degrees of freedom, or from the point of view of 'relational reference frames', our description is rendered explicitly and fully gauge invariant, even in the presence of arbitrary boundaries.

Summing up Section 3.1: back in Section 2.1 I had justified Limitations (A) and (B) from a conceptual and mathematical perspective. Here I alerted the reader that the Limitations pertain also to physics; and indeed they must be resolved when dealing with quantum systems, subsystems, and counterfactuals.

3.2 Representational schemes via gauge-fixing: general properties

At the abstract level explored in this and the next Sections, there are no substantive differences between diffeomorphisms and the symmetries of gauge theories. Indeed, even symmetries of non-relativistic particle mechanics, which mostly lie outside of the scope of this paper, can still be encompassed by the formalism. And so we treat all of these symmetries uniformly, labelling them with the group $\mathcal{G}$, which could be infinite-dimensional.

Given the space of models of a theory, Φ, we assume it has a smooth manifold structure (infinite-dimensional, in the case of field theories) and it admits an action of $\mathcal{G}$, i.e. a map $\Phi : \mathcal{G} \times \Phi \to \Phi$, that preserves the global structure on Φ (e.g. is smooth, in the topology of Φ), and preserves dynamics, e.g. the Hamiltonian, the action functional, or the equations of motion (in the language of (Gomes, 2021c), they are S-symmetries). More formally: there is a structure-preserving map, μ, on Φ that can be characterized element-wise, for $g \in \mathcal{G}$ and $\varphi \in \Phi$, as follows:

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

transformations of the boundary state, so, in order to preserve gauge-invariance at the level of the subsystems, the boundary state must be understood as 'already gauge-fixed', even without an explicit gauge-fixing condition. This would in practice fix gauge transformations to the identity at the boundary of a subsystem. In effect, fixing the gauge transformations in this way, for finite boundaries, requires the stipulation of a single value for the fields at the boundary. Of course, this may be very well for a particular subsystem, but it is not a condition we would wish to enforce for *any* subsystem, especially if we are to take subsystems of subsystems. In contrast, a representational scheme does not impose any physical limitation on the boundary states, and so restrictions to subsystems is explicitly recursive here.

The modern guise of the obstruction to gauge-invariance posed by boundaries was first highlighted in (Donnelly & Freidel, 2016) (which also introduced 'edge-modes' as new degrees of freedom that fixed this obstruction), which caused a flurry of papers on subsystems in gauge theory (cf. e.g. (Carrozza & Höhn, 2022; Donnelly & Freidel, 2016; Geiller, 2017; Geiller & Jai-akson, 2020; Gomes, 2019; Gomes et al., 2019; Gomes & Riello, 2021; Ramirez & Teh, 2019; Riello, 2021) and references therein). (Gomes & Riello, 2017) introduced relational reference frames as a tool to restore gauge-invariance in a subsystem-recursive manner. It was there fist suggested that the resolution of the gauge-invariance problem did not require new degrees of freedom at the boundary, but physical reference fields with respect to which one should describe variations *throughout* the subsystems, and not just at the boundary. See (Carrozza & Höhn, 2022; Goeller et al., 2022) for a more recent and comprehensive review of the relationship between dressings, relational reference frames, and gauge-fixings.

The symmetry group partitions the space of models into equivalence classes in accordance with an equivalence relation, $\sim$, where $\varphi \sim \varphi'$ iff for some g, $\varphi' = \varphi^g$. We denote the orbit of φ under $\mathcal{G}$ by $\mathcal{O}_\varphi := \{\varphi^g, g \in \mathcal{G}\}$, which as a set is isomorphic to the equivalence class of φ under the equivalence relation, usually denoted by square brackets, $[\varphi]$.[17] Writing the canonical

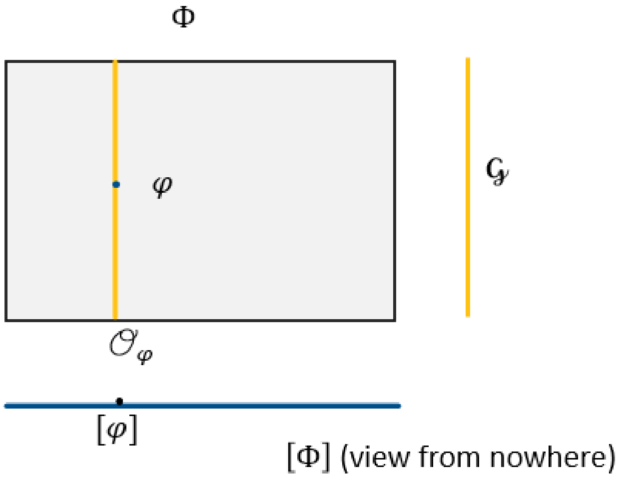

Figure 1: The space of states, 'foliated' by the action of some group $\mathcal{G}$ that preserves the value of some relevant quantity, S, and the space of equivalence classes. In field theory, S is the value of the action functional, and the spaces $\mathcal{G}$, Φ, and $[\Phi]$ are infinite-dimensional (Frechét) manifolds.

projection operator onto the equivalence classes, $\mathsf{pr} : \Phi \to \Phi/\mathcal{G} =: [\Phi]$, taking $\varphi \mapsto [\varphi]$, then the orbit $\mathcal{O}_\varphi$ is the pre-image of this projection, i.e. $\mathcal{O}_\varphi := \mathsf{pr}^{-1}([\varphi])$.

We can now improve on Definition 2 with the more precise:

Definition 3 (Representational scheme via gauge-fixing) *A representational scheme via gauge-fixing is an injective map*

$$\begin{aligned} \sigma : [\Phi] &\to \Phi \\ [\varphi] &\mapsto \sigma([\varphi]) \end{aligned} \tag{3.2}$$

that respects the required mathematical structures of Φ, e.g. smoothness or differentiability and is such that $\mathsf{pr}(\sigma([\varphi])) = [\varphi]$, where pr is the canonical projection map onto the equivalence classes.

Armed with such a choice of representative for each orbit, a generic model φ could be written uniquely as some doublet $\varphi = ([\varphi], g)_\sigma := \sigma([\varphi])^g$, which of course satisfies: $\varphi^{g'} =$

[17] I here see no important difference between $[\varphi]$ and $\mathcal{O}_\varphi$; the only difference is that the latter is usually seen as an embedded manifold of Φ, whereas the former exists only implicitly, or abstractly, outside of Φ (which is why I called it 'the view from nowhere'; see Figure 1).

$(\sigma([\varphi]))^{gg'} = ([\varphi], gg')_\sigma$.[18] Thus we identify $\Phi \simeq [\Phi] \times \mathcal{G}$ via the diffeomorphism:

$$\begin{aligned} \overline{\sigma} : [\Phi] \times \mathcal{G} &\to \Phi \\ ([\varphi], g) &\mapsto \sigma([\varphi])^g \end{aligned} \tag{3.3}$$

Now, as I mentioned, the space $[\Phi]$ is abstract, or only defined implicitly; this is why I called it "the view from nowhere". Since we cannot usually represent elements $[\varphi]$ of $[\Phi]$ intrinsically, we in practice replace σ by an equivalent projection operator that takes any element of a given orbit to the image of σ:

Definition 4 (Projection operator for σ) *A map*

$$\begin{aligned} h_\sigma : \Phi &\to \Phi \\ \varphi &\mapsto h_\sigma(\varphi) = \sigma([\varphi]), \end{aligned} \tag{3.4}$$

is called the projection operator for σ of Definition 3. It maps a model to an isomorphic copy chosen as a representation of its orbit by σ. The end result, $h_\sigma(\varphi)$, is called a dressed *φ.*

Since $[\varphi^g] = [\varphi]$, we must have

$$h_\sigma(\varphi^g) = h_\sigma(\varphi). \tag{3.5}$$

Now, $h_\sigma(\varphi)$ is an explicit map on Φ, i.e. it is a function of φ. So h_σ uniquely and concretely represents the invariant, or structural, content. In other words, two given models, φ, φ', give the same value for *all* symmetry-invariant quantities iff $h_\sigma(\varphi) = h_\sigma(\varphi')$. Thus the invariant structure of each isomorphism-class $[\varphi]$ is represented uniquely, according to σ, by the corresponding value of the map $h : \Phi \to \Phi$.

3.3 Gauge-fixing, dressing, and relational frames in field theories

The three Sections below will respectively focus on gauge-fixing, dressings, and relational frames. The common thread in this Section, that I will argue runs through each of these three topics, is the notion of representational scheme.

3.3.a Gauge-fixing

In equation (3.3), describing the product structure of the space of models, I oversimplified: finite-dimensional principal bundles need not be trivial, and, accordingly, the space of models need not be globally isomorphic to the product between the space of physical states and the group of gauge transformations, $[\Phi] \times \mathcal{G}$. Indeed, even a local product form for the space of models is only guaranteed to exist in the finite-dimensional case.

Nonetheless, it is in fact true that the space of models Φ—for both Riemannian metrics and gauge connection-forms—is mathematically very similar to a principal fiber bundle, with

[18] Our notation is slightly different than Wallace (2019, p. 9)'s, who denotes these doublets as (O, g) (in our notation $([\varphi], g)$), and labels the choice of representative (or gauge-fixing) as φ_O (our φ_σ). We prefer the latter notation, since it makes it clear that there is a choice to be made.

$\mathcal{G}$ as its structure group (resp. diffeomorphisms and vertical automorphisms). But there are important differences between the infinite-dimensional and the finite-dimensional case. In the finite-dimensional case, a free (and proper) action of the group on the manifold suffices for that manifold to have a principal G-bundle structure, usually written as $G \hookrightarrow P \to P/G$. In the infinite-dimensional case, these properties of the group action are not enough to guarantee the necessary fibered, or local product structure: one has to construct that structure by first defining a *section*. I summarise the construction of that structure and its main obstructions in Appendix B.

A choice of section is essentially a choice of embedded submanifold on the model space Φ that intersects each orbit exactly once. In the 2-dimensional figure 1 above, this would be represented by a 'nowhere-vertical' curve: i.e. a submanifold that is transversal to all the $\mathcal{G}$-orbits and intersects each orbit once.

A convenient way to define such sections is to impose further functional equations that the model in the aimed-for representation must satisfy; this is like defining a submanifold through the regular value theorem: in finite dimensions, defining a co-dimension k surface $\Sigma \subset N$ for some n-dimensional manifold N, with $n > k$, as $\mathcal{F}_\sigma^{-1}(c)$, for $c \in \mathbb{R}^k$, and $\mathcal{F}_\sigma$ a smooth and regular function, i.e. $\mathcal{F}_\sigma : N \to \mathbb{R}^k$ such that $\mathsf{ker}(T\mathcal{F}_\sigma) = 0$.

Once the surface is defined, σ can be seen as the embedding map with range $\sigma([\Phi]) = \mathcal{F}_\sigma^{-1}(0) \subset \Phi$. The next step is to find a gauge-invariant projection map, h_σ, that projects any configuration to this surface.

3.3.b Dressing

I have called $h_\sigma(\varphi)$ the dressed variables. What are its clothes? In the case of dressings associated to gauge-fixings, they are the gauge transformation $g_\sigma(\varphi)$ required to transform φ to a configuration $\varphi^{g_\sigma(\varphi)}$ which belongs to the gauge-fixing section σ.[19]

Here we encounter the dual interpretations of gauge-fixings and dressed variables. In the case of fields, the projection g_σ is a function of φ (that is usually non-local and 'relational', in the sense that it stands for a comparison of the values of different fields at different points). Conceptually, the image of this projection—the dressed field—is a symmetry-invariant 'composite' of the original degrees of freedom which uniquely describes each isomorphism-class relationally.

In more detail: a function $\mathcal{F}_\sigma : \Phi \to W$, valued on some general vector space W, defining the section (surface) $\sigma([\Phi])$ as its level surfaces $\mathcal{F}_\sigma^{-1}(0)$ is suitable as a gauge-fixing iff it gives rise to a unique dressing by satisfying two conditions:[20]

[19]This is not the most general type of dressing function. Firstly, there is the infinitesimal version, given by the relational connection-form ϖ (cf. (Gomes & Riello, 2017, 2018, 2021) and (Gomes et al., 2019); and (Gomes, 2019, 2021a) and (Gomes & Butterfield, 2023a, Appendix A) for philosophical introductions). I will briefly summarise this construction in Section 6.2. Second, there is a more complete theory of dressings and dressed fields, which only requires the right covariance properties and can restrict to subgroups of the gauge group: see (François, 2021; François et al., 2021) for a review.

[20]For these conditions to be satisfied, and for their satisfaction to be guaranteed by the regular value theorem,

• *Universality* (or existence): For all $\varphi \in \Phi$, the equation $\mathcal{F}_\sigma(\varphi^g) = 0$ must be solvable by a dressing $g_\sigma : \Phi \to \mathcal{G}$.

That is, there is a Φ-structure-preserving map,

$$g_\sigma : \Phi \to \mathcal{G}, \text{ such that } \mathcal{F}_\sigma(\varphi^{g_\sigma(\varphi)}) = 0, \text{ for all } \varphi \in \Phi. \tag{3.6}$$

This condition ensures that $\mathcal{F}_\sigma$ doesn't impose *physical* constraints on what can be represented, i.e. that each orbit possesses at least one intersection with the gauge-fixing section.

• *Uniqueness*: g_σ is the unique functional solution of $\mathcal{F}_\sigma(\varphi^{g_\sigma(\varphi)}) = 0$. This condition ensures that each orbit is represented uniquely in the section.

From these two conditions, it follows that $\varphi^{g_\sigma(\varphi)} = \varphi'^{g_\sigma(\varphi')}$ if and only if $\varphi \sim \varphi'$, meaning that models have the same projection onto the section iff they are isomorphic. In one direction, the claim is immediate: $\varphi^{g_\sigma(\varphi)} = \varphi'^{g_\sigma(\varphi')}$ implies $\varphi = \varphi'^{g_\sigma(\varphi')^{-1}g_\sigma(\varphi)}$, so $\varphi \sim \varphi'$. In the other direction, assume $\phi' = \phi^g$ and g_σ is unique. From (3.6), since $[\varphi^g] = [\varphi]$ and $\sigma : [\Phi] \to \Phi$ is a unique embedding, then:

$$\varphi^{g_\sigma(\varphi)} = \sigma([\varphi]) = \sigma([\varphi^g]) = (\varphi^g)^{g_\sigma(\varphi^g)} = \varphi'^{g_\sigma(\varphi')}, \tag{3.7}$$

which establishes the other direction, that two models have the same projection only if they are gauge related. And we obtain from (3.7) the following equivariance property for g_σ:[21]

$$g_\sigma(\varphi^g) = g^{-1}g_\sigma(\varphi). \tag{3.8}$$

In possession of a dressing function g_σ, it is convenient to rewrite the dressed field h_σ of (3.4) as $h_\sigma(\bullet) := \bullet^{g_\sigma(\bullet)}$, i.e:

$$\begin{aligned} h_\sigma : \Phi &\to \Phi \\ \varphi &\mapsto h_\sigma(\varphi) := \varphi^{g_\sigma(\varphi)} \end{aligned} \tag{3.9}$$

And of course, we can still change the representational scheme itself, i.e. act with the group on the image of h_σ: because $h : \Phi \to \Phi$ is a projection (as opposed to a reduction $\mathsf{pr} : \Phi \to [\Phi]$), and thus $h_\sigma(\varphi) \in \Phi$, this gives $(h_\sigma(\varphi))^g$. But $(h_\sigma(\varphi))^g$ no longer obeys the conditions defining the previous representational scheme, $\mathcal{F}_\sigma(\varphi^{g_\sigma(\varphi)}) = 0$. If the original representational scheme was chosen to fulfill certain (extra-empirical) desiderata, these would not longer apply to $(h_\sigma(\varphi))^g$. Of course, the new representation might fulfill a new set of desiderata.

Given any two representational schemes σ, σ', we have

$$\varphi^{g_\sigma(\varphi)} = (\varphi^{g'_\sigma(\varphi)})^{g'_\sigma(\varphi)^{-1}g_\sigma(\varphi)} \tag{3.10}$$

W is a space that is locally isomorphic to the Lie group in question: i.e. it is isomorphic to the Lie algebra of the symmetry. This guarantees W has the right dimensions for regular values to form a section of the manifold, i.e. a submanifold whose complement has the dimension of the group.

[21] These equations are schematically identical to the ones we find for Yang-Mills theory in the finite-dimensional case. In the Yang-Mills case, a section σ in the space of models is essentially a section σ of the finite-dimensional principal bundle, P that depends on the isomorphism-class.

and so we obtain

$$h_\sigma(\varphi) = (h_{\sigma'}(\varphi))^{\mathfrak{t}_{\sigma\sigma'}(\varphi)}, \tag{3.11}$$

where $\mathfrak{t}_{\sigma\sigma'}(\varphi) := g'_\sigma(\varphi)^{-1} g_\sigma(\varphi)$ is a model-dependent isomorphism (the analog of the transition map between sections of a principal bundle). Note that, from the covariance of g_σ, given in (3.8), it immediately follows that the transition map is itself gauge invariant, as expected. That is, the transition function depends only on the isomorphism-classes, or orbits $\mathcal{O}_\varphi$. Note also that the dressings of each scheme according to each other are related by an inverse:

$$g_{\sigma'}(\varphi^{g_\sigma(\varphi)}) = g_\sigma(\varphi)^{-1} g_{\sigma'}(\varphi) = (g_\sigma(\varphi^{g'_\sigma(\varphi)}))^{-1} \tag{3.12}$$

3.3.c Relational reference frames

Finally, as mentioned in Section 3.1.b, dressed variables have also been associated to 'material reference frames', where 'material' has nothing to do with 'matter'; it is meant as 'physical', or 'relational' (I'll stick to relational.) Relationism of course has a noble tradition, and the broad idea of a 'relational reference frame' is a mainstay of research in quantum gravity (cf. e.g. (Rovelli, 2007)). The idea is that we can anchor a particular representation to a physical system, with the ensuing representation being straightforwardly understood in terms of relations to this physical system.

For instance, given a set of four scalar fields obeying functionally independent Klein-Gordon equations, we can understand DeDonder gauge as using the density of these fields as coordinates on a region of spacetime (see Appendix A.3). In the case of spacetime, by locating points relative to this physical system, we make precise a frequent claim found in the literature on TS about spacetime; namely that 'spacetime points can only be specified by their web of relations to other points.'

Similarly, in electromagnetism coupled to a complex scalar field that is nowhere vanishing, the phase of the complex field selects a particular internal frame for each complex plane, and we would describe other fields relative to this frame. Unitary gauge can be understood in this way (see Appendix A.2.b). In each of these cases, g_σ can be understood as the transformation required to go from an arbitrary representation of the isomorphism-classes to a preferred representation relative to this 'relational reference frame'.

But I must issue a warning that does not usually accompany expositions of the relationship between gauge-fixing and reference frames. The warning, germane to Limitation (B), is that, when locating spacetime points or directions in an internal vector space by their physical content, we must take care not to overdo it. Suppose, for example, you define a spacetime region as one who admits a certain value of the metric tensor. Clearly, we will not be able to express any other metric tensor on this region. Heuristically, the web of relations used to identify locations cannot be too rigid, otherwise each location will bear their content essentially. In that case, no interesting counterfactuals can be expressed. In the context of the hole argument, this kind of rigidity leads to (Maudlin, 1988, 1990)'s 'metric essentialist' resolution, which indeed has been criticised for being limited to the 'actual world' (Butterfield, 1989); and in (Gomes & Butterfield, 2023a) precisely for not allowing the expression of counterfactuals.

The use of relational reference frames to solve Limitation (B), about counterfactuals, is thus contingent on their association with a gauge-fixing, and in particular, its property of Universality. It is this property of gauge-fixings that enable their association to reference frames: we must employ just enough relations between parts of the fields so as to fix the representation without limiting the physical content.[22] I have detailed specific examples of this association between reference frames and gauge-fixings in Appendix A (in particular, in Sections A.2.b and A.3).

We are now ready to deploy gauge fixings to obtain a counterpart theory that can be used to overcome Limitation (B). This will be done in Section 5. Before we get to that, we must still iron out two conceptual wrinkles: about the frequent non-locality of gauge fixings (Section 3.4) and about how to choose representational conventions (Section 4).

3.4 Consistency of partial descriptions and gluing: the advantage of gauge-fixing

The kind of non-locality that emerges via dressings in gauge theory forces us to again face the question of *consistency* of partial descriptions of an isomorphism-class.[23] This question was introduced in Section 2.1, where I argued that descriptions via a denumerable list of invariants was generally inconsistent. And here I will show that, for representational schemes via gauge-fixing, it has a neat answer.

How does the non-locality due to dressings differ from the type of holism illustrated by denumerable lists of geometric invariants, as discussed in Section 2.1? Those lists will often lead to an inconsistent description, but here I will show that, in the case of dressings obtained via gauge-fixing of field theories, the consistency of piecewise-invariant descriptions is straightforward.

Briefly, suppose two regions, R_+ and R_- (or $R_\pm$ for short) overlap on a region $R_0 := R_+ \cap R_-$. Further suppose that the union of all three regions is a manifold, and that each region is an embedded submanifold within this union. Consider the fields intrinsic to each region, and their intrinsic isomorphisms; call these $\Phi_\pm, \Phi_0$ and $\mathcal{G}_\pm, \mathcal{G}_0$, respectively. For given gauge-fixings $\sigma_\pm$ of $\mathcal{G}_\pm$ on $R_\pm$, we have $h^+(\varphi^+) \in \Phi_+$, and $h^-(\varphi^-) \in \Phi_-$; which we again call $h^\pm$ for short. Then the $h^\pm$ are consistent iff given a representational scheme σ_0 (not necessarily via gauge-fixing) on R_0,

$$h^0(h^+_{|R_0}) = h^0(h^-_{|R_0}), \tag{3.13}$$

where $h^\pm_{|R_0}$ represents the restriction of the $h^\pm$ to the region R_0, and so $h^\pm_{|R_0} \in \Phi_0$. Equation (3.13) means that the $h^\pm_{|R_0}$ are related by a (isomorphism-class-dependent) transformation $\mathfrak{t}_\pm$, as described in (3.11). Thus $h^\pm$ are consistent, or can be 'glued', or composed into a single global model, iff (3.13) holds, for any σ_o.

[22] Indeed, it is for this reason that relational reference frames associated to gauge-fixings are able to restore a gauge-invariant notion of recursivity for arbitrary subsystems: see footnote 16.

[23] In Appendix A, I will give examples of dressings obtained via representational schemes, most of which are gauge-fixings. And most, whether obtained via a gauge-fixing or not, will display some degree of non-locality. This feature is expected for generic states of gauge theory and general relativity.

For a local dressing, $h(\varphi)(x) = h(\varphi(x))(x)$: the dressing at x depends only on the values of the fields being dressed at x. So, for the same representational scheme chosen throughout the manifold:

$$h^+(\varphi_{|R_+}) = h(\varphi)_{|R_+}, \quad \text{thus} \quad h^+_{|R_0} = h_{|R_0} = h^-_{|R_0}. \tag{3.14}$$

So, although (3.13) applies to both local and non-local dressings, it is only non-trivial in the non-local case. In the non-local case, even if the choice of gauge-fixing on R_+ was the same as that in R_- (e.g. Coulomb gauge), the $h^\pm$, when restricted to R_0, may differ.[24] And yet, even if $h^\pm_{|R_0}$ differ, they may still be different representations of the same invariant structure for that region. This is why, in the non-local case, it is important that the dressed quantities $h^\pm$ are also models of the theory. For, in order to assess consistency of the piecewise dressed quantities, these quantities must enter as arguments of the representational scheme h_0 in (3.13). When σ_0 is also a representational scheme via gauge-fixing, the consistency condition given in (3.13) is checked by composing the solution of two differential equations. The 'gluing' and composition of gauge-fixed quantities is described in detail for Yang-Mills theory in (Gomes & Riello, 2021, Sec. 6).

4 How to choose a representational scheme?

There are infinitely many possible choices of representational schemes, even if we just count those via gauge-fixing. How do we choose among them? In his book on the philosophical interpretation of gauge theories, Healey (2007) argues that here there is an important difference between gauge theory and general relativity. The argument occurs in his Chapter on classical gauge theories (Chapter 4) and it says that we can use features of the world to orient choices of representation for one theory but not for the other. Though Healey's alleged difference is non-existent, his claim brings to the fore a good and relevant question: how can we make choices between representations that equally well represent a given situation? By answering this question, I will (superficially) connect representational schemes to a popular topic in analytic philosophy: functionalism.

I will start in Section 4.1 by recapitulating Healey's argument. Then in Section 4.2 I

[24] The non-locality of dressings is equivalent to the non-commutativity between regional restrictions and the dressing map. That is, suppose we are given a representational scheme σ, which is applicable to any region, e.g. by stipulating a gauge-fixing and boundary condition. Suppose further that the dressing function requires solving an elliptic equation (as in e.g. Coulomb gauge, cf. Section A.2.a). Then for a generic submanifold R, with boundary ∂R:

$$h(\varphi)_{|R} \neq h(\varphi_{|R}). \tag{3.15}$$

This is easy to see in terms of boundary conditions: although $h(\varphi_{|R})$ will necessarily obey the stipulated boundary condition on ∂R, the values of $h(\varphi)_{|\partial R}$ will be rather arbitrary, and highly dependent on which surface we choose. And the difference won't be constrained to ∂R of course, since solutions of elliptic equations are globally dependent on the boundary condition. This is discussed at lengh, including complications about whether the region-intrinsic isomorphisms are also dynamical symmetries of the theory, in (Gomes, 2022, Ch. 6).

will show that, *contra* Healey, intra-theoretic resources enable us to pick out representational schemes for gauge theory.

4.1 Healey's argument from functional roles

(Healey, 2007, Section 4.2)'s describes what he takes to be a fundamental difference between the tenability of gauge-fixing in general relativity and in Yang-Mills theory. That is, a difference between specifying, amongst the infinitely many physically equivalent representatives, a particular spacetime distribution of the gauge potentials or of the metric. As I understand Healey, he argues that this specification is easy for the metric, but impossible for the gauge potential. It is easy for the metric because Lewis's ideas about functionalism apply to it, but, allegedly, they don't apply to gauge potentials. And with this alleged contrast I will disagree. My disagreement will shed light on how representational schemes are chosen.

To spell out Healey's argument in more detail, I will now indulge in a bit of 'Healey exegesis'. Healey admits that within a theory there may be many terms standing for unobservable items, such as those variables that are not gauge-invariant. But unobservability by itself is not bad news, since Lewis (1970, 1972)'s construction, which I will summarize shortly below, extends over any putative observable-unobservable gap, and indeed, over any bipartite distinction of theoretical vocabulary, usually taken to be between accepted, or better-undertood terms—labeled O-terms—and those that are 'troublesome', or less understood—labeleled T-terms. The construction allows us to specify the meaning of T-terms via their relations to O-terms and to each other. In more detail, the idea is that we can simultaneously specify what several theoretical items refer to, by their each uniquely satisfying some description (usually called "functional role") that can be formulated in terms of each other and of the better understood—perhaps even observable—terms.

But how could we use a theory to functionally single out any item if that theory treats that item and others, at least in certain respects, as being on a par, i.e. related by some symmetry? As described by Lewis (2009), the response is to appeal to patterns of facts of "geography" to break the underdetermination:

> Should we worry about symmetries, for instance the symmetry between positive and negative charge? No: even if positive and negative charge were exactly alike in their nomological roles, it would still be true that negative charge is found in the outlying parts of atoms hereabouts, and positive charge is found in the central parts. O-language has the resources to say so, and we may assume that the postulate mentions whatever it takes *to break* such symmetries. Thus the theoretical roles of positive and negative charge are not purely nomological roles; *they are locational roles as well.* [my italic]

But Healey argues that, in gauge theories, even this Lewisian strategy is bound to be plagued by under-determination. I interpret Healey as saying that one can functionally specify a representational scheme for a spacetime metric using only O-terms, but cannot similarly specify a representational scheme for a gauge potential. He says:

> The idea seems to be to secure unique realization of the terms in face of the assumed symmetry of the fundamental theory T in which they figure by adding one or more sentences stating what might be thought of as "initial conditions" to the laws of that theory [...] to break the symmetry of how these terms figure in T. They would do this *by applying further constraints* [...] Those constraints would then fix the actual denotation of the [...symmetry-related terms...] in T so that, subject to these further constraints, T is uniquely realized. [...But] *The gauge symmetry of the theory would prevent us from being able to say or otherwise specify which among an infinity of distinct distributions so represented or described is realized in that situation. This is of course, not the case for general relativity.* (Healey, 2007, p. 93) [my italics]

But I would ask: why does Healey see here a contrast between general relativity and gauge theory? I do not see in this entire passage an attempt to draw a distinction involving the applicability of TS to the two cases. So let us assume that TS applies to both, and so that there is no physical difference represented by isomorphic distributions of either the potential or the metric. This gives an uninteresting answer to an uninteresting question.

But there is a more interesting question being alluded to here. Namely, whether we can use features of the world around us to single out a unique model of the theory in both general relativity and gauge theory. Thus I take the more interesting interpretation of Healey's passage here to be that this 'singling out' of particular models is possible for gravity, but not for gauge theory, where the choice is unthethered from any feature of the physical world: it is entirely arbitrary, according to him.

This question is generally interesting because, in practice, we *do* select some model over others when we represent a given physical situation, and therefore we must somehow 'break the symmetry' between all of the isomorphic models. In other words, we are faced with a tension between: (a) taking all physical properties as invariant under the symmetries in question, and (b) in practice selecting one among the infinitely many isomorphic representatives of the same situation. At first sight, these two requirements, (a) and (b), are inimical, if not contradictory, for (a) implies we can have no physical guidance for accomplishing (b)! Getting representational schemes off the ground of course depends on a positive resolution of this paradox.

4.2 Refuting Healey's alleged distinction using representational schemes

Contra Healey, I say that having some physical "hook" with which to choose representatives does *not* imply that something in the world breaks the symmetry. Different choices of representational schemes *are* equally capable of representing a given state of affairs. But some representations may be more cumbersome, or they may obscure matters for the purposes at hand. While other representations may shine a light on a particular question we might have about the system. By choosing different representation schemes we shift our focus to different features of the world, according to our interest.

So we construct a particular representative of the gauge potential as fulfilling a given non-empirical role. We dissolve the tension between (a) and (b) described above by, in Healey's words: '*breaking* the symmetries,' by providing '*further constraints*'. These further constraints are non-empirical, but they force the representative to clearly exhibit a theoretical property of interest.

To give explicit examples, I note, to begin with that, according to Healey's standards, we are justified in including in our O-vocabulary all the 'locational roles', which describe contingent, happenstantial facts about 'where and when' specified events happen; and which I will loosely interpret as 'referring to spacetime'. Thus I free myself to include in Lewis' O-vocabulary the differential geometry of spacetime. In short, I will assume, in the case of gauge theory, that reference to spacetime is 'old' or already understood.

Here are three explicit examples of how to relate a theoretical interest to a choice of gauge, using only electromagnetism coupled to scalar fields. In our hierarchy of extra-empirical theoretical virtues, we could place Lorentz covariance very highly (cf. (Mulder, 2021) and (Mattingly, 2006) for advocacy of this criterion and choice of gauge) and so opt for Lorenz gauge, which maintains that property. Or we might want to count physical degrees of freedom of the electromagnetic field, in which case it would be better to impose Coulomb gauge (cf. Appendix A.2.a and (Gomes & Butterfield, 2022) for more about the properties of this choice). And there are more abstract theoretical roles we might want to focus on: e.g. the unitarity of the S-matrix in each order of perturbation theory is explicit in unitary gauge (cf. Appendix A.2.b and (Wallace, 2024) for more about the properties of this gauge). In the first two cases, the only extra constraints are that the spacetime (or, resp. spatial) divergence of the gauge potential vanishes. So both use only the O-vocabulary that Healey would grant us. The third case asks only that the internal phase of the scalar field be unity everywhere

Of course, it is still true that no empirical fact can pick out a unique choice of σ. But a different choice would not satisfy the original condition that uniquely specified σ. Thus different choices merely represent different lenses through which we capture the invariant structure of the states.

We should not be disappointed by this lack of uniqueness. It should be seen instead as a *flexibility* that is also explanatory. Just as it is easy to explain the Larmor effect by a Lorentz boost between different frames, we use different gauges to explain that a given process in quantum electrodynamics involves two physical polarisation states and that it is Lorentz invariant. In both scenarios two different 'frames' are necessary to more easily explain two different aspects of a given phenomenon. In the words of Tong (2018, p. 1):

> The [gauge] redundancy allows us to make manifest the properties of quantum field theories, such as unitarity, locality, and Lorentz invariance, that we feel are vital for any fundamental theory of physics but which teeter on the verge of incompatibility. If we try to remove the redundancy by fixing some specific gauge, some of

these properties will be brought into focus, while others will retreat into murk. By retaining the redundancy, we can flit between descriptions as is our want, keeping whichever property we most cherish in clear sight.

In sum, the representational scheme can be seen as specifying a functional role, using different pragmatic, explanatory, and theoretical requirements. Needless to say, there is no empirical breaking of gauge-symmetry. This concludes my response to what I called Healey's 'interesting challenge': showing that intra-theoretic resources enable us to choose representational schemes.

5 Representational schemes overcoming Limitation (B)

I take Limitation (A)—the issue described in Section 2, that TS left open the matter of providing sufficient and tractable conditions of identity for the isomorphism-classes—to have been essentially overcome in Section 3, clarified in Section 4, and exemplified in Appendix A.

But I still have not explicitly addressed how representational schemes also overcome Limitation (B). This is the aim of this Section.

To recap, the second Limitation is that (cf. Section 2.1.b):

B. TS denies that objects have *primitive* identity across physical possibilities, or, equivalently, across different isomorphism-classes. So can we find an alternative counterpart relation for objects belonging to non-isomorphic models?

For concreteness, and since most of the literature focuses on Limitation (B) in the case of spacetime, I will here also focus on that application; but most conclusions extend to the gauge case (by replacing 'points' by 'internal frames' of a vector space).

More specifically, I will address Limitation(B) through another Lewisian topic: countepart theory. Counterpart theory is the philosophical doctrine, due to David Lewis ((Lewis, 1968), (Lewis, 1973, p. 38-43) and (Lewis, 1986, Ch. 4)) that any two objects—in particular, spacetime points—in two different possibilities (in philosophical jargon: possible worlds) are never strictly identical.[25] They are distinct, though of course similar to each other in various, perhaps many, respects. Representational schemes, specified by extra-empirical criteria as in Section 4, can be construed as picking out which respects are important for the issue at hand.[26]

As to counterfactuals, what makes true a proposition that the object a of the actual world could have had property F (though in fact it lacks F) is not that in another possible world, a *itself* is F, but that in another possible world, which is sufficiently similar to the actual world, an object appropriately similar in certain respects to a, is F. That object is called the *counterpart*, at this other possible world, of a.

The ideas is that each representational scheme via gauge-fixing provides explicit and invariant relations between non-isomorphic models which define local counterparts. Each gauge

[25]For Lewis, 'possibilities' and 'possible worlds' are not exactly the same thing: possibilities require the stipulation of a world and a choice of counterpart relation.

[26]This idea of using gauge-fixings as counterpart relations was introduced in (Gomes, 2022) and developed in fullness in (Gomes & Butterfield, 2023a); here I will give only a sketch.

fixing fixes a web of relations that is maximally rigid with respect to locations and simultaneously maximally loose with respect to the physical content of those locations. Thus we can specify spacetime points in different possibilities—in non-isomorphic models—by their location within such a web of relations and compare the values of these physical quantities that are not fixed by that web.[27]

In Section 5.1 I provide the basic formalism for conceiving of counterpart relations as given by representational schemes and discuss how changing representational schemes affects the counterpart relations; in Section 5.2 I describe obstructions to construing counterpart relations in this manner—obstructions due to homogeneous models.

5.1 Basic formalism

Mathematically, I will construe counterpart relations very narrowly: as group elements $g \in \mathcal{G}$ that relate two different, not necessarily isomorphic, models φ_1 and φ_2. For a defence of this narrow mathematical interpretation of counterparthood, see (Gomes & Butterfield, 2023a).

Again, for concreteness, I will focus on the case of general relativity. Given two models for vacuum general relativity, $\langle M, g^1_{ab}\rangle$ and $\langle M, g^2_{ab}\rangle$, a diffeomorphism $f \in$ $maths f Diff(M)$ will give us a counterpart relation between the spacetime points of M in each of the two models.[28]

Given a section $\mathcal{F}_\sigma$, as described in Section 3.3, we have just such an element: the dressing $g_\sigma : \Phi \to \mathcal{G}$, given in Equation (3.6). Using this convenient mathematical operator, the counterpart relation between φ_1 and φ_2 is given by:

$$\text{Counter}_\sigma(\varphi_1, \varphi_2) := g_\sigma(\varphi) g_\sigma(\varphi_2)^{-1}. \tag{5.1}$$

Generally, i.e. for models that are not along the section, the relation is given by the group element that takes the first model down to the gauge-fixing section and then back up towards the second model. Importantly, by being associated to a gauge-fixing, we don't limit the physical content that we seek to compare via counterparthood.

Since g_σ is generally a non-local functional of the model, the counterpart relation may also depend non-locally on the models. But the counterpart relation between any two states that already lie at the gauge-fixing section is just the identity: they are both already in their preferred representation relative to the reference frame that is associated to the gauge-fixing, so these counterpart relations are trivially local. In other words, it is the reference frame associated to the gauge-fixing that defines locality.

[27]In this Section, since we will be comparing objects across different possibilities, and using different comparisons, I prefer to talk about *specifying* a point or region, rathern than individuating it, since the latter term's connotation of identity is stronger than is required here, and could lead to regions bearing their content essentially, and not leading to an interesting notion of counteparthood; cf. Section 2.3.

[28]Similarly, given a vector bundle, E, the bundle of admissible frames over E, $P \simeq L_G(E)$, and two sections of E, $\varphi, \varphi' \in \Gamma(E)$, the counterpart relations tells which frame over x for φ corresponds to which frame over x for φ'. It thus can be seen as a vertical automorphism between two models of Yang-Mills theory: $\langle P, \omega_1\rangle$ and $\langle P, \omega_2\rangle$.

Such a notion of spacetime counterparts has interesting properties. Due to covariance (see (3.8)),

$$\mathrm{Counter}_\sigma(\varphi_1^g, \varphi_2^g) = g^{-1} g_\sigma(\varphi_1)(g^{-1} g_\sigma(\varphi_2))^{-1} = g^{-1}\,\mathrm{Counter}_\sigma(\varphi_1, \varphi_2) g\ . \tag{5.2}$$

Moreover, two models that lie in the same orbit will always be related by the unique isomorphism that connects them. That is: if $\varphi_2 = \varphi_1^g$, then $\mathrm{Counter}_\sigma(\varphi_1, \varphi_2) = g$, even if neither φ_1 nor φ_2 lie in the section σ. For with the definitions above,

$$\mathrm{Counter}_\sigma(\varphi_1, \varphi_1^g) = g_\sigma(\varphi_1)(g^{-1} g_\sigma(\varphi_1))^{-1} = g \tag{5.3}$$

Thus, for example, in the spacetime case, if the unique counterpart of p in model $\langle M, g_{ab}\rangle$ is q in model $\langle M, g'_{ab}\rangle$, then the counterpart of p in model $\langle M, f^* g'_{ab}\rangle$ will be $f(q)$. And this particular property is independent of which scheme σ we choose.

In other words, though there are several distinct choices of counterparts, each choice must identify the same spacetime points across isomorphic models. In the spacetime context, this is essentially Field (1984, p. 77)'s observation that

> "individuation of objects across possible worlds" is sufficiently tied to their qualitative characteristics so that if there is a unique 1-1 correspondence between the space-time of world A and the space-time of world B that preserves all geometric properties and relations (including geometric relations among the regions, and occupancy properties like being occupied by a round red object), then it makes no sense to suppose that identification of space-time regions across these worlds goes via anything other than this isomorphism.

I will call the existence of such a 1-1 correspondence 'Field's criterion' for short. It was the centerpiece of Weatherall (2018): the paper that sparked a recent surge of interest in the hole argument in general relativity. I will call Field's suggestion (for the interpretation of isomorphic models that satisfy Field's criterion) 'the drag-along' interpretation of isomorphic models (cf. Butterfield & Gomes (2022) for more details).[29] Here we see that the drag-along interpretation of isomorphic models is automatically implemented by our notion of counterpart relations using representational schemes. But we do not always have universal, unique, counterpart relations, and so Field's criterion does not always apply. We now turn to this.

5.2 A snag for counterparts

As we have just seen, given any representational scheme via gauge-fixing, as long as the web of properties and relations in each of our models is sufficiently complex, it suffices to uniquely specify points across models, both isomorphic or not.

[29] Bradley & Weatherall (2022); Weatherall (2018) associate the 'the drag-along' idea less with Field and more with Mundy (1992), who distinguishes a theory's synthetic language—that is only able to express 'qualitative' facts—from a theory's metalanguage, in which we are able to talk about points singularly, i.e. without a definite description (see (Gomes, 2021c, Sec. 3.A)). But the gist is the same: isomorphisms— expressed in the metalanguage in Mundy's case—will map objects singled out by the same description into each other.

But if this web is not sufficiently complex, the specification will fail. To be blunt: we cannot stipulate a scheme for some region in spacetime or frame in a vector space if there aren't enough specific physical features to single out that region or frame.

As described in Appendix B (see also end of Section 3.2), models that are 'too homogeneous' or symmetric have stabilizers, which represent a degeneracy in the web of relations that they can furnish. Sometimes, moving from point x to point y makes no difference to the web. That is, on homogeneous states, $\tilde{\varphi}$, a representational scheme—seen as a map $\sigma : [\Phi] \to \Phi$—will fail to uniquely specify the g_σ of (3.6). The reason is that g_σ fails to satisfy the property of *Uniqueness* of Section 3.3. There are particular $\tilde{g} \in \mathcal{G}$ such that $\tilde{\varphi}^{\tilde{g}} = \tilde{\varphi}$, we call this the *stabiliser* group, I_φ. The existence of stablisers implies, in particular, that we cannot fulfill the covariance Equation (3.8). In the words of Gomes et al. (2019, Sec. 7):[30]

> The fact that [the reference frame] is left undetermined in the directions belonging to the stabilizer should have been expected from the simple fact that a field insensitive to some transformations cannot be used as a reference to [discern those transformations].

More broadly, this type of obstruction due to homogeneity is well known, and related to Black (1952)'s criticism of Leibniz's Principle of the Identity of Indiscernibles—the principle that motivates Leibniz equivalence (see (Pooley, 2022, Sec. 3.3) for a thorough exposition).

One might have thought that, precisely because the degeneracy in g_σ acts trivially on the respective models, we could ignore it entirely. If its action on the models is trivial, why consider it at all? However, by being more mathematicaly rigorous about counterpart relations, we uncover a snag which has not been noticed in the philosophical literature. The snag is only for isomorphism groups that are non-Abelian, which is the case for general relativity and Yang-Mills theory; and it occurs only for very homogeneous models, on which these groups don't act freely. For such groups the quotient of the entire gauge group by the stabiliser subgroup may not form a group. This will occur if the stabiliser subgroup is not a normal subgroup of the gauge group. And so, in these cases, we can't simply define a counterpart relation 'up to stabilisers' with all of the properties described above. Here is a familiar example: the Minkowski metric admits a non-trivial action of any diffeomorphisms that is not a Poincaré transformation, but there is no well-defined quotient group of "diffeomorphisms up to Poincaré" which we could use for a representational scheme or gauge-fixing.

What should we do in these cases? There are two alternatives. First, since stabilisers themselves form a group, one could still plausibly construct a counterpart relation for that group, valid for all configurations that share that stabiliser group (up to conjugacy). So, in the Minkowski example, one 'switches' to the theory of special relativity, with a different notion of isomorphism, and a corresponding notion of representation scheme. The second alternative is a bit trickier and it uses the slice theorem, expounded in Appendix B. This would require us to

[30](Gomes et al., 2019)'s discussion is in the context of the 'relational connection-form', to be briefly described in Section 6.2. But in Section 7 (ibid), they are in particular describing a case in which that connection is integrable, and thus recovers a gauge-fixing section.

pick a section of the gauge group, seen as a principal fiber bundle whose structure group is the stabliser group! So, in the case of diffeomorphisms, additionally to the standard representation scheme given in (3.3), we would have a local section (around the identity coset) of the bundle $\mathcal{I}_\varphi \hookrightarrow \mathcal{G} \to \mathcal{G}/\mathcal{I}_\varphi$, i.e. :

$$\begin{aligned} \overline{\rho} : [\mathcal{G}] \times \mathcal{I}_\varphi &\to \mathcal{G} \\ ([g], \eta) &\mapsto \eta\sigma([g]). \end{aligned} \tag{5.4}$$

In the example of the Minkowski spacetime, we would have to choose, using further extra-empirical criteria, some representational scheme for the diffeomorphisms relative to the Poincaré transformations.[31]

In sum, for very homogeneous states, g_σ is not at first glance well-defined, so the counterpart relation (5.1) is no longer valid and will not recover Field's criterion. Nonetheless, the interpretation of isomorphic models through TS survives unscathed and, when the stabiliser group is a normal subgroup of $\mathcal{G}$, this interpretation is preserved by representational schemes, which are still well-defined (but now the representational schemes apply to the quotient group). The investigation of counterpart relations for non-normal stabiliser groups can still be accomplished however, via two different routes. But a further examination of these routes will be left for future work.

6 Summing up

In this Section, I conclude, by first, in Section 6.1, offering a very brief summary. Then, in Section 6.2, I list two idealizations involved in representational schemes, which I have so far skatted over.

6.1 Representational schemes conceptually summarised

In both the gravitational and the gauge cases, there are many ways in which we can represent structurally identical patterns; each of these ways correspond to one of many *isomorphic models* of the theory. But we can pick particular representations by resorting to a *representational scheme.*

Representational schemes introduce a notion of 'relationism' into the formalism. This is not necessarily the notion that is usually attributed to Mach, Poincaré, Barbour, etc, although there are similarities. For instance, though representational schemes also require relative values

[31]Let us look at this example in more detail, let $M \simeq \mathbb{R}^4$, $\mathcal{P} \subset \mathsf{Diff}(M)$ be the Poincaré group, where $\mathsf{Diff}(M)$ is the group of diffeomorphisms of M. Assume an action $\mathcal{P} \times \mathsf{Diff}(M) \to \mathsf{Diff}(M)$. A gauge-fixing of $\mathcal{P}$ then can be represented by some functional $\mathcal{F}$ on $\mathsf{Diff}(M)$, valued in the Lie algebra of $\mathcal{P}$, such that, for any $d \in \mathsf{Diff}(M)$,

$$\exists! \tau \in \mathcal{P}, \quad \mathcal{F}(d^\tau) = 0. \tag{5.5}$$

This has not yet been studied, as far as I know.

of fields or particle positions, since there are many choices of schemes, these relative values are not taken to correspond uniquely to the 'fundamental ontology'.

Let us take general relativity as an example for examining the metaphysical interpretation that accompanies representational schemes. Alongside most philosophers of physics, I accept that the implied fundamental ontology is firmly about distances between spacetime points. In my view, TS and anti-haecceitism say all there is to say about that. But even being clear about the fundamental ontology, or the type of structure in question, we have at first glance no concise way of qualitatively describing how these distances are distributed, since we have no concise way of qualitatively designating the points that stand in these relations. Here then is the role of relationism as embodied by representational schemes: each representational scheme is a choice of a set of relations used to describe physical content without redundancy, i.e. satisfying Definition 1. In the case of spacetime, in metaphysical jargon, one can think of each representational scheme as providing a particular qualitative description of these distributions of distances as described relative to some reference frame.

Thus one main difference between mine and the traditional notion of relationism is that I explicitly admit multiple choices of sets of relations, leading to many distinct possible reductions of the theory. One set can only be *pragmatically* preferred to others. Indeed, when Einstein famously overcame his doubts about the 'hole argument' by appealing to 'point coincidences' (cf. (Giovanelli, 2021)), he assumed there was no further question as to what they were coincidences *of*. But, as we have seen, different choices can yield very different physical descriptions of a given state of affairs; very different coincidences, which equally well coordinitize a given region of spacetime.

There is then the matter of which scheme to choose. But there is no great mystery here: to pick a scheme we must employ intra-theoretic resources by the lights of which the scheme is appropriate.

In sum, representational schemes can be construed as (non-unique) choices of a relational set of quantities that:

(I) provide tractable, complete characterizations of the individual isomorphism-classes, i.e. of the structure that, given a fixed notion of isomorphism, is: common to each class of isomorphic models, and different for different classes of isomorphic models. They give a tractable criterion of individuation of isomorphism-classes.

(II) allow us to describe qualitative counterpart relations for points, regions, frames, etc, across isomorphic and non-isomorphic models, while entirely denying any primitive—i.e. scheme-independent—identification across models.

And so, to complement my first two Desiderata for TS of (Gomes, 2021c), I here introduce Desideratum (iii): that the theory's symmetries admit representational schemes. Fulfillment of this Desideratum overcomes the first two Limitations of Section 2.1.b, left over from (Gomes, 2021c). But is Desideratum (iii) a non-trivial requirement?

6.2 Limitations of representational schemes and what lies beyond them.

Representational schemes overcome some of the limitations of TS, but they are also limited in two ways I have only so far mentioned in passing.

The first way concerns the topology of the base set of independent variables. Let me illustrate with an example. Take De Donder variables as defined by four independent massless scalar fields $\phi^{(i)} : M \to \mathbb{R}, i = 1, ..., 4$, satisfying wave equations (cf. Equation (A.25)). If their domain M is a generic spacetime, that spacetime will (generically) not be homeomorphic to $\mathbb{R}^4$. As we know, there are topological limitations to the existence of such homeomorphisms, and there might be other obstructions as well. So we cannot locate an arbitrary point of spacetime by the four values of the scalar fields. The upshot is that representational schemes aren't usually global in spacetime.

Neither are they global in field-space; this is the second kind of limitation and it is due to an obstruction closely associated to the existence of 'Gribov horizons' (see Appendix B). Namely, for theories with non-Abelian structure groups, there are no gauge-fixings that cover the entire space of models (cf. (Gribov, 1978; Singer, 1978)). The principle behind the obstruction is easy to illustrate by thinking of gauge-fixings as instantiated by physical reference frames (see footnote 16): any particular choice of reference frame is thus physically instantiated, so there can be states of the world that do not admit this instantiation. What is perhaps more surprising is that any physical reference frame for non-Abelian theories is limited in this way.

Thus we see that not all theories with symmetries that satisfy the necessary conditions for TS admit representational schemes in the way I defined them here. In fact, the limitations above show that only Abelian theories in a simply connected background can satisfy Definition 3 without caveats. Nonetheless, the general lessons obtained from representational schemes remain for non-Abelian theories by thinking of theories locally in both the background manifold and on the space of models.

I'll end this paper on a rather technical comment, about a generalisation of representational schemes via gauge-fixing that overcomes the limitation due to the Gribov horizons. In the original treatment of Gomes et al. (2019); Gomes & Riello (2017, 2018, 2021), obstructions due to Gribov horizons were avoided by a more flexible tool than representational schemes, called a 'relational connection-form'. In brief, this is a distribution of infinitesimal gauge-fixing surfaces, with appropriate covariance properties. These connection forms generalize parallel transport in principal fiber bundles to parallel transport in the space of models, where the gauge group forming the orbits is the entire group of gauge transformations. In this new context, connection-forms provide counterpart relations between members of any one-parameter sets of models—i.e. along 'histories' of models—and so generalise Barbour's 'best-matching' (see e.g. (Mercati, 2017, Ch. 4)). Heuristically, this can be understood as picking out, along two neighboring orbits, the two models that are the closest, according to some choice of gauge-invariant notion of 'closeness' that applies between infinitesimally different models.

The transversal, covariant distribution defined by a relational connection form is not necessarily integrable into 'slices' (see Appendix B), which means it usually has curvature in Φ, and that it is inequivalent to any gauge-fixing. In this case, since they are history-dependent, counterpart relations given by connection-forms will not generally satisfy 'Field's criterion' (see (Gomes et al., 2019) for a pedagogical introduction).[32]

Acknowledgements

I would like to thank Jeremy Butterfield for many detailed comments on several versions of this paper, and Frank Cudek, for comments and pointing me to Curiel's paper on the problem of counterfactuals for general relativity. I would also like to thank the British Academy for financial support.

APPENDIX

A Examples of Representational Schemes

Representational schemes can be very general. Here I will provide several examples, mostly based on gauge-fixings, starting from non-relativistic particle mechanics to gauge theory and general relativity.

A.1 Non-relativistic particle mechanics

A.1.a Center of mass for non-relativistic point-particles

Take Newtonian mechanics for a system with N particles, written in configuration space as trajectories $\varphi(t) = q_\alpha(t) \in \mathbb{R}^{3N}$ (with $\alpha = 1, \dots, N$ labeling the particles), and with shift symmetry:

$$\varphi^g = (q^\alpha(t) + g), \quad \text{with } g \in \mathbb{R}^3. \tag{A.1}$$

Now we take the section $\mathcal{F}_\sigma$ to be defined by the center of mass:

$$\mathcal{F}_\sigma(q) = \sum_\alpha q^\alpha m^\alpha = 0. \tag{A.2}$$

The most straightforward way to find the projection $h(q)$ is displayed in (3.9). It requires us to find the dressing as a functional of the configuration. That is, take

$$h^\alpha(q) = q^\alpha + g_\sigma(q), \tag{A.3}$$

and solve

$$\mathcal{F}_\sigma(h(q)) = 0 \quad \text{for } g_\sigma : \Phi \to \mathbb{R}^3 \text{ and arbitrary } \; q. \tag{A.4}$$

[32]This more flexible tool (which was technically developed in (Gomes et al., 2019; Gomes & Riello, 2017, 2018, 2021) has been conceptually appraised in (Gomes, 2019, 2021a) and (Gomes & Butterfield, 2023a, Appendix A).

We obtain:

$$\sum_\alpha (q^\alpha + g_\sigma(q))m^\alpha = 0 \Rightarrow g_\sigma(q) = -\frac{\sum_\alpha q^\alpha m^\alpha}{\sum_\alpha m^\alpha}.$$

Clearly,

$$g_\sigma(\varphi^g) = g_\sigma(q+g) = g_\sigma(q) - g. \tag{A.5}$$

Thus if follows that the new set of dressed configuration variables, written as functionals of arbitrary configuration variables, is gauge-invariant (now including indices for both particles and components in $\mathbb{R}^3$, to be completely explicit):

$$h_i^\alpha(q^\alpha + g) = (q_i^\alpha(t)) + g + g_\sigma(q) - g = h_i^\alpha(q^\alpha). \tag{A.6}$$

And it is also easy to see that, if $h^\alpha(q) = h^\alpha(q')$, then $q^\alpha - q'^\alpha$ is given by a shift, and so $h^\alpha(q) = h^\alpha(q')$ iff q and q' are isomorphic.

This example is still non-local, in the sense of Section 3.4: given two particle subsystems, indexed by $N_+, N_- \subset N$, with $N_+ \cup N_- = N$,[33] the center of mass of either subsystem is generically not identical to the center of mass of their union. But, in the notation of Section 3.4, for $N_0 = N_+ \cap N_-$, it is straightforward to verify that, although $h^+_{|N_0} \neq h^-_{|N_0}$ (since the particles in N_0 are being described relative to different centers of mass), if $N_+, N_- \subset N$, with $N_+ \cup N_- = N$, equation (3.13) will hold, since both sides correspond to the description of the N_0 subsystem with respect to its own center of mass.[34]

A.1.b Inter-particle distances

This is the first example of a representational scheme which is *not* equivalent, at least *prima facie*, to any gauge-fixing. We take the same system as in the previous example, described in (A.1). And instead of (A.1), we take, following Rovelli (2014), a new complete set of relational variables:

$$\bar{q}_\alpha(t) := (q_{\alpha+1} - q_\alpha), \quad \text{for } \alpha = 1, \cdots, N-1. \tag{A.7}$$

Thus the new variables provide a shift-invariant description of the system. They are complete, since the dimension of the space of configurations of N particles in e.g. $\mathbb{R}^3$ is $\mathbb{R}^{3N}$, and modulo translations, this is $\mathbb{R}^{3(N-1)}$. So they satisfy Definition 1. But there is a clear disadvantage to this kind of description: because the invariant variables are not elements of the original space of models, Φ, it is not a representational scheme in the sense of Definition 2 (or the more

[33] I am here denoting a set of particle labels with $N_\pm, N$, whereas before N was a number. But this slight innacuracy does not justify the amount of notation that would need to be introduced in order to clarify this point.

[34] Of couse, we need not have $N_+, N_- \subset N$, with $N_+ \cup N_- = N$. That is, we need not have h_+, h_- coming from the restrictions of a consistent global model: equation (3.13) is supposed to assess consistency, not assume it. If $h_0(h^+_{|N_0}) \neq h_0(h^-_{|N_0})$ there is a bijection between the labels of the particles in N_0 according to N_- and N_+, but the intrinsic statess of the restricted subsystems are not symmetry-equivalent.

precise Definition 3). So we cannot iteratively apply representational schemes in order to assess consistency, as in (3.13).[35]

Of course, examples for non-relativistic particle mechanics abound: beyond the case of translations and boosts, we could fix a frame in $\mathbb{R}^3$ by diagonalizing the moment of inertia tensor around the center of mass; in this case g_σ would be a state-dependent rotation (see (Gomes, 2021b, Sec. 4) for details).[36]

A.2 Gauge theory

A.2.a Electromagnetism: Coulomb and Lorenz gauge

Electromagnetism in the gauge potential formalism is a gauge theory, and in it, Lorenz gauge is an explicitly Lorentz covariant choice, which has its virtues thoroughly extolled by Mattingly (2006), who argues that it should be considered as Maudlin (1998)'s "ONE TRUE GAUGE". Maudlin (2018) himself endorses a different choice, Coulomb gauge.

Coulomb gauge arises from a split of the electric field into a component that is purely Coulombic, or determined by the synchronic distribution of charges, and another component that is purely radiative. And, apart from worries with *Uniqueness* (described in Section 3.3.b, associated to the Gribov problem (discussed in Appendix B), there are straightforward extensions of these gauge-fixings to the non-Abelian case. This particular example of $\mathcal{F}_\sigma$ (and h_σ, and g_σ) in field theory is fully worked out by Gomes & Butterfield (2022), alongside its physical interpretation. Lorenz gauge is explicitly Lorentz covariant, and has the representation of the gauge potential as determined by the past distribution of currents (see footnote 37).

Although conceptually very different, these two gauges are in fact mathematically very similar: the main difference is that the Lorenz gauge concerns all the components of the potential in a Lorentzian manifold, whereas Coulomb gauge concerns only spatial components of the gauge potential, so is described in a Riemannian manifold. But here I will ignore the difference by construing Lorenz gauge in a spacetime manifold with Euclidean signature, and Coulomb gauge to be a choice of gauge in the Hamiltonian version of electromagnetism.[37]

[35] Suppose we were coupling two subsystems for which $N_+ \cap N_- = \emptyset$. Now we have $\mathbb{R}^{3(N_\pm - 1)}$ degrees of freedom in each subsystem. But, after coupling, we have $\mathbb{R}^{3(N_+ + N_- - 1)} \neq \mathbb{R}^{3(N_+ - 1)} + \mathbb{R}^{3(N_- - 1)}$. Indeed, the coupled system has 3 more degrees of freedom than the individual subsystems. Rovelli (2014) argues that these are the gauge degrees of freedom that need to be retained for the coupling of subsystems. Gomes & Riello (2021) show that the same can be applied to gauge theories. The difference is that if we restrict the division of the fields to supervene on a complete division of the manifold by complementary regions, the number of extra, holistic degrees of freedom is the dimension of the stabilliser group of the states at the boundaries between the complementary regions. See Gomes (2021a) for a conceptual exposition.

[36] Configurations that are collinear, or are spherically symmetric, etc. would be unable to fix the representation: these are *reducible* configurations; they have stabilizers that cannot be fixed by any feature of the structure that is isomorphism-invariant.

[37] The difference matters: the Lorenz choice in a Lorentzian manifold gives rise to a hyperbolic, and Coulomb always gives rise to an elliptic, partial differential equation. In the Lorentzian case with a Lorenz choice, the Maxwell equations of motion for A_a become the hyperbolic equation (which gives a well-posed initial value problem: see Gomes (2021e) for more details about the relation between the IVP and gauge freedom):

In the Coulomb case for Hamiltonian electromagnetism, the field φ is a doublet that transforms only in one component, i.e.

$$\varphi = (E^i, A_i), \quad \varphi^g = (E^i, A_i + \nabla_i g), \tag{A.8}$$

with $g \in C^\infty(M)$, where M is a spatial (Cauchy) surface, E_i is the electric field, and A_i is the spatial gauge potential, with spatial indices i, j etc. The Lorenz case would be similar, but M would be the spacetime manifold and we would have:

$$\varphi = A_\mu, \quad \varphi^g = A_\mu + \nabla_\mu g, \tag{A.9}$$

with μ, ν, etc, spacetime indices. In the following, I will subsume both cases using the abstract index notation A_a.

First, following the definition of the dressed function, in (3.9), we write:

$$h(\mathbf{A})_a := A_a + \nabla_a g_\sigma(\mathbf{A}), \tag{A.10}$$

where $g_\sigma : \Phi \to C^\infty(M)$, and M is either the space or the spacetime manifold (with Euclidean signature), and Φ is the space of gauge potentials over M satisfying appropriate boundary conditions (which I will not here discuss: for more details, about boundary conditions and the point that this gauge is not complete in the Lorentzian signature: see footnote 37 and (Gomes & Riello, 2021)). The condition to be satisfied by g_σ will be obtained from our definition of $\mathcal{F}_\sigma : \Phi \to C^\infty(M)$, with:

$$\mathcal{F}_\sigma(\mathbf{A}) = \nabla^a A_a = 0. \tag{A.11}$$

That is, we obtain:

$$\mathcal{F}_\sigma(h(\mathbf{A})) = \nabla^a (A_a + \nabla_a g_\sigma(\mathbf{A})) = 0, \tag{A.12}$$

and thus:

$$g_\sigma(\mathbf{A}) = -\nabla^{-2}\nabla^b A_b, \tag{A.13}$$

where ∇^{-2} is the inverse Laplacian (in the Lorentzian case, given sufficiently strict boundary conditions, it would be a Green's function, $\Box^{-1}$, with $\Box$ the D'Alembertian). Equation (A.13) clearly satisfies the covariance equation (3.8), namely:

$$g_\sigma(\mathbf{A}^g) = -\nabla^{-2}\nabla^b (A_b + \nabla_b g) = g_\sigma(\mathbf{A}) - g. \tag{A.14}$$

Since the inverse Laplacian is determined only up to a constant, so is g_σ, but this is a happy case in which the degeneracy in g_σ is a stabiliser of the gauge potential (and the stabilisers form a normal subgroup of $\mathcal{G}$; cf. Appendix B), and so has no effect on the dressed variable.

$\Box A_a = j_a$, This equation still has the gauge-freedom corresponding to ϕ such that $\Box g = 0$. To fix g uniquely one requires strict boundary conditions, which could still carry the non-local behavior we have alluded to in Section 3.4 (cf. Gomes (2021e)). Thus, in the Lorentzian setting, this gauge-fixing is not complete and we would require additional imput about the initial state. In contrast, the kernel of the elliptic equation $\nabla^2 g = 0$ corresponding to Coulomb gauge can be defined in the absence of spatial boundaries: it is zero. So Coulomb, but not Lorenz, define a *bona-fide* gauge-fixing condition without the need of auxiliary conditions.

The corresponding dressed functional is:

$$h(\mathbf{A})_a := A_a^{g_\sigma(A)} = A_a - \nabla_a g_\sigma(A) = A_a - \nabla_a(\nabla^{-2}\nabla^b A_b) \tag{A.15}$$

Gauge-invariance of the dressed function is implied by (A.14) and can be immediately verified in (A.15). And of course, from (A.10), it is clear that if two dressed potentials match, they are dressings of gauge-related potentials. Thus two dressed variables match iff what is being dressed is related by an isomorphism.

A.2.b Maxwell Klein-Gordon: unitary gauge

This gauge is only available for some sectors of Abelian theories like electromagnetism with a nowhere vanishing charged scalar (Maxwell Klein-Gordon). As we will see, this choice of gauge is *sui generis* for being local (see Wallace (2024) for an in-depth analysis of the merits and shortcomings of this gauge).

The Maxwell-Klein-Gordon theory has models $\varphi = (A_\mu, \psi)$, with A_μ the gauge potential and ψ a complex scalar function. The gauge transformation of the models is defined as:

$$\varphi^g = (A_\mu + \partial_\mu g, \psi e^{ig}), \tag{A.16}$$

with $g \in C^\infty(M)$, and M representing spacetime.

In the unitary gauge, we have the condition

$$\mathcal{F}_\sigma(\varphi) = |\psi| - \psi = 0. \tag{A.17}$$

The dressed variables are:

$$h(A, \psi) = (A_\mu + \partial_\mu g_\sigma(\psi), \psi e^{ig_\sigma(\psi)}), \tag{A.18}$$

where we solve $F(h(\varphi)) = 0$ for $g_\sigma : \Phi \to C^\infty(M)$ and arbitrary φ. Assuming $|\psi|(x) \neq 0, \forall x \in M$, we write:

$$\psi = |\psi| e^{i\theta}, \text{ with } \theta = -i \ln \frac{\psi}{|\psi|}.\text{[38]} \tag{A.19}$$

Finally, (A.17) and (A.18) imply that:

$$|\psi| e^{i(\theta + g_\sigma(\psi))} = |\psi|, \tag{A.20}$$

so we find

$$g_\sigma(\psi) = -\theta, \tag{A.21}$$

which clearly has the right gauge-covariance properties. Thus,

$$h(A, \psi) = h(A, |\psi|) = (A_\mu - \partial_\mu \theta, |\psi|), \tag{A.22}$$

[38] There is an issue here that the logarithm of complex functions is not well-defined over the entire $\mathbb{C} - \{0\}$, because of the periodicity of solutions. Thus we should choose a subset of $\mathbb{C} - \{0\}$ that contains a single branch; but nothing depends on the branch.

which is gauge-invariant. But subtracting $\partial_\mu\theta$ from A_μ imposes no constraint on the values that $A_\mu - \partial_\mu\theta$ can take. Thus we can redefine $A_\mu - \partial_\mu\theta \to \tilde{A}_\mu$ as an unconstrained, invariant 1-form. We thus rewrite the dressed fields as $(\tilde{A}_\mu, \rho)$, where $\rho \in C^\infty_+(M)$ (it is an everywhere positive smooth scalar function) and $\tilde{A}_\mu \in C^\infty(T^*M)$. So this is an explicitly local representation of the isomorphism-class; in particular, the consistency of piecewise-invariant descriptions as in (3.14), becomes trivial to assess.

A natural extension of this unitary gauge to the non-Abelian case takes a nowhere vanishing matter field to define an internal direction in the vector spaces where the fields take their values.[39]

In closing, this list of gauge-fixings for Yang-Mills theory is of course not exhaustive. There are many other, physically motivated choices we can make. For instance, we may also want to highlight the helicity degrees of freedom of the theory, in which case we would use temporal gauge. But temporal gauge, like the Coulomb one, depends on a foliation of spacetime by spacelike surfaces, and so, strictly speaking, cannot be given a univocal physical interpretation unless coupled to a representational scheme about spacetime foliations. We turn to this in Section A.3. Before getting there, I will now exhibit a choice of dressed variables, that, like the choice of inter-particle distances for non-relativistic mechanics described in Section A.1.b, is *not* obtained by a gauge-fixing section.

A.2.c Holonomies

A holonomy basis for Abelian gauge theories associates to each loop in spacetime a gauge-invariant phase, namely:

$$hol_\gamma(A) = \exp i \int_\gamma A, \quad \text{for each} \quad \gamma : S^1 \to M. \tag{A.23}$$

More generally, the holonomy thus defined is Lie-group-valued, and, in the non-Abelian case, it is not entirely gauge-invariant, but transforms covariantly in the adjoint representation. Nonetheless, there is a straightforward way to regain gauge-invariant variables by taking the trace of the holonomy: these are called *Wilson-loops*.[40]

The Wilson-loop basis for gauge-invariant quantities is a representational scheme in the sense of Definition 1, but not in the sense of 3: although they 'dress' the original variables, they do not take values in the original space of models (they assign numbers to loops in spacetime). Thus the consistency of the composition of piecewise-invariant descriptions (discussed in Section

[39] In the context of covariant perturbative gauge-fixings, this was labeled 'the Higgs relational connection' in (Gomes et al., 2019, Sec. 7). In the case of SU(2), a closely related decomposition of an arbitrary field using an everywhere non-zero internal direction is known as the The Cho-Duan-Ge (CDG) decomposition (Duan & Ge, 1979). It was rediscovered at about the turn of the century by several groups who were readdressing the stability of the chromomonopole condensate: the ensuing decomposition of the gauge potential helps in identifying the topological structures like monopoles and vortices in the gauge field configuration (see (Walker & Duplij, 2015) for a more recent concise summary of the literature and an application to QCD).

[40] Closed loops suffice in the vacuum case, but charges can be accomodated by integrating a gauge potential along a segment with charges at both its ends.

2.1 in the case of denumerable lists of invariants), cannot be settled by a simple criterion such as (3.13), which applies only for representational schemes via gauge-fixing.

In more detail: unlike dressed quantities obtained from gauge-fixings, the gauge-invariant variables derived from holonomies is vastly overcomplete. Thus the composition of these variables must obey certain constraints (cf. footnote 8). And whereas the consistency of the composition of piecewise-invariant descriptions via dressed variables obtained via gauge-fixing is settled by numerical coincidence of solutions of partial differential equations (cf. (3.13)), the constraints to be satisfied by the composition of holonomies are known as *Mandelstam identities*, and they can be rather complicated in the case of non-Abelian gauge theories. Thus, in order to assess the consistency of piecewise invariant descriptions via holonomies we need to know whether all the values of two distinct sets of Wilson-loops, beloging to two regions that partially overlap, can be obtained from the holonomies of a single connection.[41] And to assess that we need to be able to solve these constraints. This is, again, if not intractable, much harder than solving the differential equations in the gauge-fixing case.

Here we see complications for describing gluing or composition for representational schemes according to Definition 1 (but not satisfying Definition 3).

A.3 General relativity

In the case of general relativity in the Lagrangian formalism, I will start with De Donder gauge, which fixes coordinates so that the densitized metric is divergence-free:

$$\mathcal{F}_\sigma(g) = \partial_\mu(g^{\mu\nu}\sqrt{g}) = 0. \tag{A.24}$$

I will not bore the reader with yet another demonstration that this gauge gives rise to dressed gravitational fields with the expected properties, described in Section 3.3.[42]

One can also see De Donder gauge as anchoring coordinate systems on waves of massless scalars: four different solutions of the relativistic wave equation. That is, given coordinates in an initial hypersurface Σ (or a portion thereof), we *define* coordinates in some region of spacetime using four non-local scalar functionals of the metric, that solve four wave-equations:

$$\Box\mathcal{R}^\mu_g(p) = 0; \tag{A.25}$$

[41] Going in the other direction, if one assumes that all Wilson loops already satisfy the Mandelstam constraints, it *is* possible to reconstruct the corresponding isomorphism-classes, cf. (Barrett, 1991).

[42] We are here also in effect assuming the active-passive correspondence of (Gomes, 2021d, Sec. 5), to say that fixing a choice of coordinates is equivalent to picking out a unique model within an isomorphism class. In fact, as in the case of Lorenz gauge in gauge theory (cf. footnote 37), De Donder or harmonic gauge does not completely fix coordinate freedom in general relativity. Given the coordinates in an initial surface, the gauge uniquely defines coordinates to the past and future of that surface. Moreover, the use of coordinates implies that in general such conditions are only local. Giving a global, or geometric version of deDonder gauge can accomplished using an auxiliary metric, a measure of distance between metrics, and a diffeomorphism that is used to compare them (see (Landsman, 2021, Sec. 7.6)). This strategy—of using auxiliary metrics—is often employed in converting results originally expressed in particular coordinate systems into explicitly covariant results.

the variables depend functionally on the metric (possibly non-locally), but take values at each point of the manifold. In a less coordinate-centric language, the idea here is to define point $p_x(g_{ab})$ as the point in which a given list of scalars $\mathcal{R}^\mu_g(p)$, $\mu = 1, \cdots, 4$ takes a specific list of values, $(x_1, \cdots, x_4)$'). In other words, fixing the metric and initial values, the four scalar quantities, $\mathcal{R}^\mu_g$ define a map:

$$\mathcal{R}_g = (\mathcal{R}^1_g, \cdots \mathcal{R}^4_g) : M \to \mathbb{R}^4. \tag{A.26}$$

To pick out points $p \in M$ by the value of the quadruple we invert the map (A.26). Assuming that the map (A.26) is a diffeomorphism—in general it is only locally one—there is a unique value, for all of the models, of $g_{ab}(\mathcal{R}^{-1}_g)(x)$, for any $x \in \mathbb{R}^4$.[43] That is, applying the chain rule for the transformations of $\mathcal{R}_g$,

$$\forall f \in \mathsf{Diff}(M), \quad g_{ab} \circ \mathcal{R}^{-1}_g = f^* g_{ab} \circ \mathcal{R}^{-1}_{f^*g}. \tag{A.27}$$

So, using this quadruple, we have a unique representation of the metric on $\mathbb{R}^4$. Given some metric tensor $g^{\kappa\gamma}$ in coordinates x^κ, we can compute the metric in the new, $\mathcal{R}^\mu$ coordinate system as:

$$h^{\mu\nu} = \frac{\partial \mathcal{R}^\mu_g}{\partial x^\kappa} \frac{\partial \mathcal{R}^\nu_g}{\partial x^\gamma} g^{\kappa\gamma}. \tag{A.28}$$

This is just a family of 10 scalar functions indexed by μ and ν: the left-hand-side is the dressed variable, $h^{\mu\nu}_\sigma$, and the $\mathcal{R}^\mu$ are the dressing functions.

Of course, equations (A.26) and (A.27) would hold for any viable gauge-fixing: those choices also describe a local physical coordinate system, or a relational reference frame, on a patch of spacetime. And indeed, any such relation specifying points in terms of their 'qualitative properties' will be seen to explicitly furnish a specific counterpart relation across isomorphic and non-isomorphic models, thus resolving Limitation (B); cf. (Gomes & Butterfield, 2021).[44]

And there are many such examples. One that is widely used to study black hole mergers and initial value problems in general, is (the partial) 'CMC gauge' (cf. e.g. respectively (Pretorius, 2005; York, 1971)): this choice picks out clocks and simultaneity surfaces so that simultaneous observers measure the same local expansion of the universe. In this case, $\mathcal{F}_\sigma(\gamma_{ij}, \pi^{ij}) = \gamma^{ij}\pi_{ij}$ =const, where γ_{ij} is the spatial metric and π^{ij} is its conjugate momentum, obtained in the Hamiltonian (3+1) formulation of general relativity (Arnowitt et al., 1962).

Another example that does not require any fields other than the gravitational ones, is known as *Komar-Kretschmann* variables. For this example, we must first restrict our attention to spacetimes that are not homogeneous, i.e. generic spacetimes (i.e. excluding Pirani's type II

[43] More commonly, for each g_{ab}, there will be only a subset $U \subset M$ that is mapped diffeomorphically to $\mathbb{R}^4$.

[44] Curiel (2018, p. 468) construes qualitative identity of points similarly:

> Once one has the identification of spacetime points with equivalence classes of values of scalar fields, one can as easily say that the points are the objects with primitive ontological significance, and the physical systems are defined by the values of fields at those points, those values being attributes of their associated points only per accidens.

But he does not consider the element of choice of the schemes, as we do.

and III spaces of pure radiation, in addition to excluding symmetric type I spacetimes). Once this is done, we consider $\mathcal{R}^{\mu}_{g}(p))$, $\mu = 1, \cdots 4$, formed by certain real scalar functions of the Riemman tensor. Since the spacetimes considered here are suitably inhomogeneous, they all contain points in which these functions are linearly independent, and so can be used to specify location without limiting the physical content of the spacetime region *as a whole*.[45] This choice is exceptional for not requiring auxiliary structure: no fields other than the metric, nor fixed spacetime curves or points; the relational reference frame that it provides supervenes only on properties of the metric tensor.[46]

But given particular spacetime points or curves—taken to have, in Kripke's famous terminology, a *rigid designation*, across a suitably restricted set of models[47]—we can find adapted coordinate systems or reference frames. Perhaps the most famous choice of coordinates for general relativity, called 'Riemann-normal' coordinates, is obtained in this way: having rigidly fixed a tangent frame at an event across a certain set of isomorphism-classes (or physical possibilities), the local coordinates are obtained by applying the Riemann exponential map for the tangent space of that event. This choice is particularly useful for a set of possibilities contaning material systems whose scale is small compared to the curvature scale: in these coordinates the metric is described as almost flat along the trajectory of the system; the Riemann curvature appears only at second order in the proper distance to the freely-falling trajectory of the material system.

Another common example, indeed the most applied for navigation on Earth, is a coordinate choice based on GPS sattelites. This requires four sattelites to be rigidly designated (taken to be small, and non-back-reacting, and thus following timelike geodesics), which cross at an initial event and emit timed light-like signals. Those signals cover a spacetime region, and their values provide it with a coordinate system (cf. (Rovelli, 2002) for details). [48]

B Infinite-dimensional principal fiber bundles

In the case of field theories, such as general relativity and Yang-Mills, even admitting a free and proper action of an infinite-dimensional Lie group on an infinite-dimensional smooth manifold,

[45] Komar (1958) finds these real scalars through an eigenvalue problem:

$$(R_{abcd} - \lambda(g_{ac}g_{bd} + g_{ad}g_{bc}))V^{cd} = 0,$$

where V^{cd} is an anti-symmetric tensor. The requirement ensures solutions λ, whose existence we assume, are independent real scalar functions. Komar (1958, p.1183) takes these scalars to be preferred, "since they are the only nontrivial scalars which are of least possible order in derivatives of the metric, thus making them the simplest and most natural choice."

[46] See (Bamonti, 2023) for a classification of different types of reference frames in general relativity, according to their coupling to the metric and to the inclusion of back-reaction.

[47] A term is said to be *a rigid designator* when it designates (picks out, denotes, refers to) the same thing in all possible worlds in which that thing exists.

[48] Similarly to GPS coordinates, the 'dressed' diffeomorphism-invariant observables of Donnelly & Giddings (2016) are anchored on non-null spacetime curves.

it is not guaranteed that a local product structure exists everywhere: so Φ is not necessarily a *bona-fide* infinite-dimensional principal $\mathcal{G}$-bundle, i.e. $\mathcal{G} \hookrightarrow \Phi \to \Phi/\mathcal{G}$. The obstacle is that there are special states—called *reducible*—that have stabilizers, i.e. elements $\tilde{g} \in \mathcal{G}$ such that $\tilde{\varphi}^{\tilde{g}} = \tilde{\varphi}$. Moreover, the statement holds for the entire group orbit, as it is easy to show that for some $\tilde{\varphi}' := \tilde{\varphi}^g \in \mathcal{O}_{\tilde{\varphi}}$,

$$\tilde{\varphi}'^{g^{-1}\tilde{g}g} = \tilde{\varphi}^{\tilde{g}g} = \tilde{\varphi}^g = \tilde{\varphi}'$$

and so all the elements of the orbit are also reducible (with stabilizers related by the co-adjoint action of the group). And so orbits are of 'different sizes', and not isomorphic to the structure group. Nonetheless, there is a generalization of a section, called a slice, that provides a close cousin of the required product structure (see (Gomes & Butterfield, 2023a, Sec. 2.1, and footnotes 4,5)). As has been shown using different techniques and at different levels of mathematical rigour, Diez & Rudolph (2019); Ebin (1970); Isenberg & Marsden (1982); Kondracki & Rogulski (1983); Mitter & Viallet (1981); Palais (1961); Wilkins (1989) both the Yang-Mills configuration space and the configuration space of Riemannian metrics (called $\mathrm{Riem}(M)$), admit slices. These slices endow the quotient space with a *stratified* structure. That is, the space of models can be organised into orbits of models that possess different numbers of stabilizers; with the orbits with more stabilizers being at the boundary of the orbits of models with fewer stabilizers. For each stratum, we can find a section and form a product structure as in the standard picture of the principal bundle.

For both general relativity and non-Abelian gauge theories, reducible configurations form a meagre set. *Meagre* sets are those that arise as countable unions of nowhere dense sets. In particular, a small perturbation will get you out of the set (and this is true of the reducible states in the model spaces of those theories, according to the standard field-space metric topology (the Inverse-Limit-Hilbert topology cf. e.g. Fischer & Marsden (1979); Kondracki & Rogulski (1983)). In this respect, Abelian theories, such as electromagnetism, are an exception: *all* their configurations are reducible, possessing the constant gauge transformation as a stabilizer.

And apart from this obstruction to the product structure—i.e. even if we were to restrict attention to the generic configurations in the case of non-Abelian field theories—one can have at most a *local* product structure: no representational scheme, or section, giving something like (3.3), is global (this is known as the *Gribov obstruction*; see Gribov (1978); Singer (1978)). Unfortunately, the space of Lorentzian metrics is not known to have such a structure: it has only been shown for the space of Einstein metrics that admit a constant-mean-curvature (CMC) foliation.

In the infinite-dimensional case, both the dimension and the co-dimension of a regular value surface can be infinite, and it becomes trickier to construct a section: roughly, one starts by endowing Φ with some $\mathcal{G}$-invariant (super)metric, and then finds the orthogonal complement to the orbits, $\mathcal{O}_\varphi$, with respect to this supermetric. But here the intersection of the orbit with its orthogonal complement cannot be assumed to vanish, as it does in the finite-dimensional case. Nonetheless, in the cases at hand, that intersection is given by the kernel of an elliptic operator, and one therefore can invoke the 'Fredholm Alternative' (see (Gilbarg & Trudinger, 2001, Sec.

5.3 and 5.9)) to show that that intersection is at most finite-dimensional, but generically is zero, and thus the generic orbit has the 'splitting' property: the total tangent space decomposes into a direct sum of the tangent space to the orbit and its orthogonal complement. Now we must extend the directions transverse to the orbit, so as to construct a small patch that intersects the neighboring orbits only once. But the space of Riemannian metrics is a cone inside a vector space, so it is not even affine and we cannot just linearly extend the directions normal to the orbit and hope for the best. And so we extend the normal directions by using the Riemann normal exponential map with respect to the supermetric (cf. Gil-Medrano & Michor (1991)), and thus conclude that, for a sufficiently small radius, the resulting submanifold is transverse to the neighboring orbits and has no caustics. Finally, to show that this 'section' is not only transverse to the orbits, but indeed that it intersects neighboring orbits only once, the orbits must be embedded manifolds, and not just local immersions: this is guaranteed if the group action is proper, cf. Ebin (1970).

References

Arnowitt, R., Deser, S., & Misner, C. (1962). The dynamics of general relativity, pp. 227-264. In *in gravitation: an introduction to current research, l. witten, ed.* Wiley, New York.

Bagan, E., Lavelle, M., & McMullan, D. (2000). Charges from dressed matter: construction. *Annals of Physics*, *282*(2), 471–502.

Bamonti, N. (2023). *What is a reference frame in general relativity?*

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. (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

Bergmann, P. G. (1961). Observables in General Relativity. *Rev. Mod. Phys.*, *33*, 510–514. doi: 10.1103/RevModPhys.33.510

Black, M. (1952). The Identity of Indiscernibles. *Mind*, *61*(242), 153–164. Retrieved 2022-07-15, from http://www.jstor.org/stable/2252291

Bradley, C., & Weatherall, J. O. (2022). Mathematical responses to the hole argument: Then and now. *Philosophy of Science*, *89*(5), 1223–1232. doi: 10.1017/psa.2022.58

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

Butterfield, J., & Gomes, H. (2022). The Hole Argument: a Drama in Five Acts. *in preparation*.

Carrozza, S., & Höhn, P. A. (2022, feb). Edge modes as reference frames and boundary actions from post-selection. *Journal of High Energy Physics*, *2022*(2). doi: 10.1007/jhep02(2022)172

Curiel, E. (2015). *If metrical structure were not dynamical, counterfactuals in general relativity would be easy.*

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

Diez, T., & Rudolph, G. (2019, Aug). Slice theorem and orbit type stratification in infinite dimensions. *Differential Geometry and its Applications*, *65*, 176–211. Retrieved from http://dx.doi.org/10.1016/j.difgeo.2019.03.005 doi: 10.1016/j.difgeo.2019.03.005

Donnelly, W., & Freidel, L. (2016). Local subsystems in gauge theory and gravity. *JHEP*, *09*, 102. doi: 10.1007/JHEP09(2016)102

Donnelly, W., & Giddings, S. B. (2016, Jan). Diffeomorphism-invariant observables and their nonlocal algebra. *Physical Review D*, *93*(2). Retrieved from http://dx.doi.org/10.1103/PhysRevD.93.024030 doi: 10.1103/physrevd.93.024030

Duan, Y.-S., & Ge, M.-L. (1979). SU(2) Gauge Theory and Electrodynamics with N Magnetic Monopoles. *Sci. Sin.*, *9*(11). doi: 10.1142/9789813237278_0001

Earman, J. (1989). *World Enough and Spacetime.* MIT press.

Ebin, D. (1970). The manifold of riemmanian metrics. *Symp. Pure Math., AMS, 11*, *11,15*.

Field, H. (1984). Can we dispense with space-time? *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, *1984*, 33–90. Retrieved 2022-08-30, from http://www.jstor.org/stable/192496

Fischer, A. E., & Marsden, J. E. (1979). The initial value problem and the dynamical formulation of general relativity. In *General relativity : an einstein centenary survey. cambridge university press , new york, pp. 138-211.*

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

François, J. (2021, mar). Bundle geometry of the connection space, covariant Hamiltonian formalism, the problem of boundaries in gauge theories, and the dressing field method. *Journal of High Energy Physics*, *2021*(3). doi: 10.1007/jhep03(2021)225

François, J., Parrini, N., & Boulanger, N. (2021, dec). Note on the bundle geometry of field space, variational connections, the dressing field method, presymplectic structures of gauge theories over bounded regions. *Journal of High Energy Physics*, *2021*(12). doi: 10.1007/jhep12(2021)186

Geiller, M. (2017). Edge modes and corner ambiguities in 3d Chern–Simons theory and gravity. *Nucl. Phys.*, *B924*, 312-365. doi: 10.1016/j.nuclphysb.2017.09.010

Geiller, M., & Jai-akson, P. (2020, Sep). Extended actions, dynamics of edge modes, and entanglement entropy. *Journal of High Energy Physics*, *2020*(9). Retrieved from http://dx.doi.org/10.1007/JHEP09(2020)134 doi: 10.1007/jhep09(2020)134

Gilbarg, D., & Trudinger, N. (2001). *Elliptic Partial Differential Equations of Second Order*. Springer.

Gil-Medrano, O., & Michor, P. W. (1991). The Riemannian Manifold Of All Riemannian Metrics. *Quarterly Journal of Mathematics, (42), 183-202*.

Giovanelli, M. (2021, may). Nothing but coincidences: the point-coincidence and Einstein's struggle with the meaning of coordinates in physics. *European Journal for Philosophy of Science*, *11*(2). doi: 10.1007/s13194-020-00332-7

Giulini, D. (1995). Asymptotic symmetry groups of long-ranged gauge configurations. *Modern Physics Letters A*, *10*(28), 2059-2070. Retrieved from https://doi.org/10.1142/S0217732395002210 doi: 10.1142/S0217732395002210

Goeller, C., Hoehn, P. A., & Kirklin, J. (2022). *Diffeomorphism-invariant observables and dynamical frames in gravity: reconciling bulk locality with general covariance.*

Gomes, H. (2019). Gauging the boundary in field-space. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*. Retrieved from http://www.sciencedirect.com/science/article/pii/S1355219818302144 doi: https://doi.org/10.1016/j.shpsb.2019.04.002

Gomes, H. (2021a). Holism as the significance of gauge symmetries. *European Journal of Philosophy of Science, vol 11, 87*.

Gomes, H. (2021b). 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. (2021c). Same-diff? Conceptual similarities between gauge transformations and diffeomorphisms. Part II: Challenges to sophistication. *Arxiv: 2110.07204. Submitted.*.

Gomes, H. (2021d). Same-diff? Conceptual similarities between gauge transformations and diffeomorphisms. Part I: Symmetries and isomorphisms. *Arxiv: 2110.07203. Submitted.*.

Gomes, H. (2021e). Same-diff? Part IV: A compendium of similarities between gauge transformations and diffeomorphisms. *Arxiv: 2110.07204. Submitted.*.

Gomes, H. (2022). *Why gauge? Conceptual Aspects of Gauge theories.* arXiv. Retrieved from https://arxiv.org/abs/2203.05339

Gomes, H. (2024). Gauge theory as the geometry of internal vector spaces.

Gomes, H., & Butterfield, J. (2021). Counterpart relations in general relativity (forthcoming).

Gomes, H., & Butterfield, J. (2022, aug). How to Choose a Gauge? The Case of Hamiltonian Electromagnetism. *Erkenntnis*. doi: 10.1007/s10670-022-00597-9

Gomes, H., & Butterfield, J. (2023a, jun). The Hole Argument and Beyond: Part II: Treating Non-isomorphic Spacetimes. *Journal of Physics: Conference Series*, *2533*(1), 012003. Retrieved from https://dx.doi.org/10.1088/1742-6596/2533/1/012003 doi: 10.1088/1742-6596/2533/1/012003

Gomes, H., & Butterfield, J. (2023b, jun). The Hole Argument and Beyond: Part I: The Story so Far. *Journal of Physics: Conference Series*, *2533*(1), 012002. Retrieved from https://dx.doi.org/10.1088/1742-6596/2533/1/012002 doi: 10.1088/1742-6596/2533/1/012002

Gomes, H., Hopfmüller, F., & Riello, A. (2019). A unified geometric framework for boundary charges and dressings: Non-abelian theory and matter. *Nuclear Physics B*, *941*, 249 - 315. Retrieved from http://www.sciencedirect.com/science/article/pii/S0550321319300483 doi: https://doi.org/10.1016/j.nuclphysb.2019.02.020

Gomes, H., & Riello, A. (2017). The observer's ghost: notes on a field space connection. *Journal of High Energy Physics (JHEP)*, *05*, 017. Retrieved from https://link.springer.com/article/10.10072FJHEP0528201729017 doi: 10.1007/JHEP05(2017)017

Gomes, H., & Riello, A. (2018, Jul). Unified geometric framework for boundary charges and particle dressings. *Physical Review D*, *98*, 025013. Retrieved from https://link.aps.org/doi/10.1103/PhysRevD.98.025013 doi: 10.1103/PhysRevD.98.025013

Gomes, H., & Riello, A. (2021). The quasilocal degrees of freedom of Yang-Mills theory. *SciPost Phys.*, *10*, 130. Retrieved from https://scipost.org/10.21468/SciPostPhys.10.6.130 doi: 10.21468/SciPostPhys.10.6.130

Gribov, V. N. (1978). Quantization of Nonabelian Gauge Theories. *Nucl. Phys.*, *B139*, 1. ([,1(1977)]) doi: 10.1016/0550-3213(78)90175-X

Hall, N. (2015). Humean Reductionism about Laws of Nature. In *A companion to david lewis* (pp. 262–277). John Wiley and Sons, Ltd. Retrieved from https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118398593.ch17 doi: https://doi.org/10.1002/9781118398593.ch17

Harlow, D., & Wu, J.-Q. (2021). *Algebra of diffeomorphism-invariant observables in jackiw-teitelboim gravity.*

Healey, R. (2007). *Gauging What's Real: The Conceptual Foundations of Gauge Theories.* Oxford University Press.

Henneaux, M., & Teitelboim, C. (1992). *Quantization of gauge systems.* Princeton University Press.

Isenberg, J., & Marsden, J. E. (1982). A slice theorem for the space of solutions of einstein's equations. *Physics Reports*, *89*(2), 179–222.

Isham, C. J. (1992). Canonical quantum gravity and the problem of time. In *19th International Colloquium on Group Theoretical Methods in Physics (GROUP 19) Salamanca, Spain, June 29-July 5, 1992.*

Jacobs, C. (2022). Invariance or Equivalence: A Tale of Two Principles. *Synthese*, 1–21. doi: 10.1007/s11229-021-03205-5

Jaramillo, J. L., & Lam, V. (2021). Counterfactuals in the Initial Value Formulation of General Relativity. *The British Journal for the Philosophy of Science*, *72*(4), 1111-1128. Retrieved from https://doi.org/10.1093/bjps/axy066 doi: 10.1093/bjps/axy066

Kabel, V., de la Hamette, A.-C., Apadula, L., Cepollaro, C., Gomes, H., Butterfield, J., & Časlav Brukner. (2024). *Identification is pointless: Quantum reference frames, localisation of events, and the quantum hole argument.*

Komar, A. (1958, Aug). Construction of a Complete Set of Independent Observables in the General Theory of Relativity. *Physical Review*, *111*, 1182–1187. Retrieved from https://link.aps.org/doi/10.1103/PhysRev.111.1182 doi: 10.1103/PhysRev.111.1182

Kondracki, W., & Rogulski, J. (1983). *On the stratification of the orbit space for the action of automorphisms on connections. on conjugacy classes of closed subgroups. on the notion of stratification.* Inst., Acad. Retrieved from https://books.google.co.uk/books?id=LK0JrgEACAAJ

Landsman, K. (2021). *Foundations of General Relativity.* Radboud University Press.

Lavelle, M., & McMullan, D. (1995). Observables and gauge fixing in spontaneously broken gauge theories. *Phys. Lett.*, *B347*, 89-94. doi: 10.1016/0370-2693(95)00046-N

Lavelle, M., & McMullan, D. (1997). Constituent quarks from QCD. *Phys. Rept.*, *279*, 1-65. doi: 10.1016/S0370-1573(96)00019-1

Lewis, D. (1968). Counterpart Theory and Quantified Modal Logic. *Journal of Philosophy*, *65*(5), 113–126. doi: 10.2307/2024555

Lewis, D. (1970). How to Define Theoretical Terms. *Journal of Philosophy,* **67***,* pp. 427-446..

Lewis, D. (1972). Psychophysical and theoretical identifications. *Australasian Journal of Philosophy*.

Lewis, D. (1973). *Counterfactuals.* Cambridge, MA, USA: Blackwell.

Lewis, D. (1986). *On the Plurality of Worlds.* Blackwell Publishers.

Lewis, D. (2009). Ramseyan Humility. In D. Braddon-Mitchell & R. Nola (Eds.), *Conceptual analysis and philosophical naturalism* (pp. 203–222). MIT Press.

Martens, N., & Read, J. (2020). Sophistry about symmetries? *Synthese*. Retrieved from http://philsci-archive.pitt.edu/17184/

Mattingly, J. (2006). Which gauge matters? *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, *37*(2), 243-262. Retrieved from https://www.sciencedirect.com/science/article/pii/S1355219806000037 doi: https://doi.org/10.1016/j.shpsb.2005.08.001

Maudlin, T. (1988). The Essence of Space-Time. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, *1988*, 82–91.

Maudlin, T. (1990). Substances and space-time: What Aristotle would have said to Einstein. *Studies in History and Philosophy of Science Part A*, *21*(4), 531-561. Retrieved from https://www.sciencedirect.com/science/article/pii/0039368190900324 doi: https://doi.org/10.1016/0039-3681(90)90032-4

Maudlin, T. (1998). Healey on the aharonov-bohm effect. *Philosophy of Science*, *65*(2), 361–368. Retrieved from http://www.jstor.org/stable/188266

Maudlin, T. (2018). Ontological clarity via canonical presentation: Electromagnetism and the aharonov–bohm effect. *Entropy*, *20*(6). Retrieved from https://www.mdpi.com/1099-4300/20/6/465 doi: 10.3390/e20060465

Mercati, F. (2017). *Shape Dynamics: Relativity and Relationalism.* Oxford University Press.

Mitter, P. K., & Viallet, C. M. (1981). On the Bundle of Connections and the Gauge Orbit Manifold in Yang-Mills Theory. *Commun. Math. Phys.*, *79*, 457. doi: 10.1007/BF01209307

Mulder, R. (2021). Gauge-Underdetermination and Shades of Locality in the Aharonov–Bohm Effect. *Foundations of Physics*.

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.

Palais, R. (1961). On the existence of slices for the actions of non-compact groups. *Ann. of Math.*, *73*, 295-322.

Penrose, R. (1996, May). On Gravity's role in Quantum State Reduction. *General Relativity and Gravitation*, *28*(5), 581-600. doi: 10.1007/BF02105068

Pooley, O. (2022). *The Reality of Spacetime.* Unpublished manuscript.

Pooley, O. (2023). The Hole Argument. In E. Knox & A. Wilson (Eds.), *The routledge companion to the philosophy of physics.* Routledge. Retrieved from http://philsci-archive.pitt.edu/18142/

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

Pretorius, F. (2005, sep). Evolution of binary black-hole spacetimes. *Physical Review Letters*, *95*(12). Retrieved from https://doi.org/10.1103%2Fphysrevlett.95.121101 doi: 10.1103/physrevlett.95.121101

Ramirez, S., & Teh, N. (2019). Abandoning galileo's ship: The quest for non-relational empirical signicance. *preprint*.

Riello, A. (2021). Symplectic reduction of Yang-Mills theory with boundaries: from superselection sectors to edge modes, and back. *SciPost Phys.*, *10*, 125. Retrieved from https://scipost.org/10.21468/SciPostPhys.10.6.125 doi: 10.21468/SciPostPhys.10.6.125

Rovelli, C. (2002, January). GPS observables in general relativity. *Physical Review D*, *65*(4). Retrieved from http://dx.doi.org/10.1103/PhysRevD.65.044017 doi: 10.1103/physrevd.65.044017

Rovelli, C. (2007). *Quantum Gravity.* Cambridge University Press.

Rovelli, C. (2014). Why Gauge? *Found. Phys.*, *44*(1), 91-104. doi: 10.1007/s10701-013-9768-7

Singer, I. M. (1978). Some Remarks on the Gribov Ambiguity. *Commun. Math. Phys.*, *60*, 7-12. doi: 10.1007/BF01609471

Thiemann, T. (2003). Lectures on loop quantum gravity. *Lecture Notes in Physics*, 41–135. Retrieved from http://dx.doi.org/10.1007/978-3-540-45230-0_3 doi: 10.1007/978-3-540-45230-0_3

Tong, D. (2018). *Lecture notes on gauge theory.* (Available online at http://www.damtp.cam.ac.uk/user/tong/gaugetheory.html)

Walker, M. L., & Duplij, S. (2015). Cho-Duan-Ge decomposition of QCD in the constraintless Clairaut-type formalism. *Phys. Rev. D*, *91*(6), 064022. doi: 10.1103/PhysRevD.91.064022

Wallace, D. (2019). Observability, redundancy and modality for dynamical symmetry transformations. *Forthcoming*. Retrieved from http://philsci-archive.pitt.edu/18813/ (Revised 3/2021 to correct a few typos and add a section on Noether's Theorem.)

Wallace, D. (2024). *Gauge invariance through gauge fixing.*

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

Wilkins, D. R. (1989). Slice theorems in gauge theory. *Proceedings of the Royal Irish Academy. Section A: Mathematical and Physical Sciences*, *89A*(1), 13–34. Retrieved from http://www.jstor.org/stable/20489307

Witten, E. (2023). *A background independent algebra in quantum gravity.*

York, J. W. (1971). Gravitational degrees of freedom and the initial-value problem. *Phys. Rev. Lett.*, *26*, 1656–1658. doi: 10.1103/PhysRevLett.26.1656