

The Metaphysics of Fibre Bundles

Forthcoming in *Studies in History and Philosophy of Science*

Caspar Jacobs*

December 1, 2022

Abstract

Recently, Dewar (2019) has suggested that one can apply the strategy of ‘sophistication’—as exemplified by sophisticated substantivalism as a response to the diffeomorphism invariance of General Relativity—to gauge theories such as electrodynamics. This requires a shift to the formalism of fibre bundles. In this paper, I develop and defend this suggestion. Where my approach differs from previous discussions is that I focus on the *metaphysical* picture underlying the fibre bundle formalism. In particular, I aim to affirm the physical reality of gauge properties. I argue that this allows for a local and separable explanation of the Aharonov-Bohm effect. Its puzzling features are explained by a form of holism inherent to fibre bundles.

Acknowledgements: I would like to thank Jeremy Butterfield, Adam Caulton, Neil Dewar, Marian Gilton, Henrique Gomes, Bixin Guo, Tushar Menon, Ruward Mulder, Oliver Pooley and James Read for their valuable comments on earlier drafts of this paper. This paper was written as part of a research project funded by the British Society for the Philosophy of Science.

1 Introduction

Gauge symmetries—by which I mean *local* symmetries that act on the theory’s *internal* degrees of freedom¹—pose many of the same problems that

*Merton College, University of Oxford, Merton Street, Oxford, OX1 4JD, caspar.jacobs@merton

¹ The terminology here is particularly confusing: among philosophers, it is not uncommon to use the term ‘gauge symmetry’ for any transformation that relates physically equivalent states, whereas physicists reserve the term for local symmetries; see Weatherall (2016b).

other symmetries pose in similar contexts. Firstly, it seems that the gauge transformations of electrodynamics relate empirically equivalent yet physically distinct states of affairs, implying an underdetermination of empirical facts by the theory’s dynamics. Secondly, the fact that gauge symmetries are local means that one can construct analogues of the infamous Hole Argument: transformations that act trivially before some time t , but non-trivially thereafter. This implies a failure of indeterminism.

Broadly, there are three strategies for interpreting symmetry-related models (SRMs). The first is literalism: SRMs represent physically distinct states of affairs. In order to avoid indeterminism, however, one has to supplement literalism with some further claim. For example, one could claim that only one out of an equivalence class of SRMs represents a *possible* state of affairs. This claim in effect elevates a particular gauge fixing condition to become an additional law. But such strategies face various problems—most importantly, that the particular choice of ‘gauge law’ is essentially arbitrary. Therefore, I will not further consider literalism here. Instead, I will consider the interpretation of gauge symmetries from the perspective of the debate between *reduction* and *sophistication*.² Recall that reduction aims at a *reduced* theory formulated solely in terms of invariant quantities, such that each model of the reduced theory uniquely corresponds to an equivalence class of SRMs of the old theory. Sophistication, on the other hand, aims at a *restructured* theory such that the new theory’s SRMs are isomorphic. Since isomorphic models are structurally equivalent, the latter method allows one to interpret SRMs *anti-quidditistically* as physically equivalent (more on this below). I will defend sophistication as the correct interpretation of gauge theories. Specifically, I will argue that sophistication makes most sense of physicists’ use of the *fibre bundle formalism* in modern formulations of gauge theories, which for a reductionist seems to possess excess structure.

The suggestion that one can illuminate electrodynamics with fibre bundles is not novel; I discuss several related proposals in §4. But my account differs from previous ones in two important ways. Firstly, I intend to dissolve the issue of underdetermination: once gauge-related models are identified, the underdetermination disappears. This is not the case for the fibre bundle-based accounts of Leeds (1999), Nounou (2003) or Maudlin (2007), to name a few. Secondly, my solution involves an appeal to anti-quidditism,

Furthermore, there is some debate over whether external symmetries count as gauge; see Wallace (2015) and Dewar (2020).

² Dewar (2019) was the first to explicitly draw this distinction; see also Martens and Read (2020) and Jacobs (2021).

following a suggestion by Dewar (2019). This requires a closer study of the metaphysics of the fibre bundle formalism, which is absent from more mathematical treatments such as Weatherall's (2016a). For example, in §6.2 a novel distinction between 'deflationary' and 'inflationary' principal bundle realism is drawn. The latter results in a construction called the 'bundle of connections'.

The plan for the remainder is as follows. In §2, I use the Aharonov-Bohm effect to illustrate how symmetries pose a problem for the interpretation of gauge theories. In §3, I survey and criticise various reductionist responses to this problem. In §4, I introduce the fibre bundle framework. In §5, I discuss gauge symmetries in this formalism. In §6, I consider the metaphysics of fibre bundles on a sophisticated account. Specifically, I argue that anti-quidditism about gauge quantities allows us to interpret SRMs as physically equivalent. In §7, I discuss whether my account is *local* in various senses. I argue that the fibre bundle account is both local and separable, but that it also implies a form of holism. §8 concludes.

2 The Aharonov-Bohm Effect

The Aharonov-Bohm effect is essentially a modified double-slit experiment in which a solenoid is placed between the plate and the screen. We assume that the solenoid is impenetrable. The Aharonov-Bohm effect then refers to the fact that the interference pattern on the screen changes when we let a current run through the solenoid. This is the case despite the fact that the electromagnetic field vanishes outside the solenoid. Hence, we cannot simply understand the effect as a result of the force field acting locally on the matter field. This led Aharonov and Bohm to posit the *four-potential*, which does not vanish outside the solenoid, as causally responsible for the effect (Aharonov and Bohm, 1959).

In more detail,³ recall that the electromagnetic field tensor $F_{\mu\nu}$ can be expressed in terms of the electromagnetic four-potential A_μ :

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu \quad (1)$$

In terms of $F_{\mu\nu}$ and A_μ , the Lagrangian for a single particle with wave function ϕ in an electromagnetic field $F_{\mu\nu}$ is

$$\mathcal{L} = (D_\mu \phi)(D^\mu \phi)^* - m^2 |\phi|^2 - \frac{1}{4} F_{\mu\nu} F^{\mu\nu} \quad (2)$$

³ The exposition draws on Healey (2007) and Brown (2016).

where $D_\mu := \partial_\mu + iqA_\mu$ and q is a scalar quantity which denotes the field's charge. Here, ϕ is a classical matter field. The Lagrangian is invariant under the following *gauge transformation*:

$$\begin{aligned} A_\mu &\rightarrow A_\mu - \frac{1}{q}\partial_\mu\alpha(x) \\ \phi &\rightarrow e^{iq\alpha(x)}\phi \end{aligned} \tag{3}$$

where $\alpha(x)$ is a function of the spacetime coordinates x .

For the Aharonov-Bohm effect, we consider the matter field

$$\phi(x) = \phi_I(x) + \phi_{II}(x), \tag{4}$$

where ϕ_I and ϕ_{II} are the components of the field that pass through the left and right slit respectively. Let P denote the source of the field, and Q an arbitrary point on the screen. Then $\phi(Q)$ transforms under (3) as follows:

$$\phi(Q) \rightarrow \phi_I(Q) + e^{iq\Phi}\phi_{II}(Q) \tag{5}$$

where Φ is the flux through the solenoid.

The puzzling fact is that Φ depends only on $F_{\mu\nu}$, which vanishes outside the solenoid. The phase shift of the matter field seems to causally depend on a field with which it cannot interact locally. For this reason, physicists often consider the four-potential A_μ , which does not vanish outside the solenoid, as physically real (Aharonov and Bohm, 1959; Feynman et al., 1964). But the fact that A_μ is gauge-variant is problematic for familiar reasons. Configurations of the A-field related by the gauge transformations in (3) are observationally equivalent. Therefore, reifying the four-potential implies the underdetermination of the theory's models by the empirical data. Furthermore, there exist gauge transformations that act as the identity before some time t but non-trivially thereafter (i.e. $\alpha(t, \vec{x}) \neq 0$ iff $t' > t$). The existence of such transformations seems to imply that electrodynamics is indeterministic. This form of indeterminism is particularly problematic because the difference between outcomes is unobservable, so the indeterminism does not occur at the level of observables (unlike the indeterminism of quantum mechanics). The literal approach to A_μ , according to which gauge-related models represent distinct states of affairs, thus has several undesirable features.

3 Failures of Reduction

There are various alternatives to A -realism that aim to avoid this indeterminism. The three main proposals are *F-realism*, *holonomy realism* and *field monism*.⁴ I will discuss each of these in turn. As far as I am aware, no one has yet noticed that all of these approaches are instances of *reduction*. Recall that reduction aims to avoid underdetermination by constructing a new theory in terms of invariant quantities of the old theory, such that there exists a unique correspondence between equivalence classes of SRMs of the old theory on the one hand and models of the new theory on the other. I will argue that reduction is an unsatisfactory response to the presence of gauge symmetries.

3.1 F -Realism

F -realism is the view that the electromagnetic field-tensor $F_{\mu\nu}$ is fundamental. F -realism is a form of reduction, since $F_{\mu\nu} := \partial_\mu A_\nu - \partial_\nu A_\mu$ is an invariant quantity defined in terms of A_μ ; there is a unique correspondence between models in terms of $F_{\mu\nu}$ and equivalence classes of gauge-related models in terms of A_μ (Weatherall, 2016b, 1041-42).

F -realism faces two main problems. The first is the well-known fact that an explanation of the Aharonov-Bohm effect in terms of the Faraday tensor implies a violation of the principle of Local Action (Healey, 1997). Since there is no overlap between the electromagnetic field and the matter field, the former can only act on the latter at a distance. This is universally seen as sufficient reason to reject F -realism.

The second problem is that F -realism implies a ‘cosmic coincidence’ (Dewar, 2019, 498). This follows from the fact that $F_{\mu\nu}$ is in some sense a *comparative* quantity: the derivative that occurs in its definition means that its values depend on the values of A_μ at infinitesimally close points. The coincidence is the Gauss-Faraday law $\partial_{[\mu} F_{\nu\rho]} = 0$. When $F_{\mu\nu}$ is interpreted as fundamental, this is a brute law-like fact. But when A_μ is considered fundamental, the definition of $F_{\mu\nu}$ as the exterior derivative of A_μ *entails* the Gauss-Faraday law. Put differently, for the F -realist it is an unexplained fact that $F_{\mu\nu}$ behaves just *as if* it is the exterior derivative of a four-potential,

⁴ This leaves out several other proposals, such as those of DeWitt (1962), Mattingly (2006) and Mulder (2021). Since the aim of this section is not to offer a comprehensive overview of responses to the Aharonov-Bohm effect but only to show that the most prominent instances of reduction fail, I will not discuss these here. For a response to DeWitt, see Aharonov and Bohm (1962); for a criticism of Mattingly, see Healey (2007, §4.2).

even if it is in fact a fundamental quantity. Therefore, F -realism incurs an explanatory loss in addition to a violation of Local Action.

3.2 Holonomy Realism

While F -realism has virtually no advocates, *holonomy realism* is a popular interpretation of electrodynamics (Belot, 1998; Lyre, 2004; Healey, 2007). According to holonomy realism, the so-called *holonomies* H of A_μ are fundamental:

$$H(l) = \exp(-iq \oint_l A_\mu dx^\mu) \quad (6)$$

On this picture, there is a fundamental non-localised property associated to every closed curve in spacetime. Holonomies are not composed of the field-values at each spacetime point, but attach to curves as a whole. When a matter field interacts with these holonomies, it does so ‘at once’ around a loop. Since holonomies overlap with the matter field, interactions are local.

Like F -realism, holonomy realism is an instance of reduction. The holonomies are gauge-invariant quantities (in non-Abelian theories, the *Wilson loops* are the invariants), and moreover an equivalence class of gauge-related A -fields yields a unique equivalence class of assignments of holonomy values to closed curves in spacetime up to a choice of base point for each such curve (Barrett, 1991; Rosenstock and Weatherall, 2016).⁵ Moreover, holonomies are comparative quantities. The holonomy around a closed path l is a function of the integrals of A_μ over any pair of open paths l_1, l_2 that compose l :

$$H(l) = e^{-iq \oint_l A_\mu dx^\mu} = e^{-iq \int_{l_1} A_\mu dx^\mu} e^{-iq \int_{l_2} A_\mu dx^\mu} \quad (7)$$

Thus, $H(l)$ is comparative in the sense that it is defined as a function of *pairs* of paths, just as distance is defined as a function of pairs of particles. It is for this reason that Arntzenius (2012) calls Healey’s view ‘gauge relationism’.

Holonomy realism faces three main problems. The first is that the dynamics are still expressed in terms of A_μ , not in terms of H . As far as I am aware, no one has yet succeeded in writing down a Lagrangian for scalar electrodynamics in terms of holonomies directly. But if holonomies are the fundamental quantities of nature, why is it that we cannot express the laws in terms of them?

⁵ Note the erratum to Rosenstock and Weatherall’s paper (Rosenstock and Weatherall, 2018).

The second is that an ontology of holonomies is non-separable: the intrinsic facts about a region X and the intrinsic facts about another region Y don't uniquely determine all intrinsic facts about the joint region $X \cup Y$. Myrvold (2011) calls this notion 'patchy' separability. We can easily see that patchy separability fails when we consider two partially overlapping regions X and Y close to the solenoid in the Aharonov-Bohm effect. Since X does not enclose the solenoid, the flux through its surface is zero; and likewise for Y . But now consider the union $X \cup Y$. This region *does* enclose the solenoid, so it has a non-zero holonomy value. Therefore, the intrinsic facts about X and Y fail to determine the intrinsic facts about $X \cup Y$: separability fails.

But it is debatable whether this is a real issue. Perhaps the world just *is* non-separable. Quantum entanglement already gives us reason to think this is the case.⁶ I therefore find the third issue associated with holonomy realism most serious. The issue is that holonomy realism has to postulate certain structural patterns in the instantiation of holonomies as brute facts. Specifically, the holonomies of distinct loops l_1 and l_2 must satisfy the following relation:

$$H(l_1 \circ l_2) = H(l_1) H(l_2) \quad (8)$$

Here, $l_1 \circ l_2$ is the concatenation of the two loops, that is, the result of first going around l_1 and then going around l_2 . Call this feature *composite loop multiplication* (CLM).⁷ The fact that CLM holds guarantees that we can express holonomies as exponentials of loop integrals of A_μ . If CLM fails then it is no longer the case that any assignment of holonomies uniquely defines a potential field up to gauge-equivalence. Put differently, CLM makes it look *as if* holonomies supervene on local fields. Or, as Arntzenius (2012, 195) puts it, "a fairly obvious explanation of why [CLM] hold[s] is that the map H is, roughly speaking, the integration of a connection around a loop".

In more detail, if we define $H(l) := \exp[-i/\hbar \oint_l \mathbf{A} \cdot d\mathbf{x}]$, then

⁶ Although as Maudlin (1998) argues, the entanglement here is of a different nature. In addition, Dougherty (2017) argues that an 'untruncated' version of the holonomy view *is* separable.

⁷ In the case of non-Abelian theories, the gauge-invariant quantities satisfy a more complicated set of relations called the Mandelstam identities, which seem even more conspiratorial than CLM.

$$\begin{aligned}
 H(l_1) H(l_2) &= \exp[-iq \oint_{l_1} A_\mu dx^\mu] \cdot \exp[-iq \oint_{l_2} A_\mu dx^\mu] \\
 &= \exp[-iq \oint_{l_1 \circ l_2} A_\mu dx^\mu] \\
 &= H(l_1 \circ l_2)
 \end{aligned} \tag{9}$$

Hence we can explain CLM if we posit the existence of a local four-potential. On the holonomy interpretation, on the other hand, there is apparently nothing which guarantees CLM. For example, $H(l_1)$ could have been slightly lower than it actually is. In that case, $H(l_1 \circ l_2)$ would either have been different too, or it would remain the same. In the former case CLM could easily have failed, so it seems a conspiracy that the holonomies just happen to be lined up the right way in the actual world. In the second case, on the other hand, the values of $H(l_1)$ and $H(l_2)$ are counterfactually connected, despite the fact that l_1 and l_2 only partially overlap. This counterfactual action-at-a-distance is at least as puzzling as non-separability, if not more so. In either case, CLM is a cosmic coincidence.

In response to this objection, Healey appeals to the *loop supervenience* of holonomy properties: “the holonomy properties of any loop $\otimes_i L_i$ are determined by those of any loops L_i that compose it” (Healey, 2007, 123). In other words, composite loops are not fundamental since they are composed of smaller loops. Of course, the same is true for these smaller loops, which are themselves composed of even smaller ones. There is no end to this chain of supervenience, hence no smallest fundamental unit. But let’s set aside potential worries this infinite regress may invite. The question then becomes: can loop supervenience explain CLM? The answer is ‘No’. For the sense in which smaller loops ‘compose’ larger loops is unusual. It is not the case that smaller loops constitute composite loops in the same way that the mass of a composite system is determined by the masses of the parts, for example. In the latter case, we are simply concerned with mereological composition. But loop composition does not have the correct formal properties to count as mereological fusion. For example, mereological parthood is anti-symmetric: distinct objects cannot be proper parts of each other. But we can prove that loop composition *is* symmetric. Let l_1^{-1} stand for the loop that goes around l_1 in the opposite direction; then $H(l_1 \circ l_2) H(l_1^{-1}) = H(l_2)$. If loop composition is identical to mereological fusion, then this would imply that $l_1 \circ l_2$ is part of l_2 . But since l_2 is also part of $l_1 \circ l_2$, this implies that the whole is part of a part of the whole: a contradiction. Thus, loop composition fails

to determine a unique supervenience base. The thesis of loop supervenience cannot explain composite loop multiplication.

3.3 Field Monism

Wallace (2014b) introduces an interpretation of electrodynamics which is both local and separable. Since on Wallace’s account “the electromagnetic and scalar fields cannot be thought of as separate entities”, but jointly “[represent] aspects of a single entity” (15), I will call this view *field monism*. Instead of a complex matter field, Wallace’s fundamental fields are the real scalar field $\rho = |\phi|$ and the covariant derivative field $D_\mu\theta = \partial_\mu\theta - A_\mu$, where θ is the phase of ϕ .⁸ Since both ρ and $D_\mu\theta$ are invariant quantities defined in terms of the old theory’s variant quantities, field monism presents us with another example of reduction. The joint distributions of these fields uniquely correspond to equivalence classes of gauge-related models of electrodynamics. Moreover, Wallace’s account is comparative in the same sense in which *F*-realism was, since $D_\mu\theta$ is defined in terms of the partial derivative of θ . For this reason, we would expect similar cosmic conspiracies to appear.

But the main issue with field monism is that it does not easily extend to more complex gauge theories. Wallace (2014a, 17) himself admits that this is a problem, writing that “in general, I know of no comparably simple set of local gauge-invariant quantities in the non-Abelian case that can serve as a gauge-invariant representation”. This suggests that it is no more than a fortunate accident that we can represent the simple $U(1)$ gauge theory Wallace considers in terms of a unique set of gauge-invariant local quantities.

Moreover, even more complex *Abelian* theories need not have a unique gauge-invariant representations (as Wallace also admits). Consider a pair of complex-valued fields ϕ and χ with different charges, both of which are coupled to A_μ . The Lagrangian of this system is:

$$\mathcal{L} = \mathcal{L}_1 + \mathcal{L}_2 - \frac{1}{4}F^{\mu\nu}F_{\mu\nu} \quad (10)$$

where

$$\mathcal{L}_1 = (\partial^\mu + ie_1A^\mu)\phi^*(\partial_\mu - ie_1A_\mu)\phi \quad (11)$$

$$\mathcal{L}_2 = (\partial^\mu + ie_2A^\mu)\chi^*(\partial_\mu - ie_2A_\mu)\chi \quad (12)$$

⁸ For simplicity, we will only consider cases in which the matter field does not vanish. As Wallace (2014a, §4) notes, if the matter field does vanish then field monism is non-separable too.

$\mathcal{L}_1$ is just the Lagrangian of a single complex field coupled to the gauge field. Therefore, we can follow Wallace and rewrite it in terms of the gauge-invariant quantities $\rho = |\phi|$ and $D_\mu \theta_1 = \partial_\mu \theta_1 - A_\mu$, where θ_1 is the phase of ϕ . But since we have replaced A_μ with $D_\mu \theta_1$ in $\mathcal{L}_1$, we have to make the same substitution in $\mathcal{L}_2$. This results in a gauge-invariant ontology which consists of a real-valued field ρ_1 with charge e_1 , a complex-valued field χ with charge e_2 , and a connection $D_\mu \theta_1$. Alternatively, we could have started with $\mathcal{L}_2$ and written *that* Lagrangian in terms of a real-valued field and a connection. This results in a gauge-invariant ontology which consists of a complex-valued field ϕ with charge e_1 , a real-valued field ρ_2 with charge e_2 , and a connection $D_\mu \theta_2$. Crucially, these are different ontologies. Although both theories posit a real-valued field and a complex-valued field, the charges of these fields differ. On the first interpretation the charge of the real scalar field is e_1 and that of the complex field is e_2 , while on the second interpretation this is the other way around.

Therefore, on one way of understanding Wallace’s observation with respect to the unitary gauge, it implies a form of theoretical underdetermination: the choice between these two ontologies is arbitrary. This is hardly better than the underdetermination of theory by the empirical data implied by the existence of gauge symmetries. Therefore, in these more complex scenarios Wallace’s account for finding a gauge-invariant representation is inadequate.

4 Fibre Bundle Accounts

Instead of reduction, I suggest *sophistication* as an approach to gauge symmetries. Recall that the aim of sophistication is to restructure a theory’s models such that symmetry transformations become isomorphisms. This allows us to interpret those models as physically equivalent, as I will explain in §6. In the case of local symmetries the appropriate mathematical structures are fibre bundles. In this section, I introduce the fibre bundle formalism. Because my focus is on the physical interpretation of this formalism I do not aim for a comprehensive treatment; for more details, see *inter alia* Baez and Muniain (1994), Isham (1999), Healey (2007) or Weatherall (2016a).

I am not the first to suggest that the fibre bundle formalism can aid our interpretation of electrodynamics: Leeds (1999), Nounou (2003), Maudlin (2007), Arntzenius (2012), and Weatherall (2016a) all appeal to it in one way or another. But my account differs from theirs on a few points. Leeds, Nounou and Maudlin all aim to draw metaphysical conclusions from the fi-

bre bundle formalism. Nounou, for instance, argues that topological features of the fibre bundle explain the Aharonov-Bohm effect, whereas Maudlin emphasises the consequences of a fibre bundle picture for the status of universals. However, none of their accounts address the underdetermination or indeterminism problems. Leeds explicitly acknowledges that his picture “traffics heavily in non-measurable properties and quantities” (613); Nounou similarly admits that “we also part with determinism in the sense that [...] there are infinitely many gauge fields corresponding to one electromagnetic field” (193). The aim of a sophisticated account of gauge theories is to rid ourselves of underdetermination and indeterminism by interpreting SRMs as physically equivalent. The proposal thus comes closer to those of Arntzenius and Weatherall. Weatherall, for instance, notes that gauge transformations in electromagnetism are formally similar to those in General Relativity. In the latter case we already have a deflationary interpretation of gauge symmetries (sophisticated substantivalism), so Weatherall suggests that we apply the same interpretation to electrodynamics. But Weatherall remains silent on the metaphysical questions this raises: what are the fundamental quantities of electrodynamics? I aim to present a perspicuous metaphysical picture that corresponds to the fibre bundle formalism. In particular, my account appeals to anti-quidditism, analogous to Pooley’s (2006) anti-haecceitism in the case of General Relativity.

For an intuitive idea of fibre bundles, start with the concept of a ‘value space’. The value space of the four-potential, $\mathbb{V}_A$, has the structure of a four-dimensional vector space. According to A-field realism, spacetime points are mapped into this vector space via a function $A_\mu(x) : \mathcal{M} \rightarrow \mathbb{V}_A$. The essential idea of a fibre bundle account is to assign a local ‘copy’ of $\mathbb{V}_A$ to each spacetime point. In other words, instead of a single value space $\mathbb{V}_A$, there is a distinct value space $\mathcal{V}_A^x$ for each $x \in M$. These local value spaces are called *fibres*. The collection of all fibres forms a manifold, called the *fibre bundle*, defined as follows:

Definition (Fibre Bundle). A *fibre bundle* is a triple (E, π, M) where E and M are smooth manifolds and $\pi : E \rightarrow M$ is a continuous map, with a space F (called the ‘typical fibre’) such that for each $x \in M$ there exists an open neighbourhood $U \subseteq M$ and a homeomorphism $h : U \times F \rightarrow \pi^{-1}(U)$ for which $\pi(h(x, y)) = x$.

So, a bundle consists of a pair of manifolds and a projection π that defines which points on the bundle lie ‘above’ which points on the base manifold. The bundle is called a *fibre bundle* because locally E looks like the product $U \times F$. In physical terms, we can think of F as the generic representative of a localised value space. But note that there is no canonical map from

fibres of E to F : there exist many distinct structure-preserving maps—*local trivialisations*—from $\pi^{-1}(x)$ into F .

Physical fields are represented by *sections* of fibre bundles:

Definition (Section). A *section* of a fibre bundle is a map $s : M \rightarrow E$ such that $\pi(s(x)) = x$.

Put simply, a section assigns to each point x on the manifold a unique point p of the fibre above x . If the fibre over a point represents the possible field values at that point, then a section yields a determinate field value at each point. Thus, sections replace functions $f : M \rightarrow \mathbb{V}_A$ from the manifold into some universal value space.

We now come to discuss some more specific types of fibre bundles relevant in physics. The first of these is a *principal fibre bundle*:

Definition (Principal Fibre Bundle). A *principal fibre bundle* (P, π, M) is a fibre bundle whose typical fibre is homeomorphic to a Lie group G , for which there exists a smooth and free right action of G on P such that for any local trivialisation $\xi : U \times G \rightarrow \pi^{-1}(U)$, $\xi(p, g)g' = \xi(p, gg')$.

The typical fibre G is called the *structure group* of P . We require that any local trivialisation preserves the structure of the regular group action of G on P . An intuitive way to see this is that the action of G defines the ‘difference’ between points: if $p, q \in \pi^{-1}(x)$ and $q = pg$, then g is the difference between p and q . The requirement that local trivialisations preserve this structure means that if points on the typical fibre G are some ‘distance’ g away from each other, then so are their images in ξ .

As before, there is no privileged map from the fibres of P onto G . This implies that for points of P on different fibres, it is indeterminate whether these points correspond to the *same* element of G . We can, however, endow a principal bundle with additional structure that defines a notion of ‘sameness’ across fibres. This is called a *connection*:

Definition (Connection). Let $T_p P$ denote the tangent space at a point $p \in P$. The *vertical subspace* V_p of $T_p P$ is defined as $V_p P = \{\tau \in T_p P : \pi_* \tau = 0\}$, where π_* is the pushforward of π . A *connection* ω on P then assigns a *horizontal subspace* $H_p P$ of $T_p P$ to each point $p \in P$ such that (1) $T_p P = V_p P \oplus H_p P$ and (2) $R_{g*}(H_p P) = H_{pg} P$ (where R_g is the action of G on P and R_{g*} is the pushforward of R_g).⁹

Effectively, a connection determines which of the vectors tangent to p count as horizontal, such that any vector in $T_p P$ has a decomposition in terms

⁹ Alternatively, one can define a connection algebraically as a Lie algebra-valued one-form that satisfies analogous conditions. These definitions are equivalent (Isham, 1999, 255).

of $H_p P$ and $V_p P$ (where the latter consists of all vectors that point ‘along’ the fibre). The connection is compatible with the action of G on P , so the horizontal subspace has the same orientation at each point on a fibre.

I will now connect these mathematical structures to the theory of scalar electrodynamics. First, we postulate that the connection on the principal bundle represents the electromagnetic potential. In this case, the structure group of the principal bundle is $U(1)$. We can represent the connection ω as a vector field A_μ on M relative to a choice of section. This is the familiar vector potential, also called the Yang-Mills field. But A_μ depends on an arbitrary choice of section, whereas ω is intrinsic to the bundle. Therefore, I will focus on ω as the representative of the Yang-Mills field.

We do not yet have a representation of the matter field ϕ on which A_μ acts. These matter fields live on the *associated bundle*:

Definition (Associated Bundle). Define the G -product $X \times_G Y$ of two spaces X and Y on which G has a right action as the space that is obtained from the product space $X \times Y$ by identifying points (x, y) and (x', y') iff $x' = gx$ and $y' = gy$ for some $g \in G$. Let $[x, y]$ denote the equivalence classes obtained in this way.

If P is a principal G -bundle and F is a space with a right G -action, define $P_F = P \times_G F$. The *associated F -bundle* of a principal G -bundle then is a triple (P_F, π_F, M) where $\pi_F([p, v]) = \pi(p)$. If F is a vector-representation of G , then the associated bundle is a vector bundle.

The matter field ϕ is represented by sections of the associated bundle. Since the structure group of electrodynamics is $U(1)$, the associated bundle is a vector bundle whose typical fibre is isomorphic to $\mathbb{C}$. Locally, a section of this bundle is an assignment of an element of $\mathbb{C}$ to each point in M . But again, there is no canonical map from fibres of P_F to $\mathbb{C}$, so a comparison of field-values across points depends on a conventional choice of local trivialisation.

However, the connection on P endows P_F with some further structure which defines a notion of *parallel translation*. First, define the *horizontal lift* of a curve on M to P :

Definition (Horizontal Lift to Principal Bundle). Let $\gamma(t)$ be a smooth curve on M . A curve $\gamma^\uparrow(t)$ on P is a *horizontal lift* of $\gamma(t)$ iff $\pi(\gamma^\uparrow(t)) = \gamma(t)$ and $\gamma^\uparrow(t)$ is horizontal, i.e. $\dot{\gamma}^\uparrow(t) \in H_p P$. For each point $p \in \pi^{-1}(\gamma(0))$, there is a unique horizontal lift $\gamma^\uparrow(t)$ of $\gamma(t)$ such that $\gamma^\uparrow(0) = p$.

This gives us the lift of a curve on M to the principal bundle P , but we are interested in parallel translation on the associated bundle P_F . We can use the previous definition to define horizontal lifts on P_F as follows:

Definition (Horizontal Lift to Associated Bundle). Recall that points on the associated bundle are equivalence classes $[p, v]$. Let $k_v(p) = [p, v]$, and let $\gamma^\uparrow$ be the

unique horizontal lift of a curve $\gamma(t)$ on M to the principal bundle P such that $\gamma^\uparrow(0) = p$. Then, for any point $[p, v] \in \pi_F^{-1}(\gamma(0))$, the horizontal lift of $\gamma(t)$ that passes through $[p, v]$ is the curve $\gamma_F^\uparrow(t) : k_v(\gamma^\uparrow(t)) = [\gamma^\uparrow(t), v]$.

We can then define parallel translation as a ternary relation $S(\gamma, [p, v], [p', v'])$ between a path $\gamma(t)$ on M and a pair of points $[p, v], [p', v'] \in P_F$ such that $[p, v] \in \pi_F^{-1}(\gamma(a))$ and $[p', v'] \in \pi_F^{-1}(\gamma(b))$ for some $a, b \in \gamma(t)$. Then $S(\gamma, [p, v], [p', v'])$ iff the horizontal lift $\gamma_F^\uparrow(t)$ such that $\gamma_F^\uparrow(a) = [p, v]$ is such that $\gamma_F^\uparrow(b) = [p', v']$. Intuitively, this means that when one starts at $[p, v]$ and travels along γ , one ‘ends up’ at $[p', v']$. The relation of parallel translation is *path-dependent*: it is possible that $[p, v]$ and $[p', v']$ are connected via one path, but not via another. Therefore, parallel translation does not offer a well-defined universal notion of sameness across fibres.

In the next two sections, I will show how this picture facilitates a sophisticated account of electrodynamics and the Aharonov-Bohm effect.

5 Gauge Symmetries

According to sophistication, one can interpret symmetry-related models as physically equivalent *when those models are isomorphic*. This follows from an appeal to anti-quidditism—the analogue of anti-haecceitism for determinates—which I will discuss in the next section. In this section I argue that fibre bundle theories are appropriately structured in the sense that symmetry-related models *are* isomorphic.

In order to see whether fibre bundle theories are sophisticated, we first need to know their symmetries. Earlier I remarked that the local representative A_μ of the Yang-Mills field depends on a choice of section of P . It is sometimes claimed that a gauge transformation is simply a choice of a different section from which to represent the Yang-Mills field. However, this amounts to a *passive* transformation: a different choice of section yields a different coordinatisation of ω . But we are after *active* transformations. These are induced by maps between sections, also called *vertical principal bundle automorphisms*:

Definition (Vertical Principal Bundle Automorphism). A principal bundle automorphism is a diffeomorphism $u : P \rightarrow P$ such that $u(pg) = u(p)g$. A principal bundle automorphism is *vertical* iff $\pi(u(p)) = \pi(p)$.

The vertical principal bundle automorphisms are dynamical symmetries of scalar electrodynamics. But vertical bundle automorphisms are, as the name suggests, *also* symmetries of $(P, \pi, \mathcal{M})$: such automorphisms preserve both

the action of G on P and the map π from P to $\mathcal{M}$. Moreover, the same transformations are also symmetries of the associated bundle $(P_F, \pi_F, \mathcal{M})$. This is so because any principal bundle automorphism induces an associated bundle automorphism $u_F : [p, v] \rightarrow [u(p), v]$.

Because of this the symmetry-related models of electrodynamics are isomorphic. We can see this when we consider the action of gauge symmetries on the physical quantities s and ω . The transformation of s is straightforward: if $s(x) = [p, v]$, then $u_*s(x) = [u(p), v]$. The connection, meanwhile, transforms via the *pull-back* of u to $u^*\omega$. By definition, u preserves the ‘vertical’ structure of both s and ω , namely the regular group action of G on both. This leaves us with the ‘horizontal’ structure, which is captured by their behaviour under parallel translation. Consider a path γ and values of s at points $x, y \in \gamma$. The question then is: is it the case that $S(\gamma, s(x), s(y))$ iff $S(\gamma, u_*s(x), u_*s(y))$? The answer is ‘Yes’. Let $s(x) = [p, v]$ and let $\gamma^\uparrow(t)$ be the lift of $\gamma(t)$ to P which passes through p . Suppose that $\gamma^\uparrow(y) = q$. From the definition of the horizontal lift to the associated bundle, it follows that $s(y)$ is the parallel translation of $s(x)$ iff $s(y) = [q, v]$. Consider now the gauge-related section $u_*s(x)$. From the same definition as before, it follows *mutatis mutandis* that $u_*s(y)$ is the parallel translation of $u_*s(x)$ iff $u_*s(y) = [u(q), v]$. But this is the case iff $s(y) = [q, v]$, hence $s(x)$ and $s(y)$ are parallel translated via γ iff their respective gauge-transformed values are. Therefore, the relation of parallel translation is preserved under gauge transformations. Since gauge symmetries preserve both vertical and horizontal structure, it follows that gauge-related models are indeed isomorphic.

6 The Metaphysics of Fibre Bundles

The fact that gauge theories possess local symmetries is usually seen as a defect. Consider the matter field as represented by a section $s(x)$ of the associated bundle. Generally, $s(x) \neq u_*s(x)$, where u is a vertical bundle automorphism. However, since these fields are symmetry-related, the difference between them is undetectable. The presence of symmetries thus seems to entail underdetermination. Since the symmetries of gauge theories are local, this also seems to entail that such theories are indeterministic.

But these arguments are based on the implicit assumption that $s(x)$ and $u_*s(x)$ represent *distinct* fields. This follows from a literalist reading of the fibre bundle formalism: if every point $[p, v] \in P_F$ represents a field value, then an assignment of different elements of P_F to points of M must represent a distinct field. However, this claim rests on the assumption that

any point $[p, v]$ of the associated bundle represents the *same* field value across models.¹⁰ But sophistication rejects this assumption: it claims that when symmetry-related models are isomorphic, we can interpret them as physically equivalent. In the case of field theories, this requires an appeal to *anti-quidditism*: the thesis that field values are qualitatively individuated. In brief, this means that *which* field value is instantiated at some point x depends on the structure of the field over the totality of spacetime. I will elaborate on this claims below.

In this section, I discuss the matter field and the Yang-Mills field in turn. I will argue that both fields are amenable to sophistication. This will yield the result that gauge-related models of classical field theories are physically equivalent. In this way, sophistication avoids both underdetermination and indeterminism.

6.1 Matter Fields

The analysis of matter fields consists of two steps. The first step is to reify the associated bundle; the second is to identify field values with structural positions within this bundle: anti-quidditism. I will discuss each step in turn.

Firstly, I propose *associated bundle Platonism*: elements of the associated bundle P_F represents physically real entities, namely the values of the relevant matter field. In the case of scalar electrodynamics, these are the values of the matter field ϕ . The sections of the associated bundle then represent physical fields. Thus, spacetime points instantiate field values in the same way that particles instantiate masses, except that in the fibre bundle formalism each spacetime point carries its own set of field values. This implies that field values at different points are numerically distinct, so distinct points cannot possess the same field value. The structure of the associated bundle encodes the relations between field values. For instance, the ternary relation of parallel translation is part of the horizontal structure: there is a physical fact of the matter as to whether the field values instantiated at distinct points are related to each other via parallel translation along some path $\gamma(t)$.¹¹

The above claims are meant literally: field values exist and stand in second-order relations. We can contrast this with the views of Maudlin (2007) and Arntzenius (2012). For Maudlin, field values are neither universals nor tropes; he rejects the existence of field properties entirely. Maudlin

¹⁰ Jacobs (2021) calls this the ‘Value-Magnitude Link’.

¹¹ In this way, the present view resembles the one found in Leeds (1999).

argues that properties ought to induce a notion of *similarity* between objects. For example, tomatoes and strawberries are similar in that both are red. But since the fibre bundle picture implies that the field values at each point are *sui generis*, there simply is no (path-independent) notion of similarity for gauge fields. Maudlin concludes that the category of properties is superseded by that of fibre bundles. The view I call associated bundle Platonism is dismissed in a footnote: “One could suggest that there are still [gauge] properties, but that every point in the spacetime has its own set of properties, which cannot be instantiated at any other point. But since such point-confined properties could not underwrite any notion of similarity or dissimilarity [...], it is hard to see what would be gained by adopting the locution.” (Maudlin, 2007, fn. 9)

Maudlin focuses too much on *identity* here. If distinct objects possess the very same property, then this is one way in which they are similar. But consider two objects, one of which is 1 kg and the other is 2 kg. These objects do not possess the same mass value, yet they are similar in that both objects are massive. Furthermore, it is clear that both objects are more similar to each other in this respect than to some third object whose mass is 100 kg. The latter claim follows from the fact that mass value space has a certain structure for which mass ratios are well-defined: the closer a mass ratio is to 1, the more similar the objects are with respect to their masses. But just like mass value space, the associated bundle which represents the matter field possesses a highly non-trivial structure—for instance, the relation of parallel transport defined above. Although the fibre bundle picture does not allow us to say whether distinct points possess the *same* field value (or, rather, it unequivocally says that distinct points *cannot* possess the same field value), this does not mean that we cannot say anything of interest about the field properties at distinct points. In particular, just like there is a well-defined difference between mass values, the group structure of the principal fibre bundle defines a notion of closeness between field values within the same fibre. For example, let γ and γ' represent distinct paths from x to y ; and let $\gamma_F^\uparrow$ and $\gamma'^\uparrow_F$ represent their horizontal lifts onto P_F such that $\gamma_F^\uparrow(x) = \gamma'^\uparrow_F(x)$. Then $\gamma_F^\uparrow(y)$ and $\gamma'^\uparrow_F(y)$ lie on the same fibre, and hence one can meaningfully say which of these is closer to $s(y)$, the value of the matter field actually instantiated at y . Maudlin’s view simply ignores such rich structural relations between localised field values.

Arntzenius (2012) presents a different interpretation of fibre bundles, which Wolff (2020) calls *locationism*. The core idea of this view is that field values are not Platonic universals, but of the same metaphysical kind

as spacetime points. In addition to their usual spatiotemporal location, particles also have a location in ‘field value space’. The advantage of this view is that it satisfies Occam’s razor insofar as it posits fewer *kinds* of entities. But the problem with fibre bundle locationism is that it fails in the context of field theories. In field theories, spacetime points themselves possess field values. Formally, the field distribution is represented by a section: a function from the manifold into the bundle. On Arntzenius’ view, the section determines the location of each spacetime point within the fibre above that point. But while we can make sense of the idea that some discrete object has a location both in spacetime and in some distinct value space, it does not make much sense to say that spacetime points themselves are located in a value space: spacetime points *are* locations. I therefore prefer Platonism over locationism.

The second step of sophistication is to stipulate certain cross-world identity conditions for field values. In particular, we stipulate that field values are qualitatively individuated: they are nodes in a web of relations. This amounts to a rejection of *quidditism*. Black (2000) defined quidditism in terms of determinable quantities: anti-quidditism then implies that there are no distinct world in which (for instance) mass and charge are swapped. But I will use the term in a slightly different sense, namely to cover the determinate *values* of physical quantities. The idea is that field values have no primitive identities, but are qualitatively identified via their pattern of instantiation. Dewar formulates the claim as follows: “We should be anti-quidditists, and deny that physical properties are modally robust. We should not believe that there are worlds that instantiate the same structure in their laws, and differ only over which properties play which nomological roles” (Dewar, 2019, 505). The differences between isomorphic SRMs are representationally irrelevant, since they concern quidditistic facts about *which* particular field values are instantiated. But because the structural relations between field values are the same in isomorphic models, anti-quidditism implies that the *same* values of the matter field are instantiated across those models.

It is important to disambiguate this claim. On the one hand, a gauge transformation changes which element of the fibre a spacetime point is mapped into by some section s . In that sense different values of the matter field are instantiated across gauge-related models. On the other hand, it is a consequence of anti-quidditism that these different points in the fibre represent the same *physical* field value across gauge-related models. This is so because the transformations in question preserve the structural patterns between points of the bundle. The relevant structural patterns here are

just those codified in the vertical and horizontal structure of the fibre bundle. The fact that gauge-related models are isomorphic means that gauge transformations of field values preserve these structural patterns. From anti-quidditism, it thus follows that a gauge transformation does not affect which physical values of the matter field are instantiated at any spacetime point.

For a concrete example, consider a pair of gauge-related sections s and s' . On a literalist view, these sections represent distinct field configurations. But from the discussion of §5 it follows that s and s' are structurally equivalent: for any point $x \in M$, the values $s(x)$ and $s'(x)$ instantiate the same qualitative pattern. For example, the relations of parallel translation between field values at distinct points is preserved under gauge transformations: if $s(x)$ and $s(y)$ are related by parallel translation via some path γ , then so are $s'(x)$ and $s'(y)$. Therefore, $s(x)$ and $s'(x)$ are merely different representations of the same field. Notice that the claim here is *not* that either a section or the connection are *invariant* under gauge transformations. This is not the case: a gauge transformation changes which direction on the bundle counts as ‘horizontal’. But the bundle contains no background structure from which one can discern this difference, which means that the transformed connection is structurally equivalent to the untransformed one.

Consequently, gauge-related models represent the same physical possibility. On a sophisticated interpretation the theory is neither underdetermined nor indeterministic.

6.2 Yang-Mills Fields

Since the Yang-Mills field is represented by a connection on the principal bundle it is tempting to adopt *principal bundle Platonism* in analogy with associated bundle Platonism. But there is an important disanalogy between the matter field and the Yang-Mills field: while the former is represented by a *section* of the associated bundle, the latter is represented by a *connection* on the principal bundle. While we can easily interpret a section as an assignment of field values to spacetime points, the same is not the case for the connection. The connection specifies relations *between* points of the principal bundle—but what do the points of the principal bundle themselves represent?¹²

I will discuss two possible answers to this question. The first is a *deflationary approach*: neither the principal bundle nor the connection on its own represent anything physical. Rather, it is the induced connection on

¹² Dewar (2019, fn. 42) also mentions this conundrum.

the associated bundle that represents the Yang-Mills field. This approach has difficulties in accounting for distinct matter fields coupled to the same Yang-Mills field. The *inflationary approach*, on the other hand, reifies not the principal bundle but the so-called ‘bundle of connections’. The inflationary approach is preferable because it can explain the way in which distinct matter fields couple to the same Yang-Mills field.

6.2.1 Deflationary Bundle Realism

Recall from §4 that the connection on the principal bundle P induces a connection on the associated bundle P_F . According to the deflationary approach the latter represents the Yang-Mills field, which is therefore not really a *field*. Rather, the connection specifies relations *between* field values. The Yang-Mills field is thus similar to velocity, in the sense that velocities also supervene on relations between nearby points on a curve. Wallace (2014a) points out that this yields an essentially dualistic ontology of local field values on the one hand and infinitesimal field relations on the other.

On this picture, the principal bundle is a mathematical abstraction. Weatherall (2016a) defends this view. Weatherall notes that just as it is possible to define an associated bundle from a principal bundle, so one can do the reverse. Furthermore, the principal bundle is mathematically similar to the bundle of frames in General Relativity, whose role it is to coordinatise spacetime. On this view, it is only when we consider more than one field that the principal bundle becomes relevant. For if distinct matter fields couple to the same Yang-Mills field, it is useful to represent the latter ‘by itself’ on a principal bundle. The claim that both matter fields couple to the same Yang-Mills field then translates into the fact that both vector bundles are associated to the same principal bundle.

But it is a problem for this approach that the two fields survey the same connection as a matter of brute fact. There really are two connections: one defined over the first associated bundle, and one defined over the second. These connections are the same only in the sense that we can represent both with the same connection on a single principal bundle. But on the deflationary approach there is no independent Yang-Mills field that the associated bundle connections supervene on. This makes it seem somewhat mysterious that these connections are equivalent. The coordination between associated bundles begs for a ‘common cause’ in the form of an independently existing Yang-Mills field.¹³ I will therefore consider an alternative view which can

¹³ There is a similarity here to the debate between the dynamical and geometrical approach

explain the coincidence of distinct connections in terms of a physically real Yang-Mills field which interacts with both matter fields.

6.2.2 Inflationary Bundle Realism

Recall that a connection defines a horizontal subspace at each point $p \in P$. The function of the connection is to determine which direction counts as horizontal. This strongly suggest that the ‘values’ of the Yang-Mills field are not elements of P , but of the tangent bundle TP . Specifically, at each point p the value of the Yang-Mills field consists of a subspace H_pP of T_pP which counts as horizontal.

Yet this picture is redundant, in two ways. Firstly, the connection assigns an entire subspace of T_pP to points of M , rather than a single vector of it. In this sense, the connection does not yield a unique value of the Yang-Mills field. Secondly, the connection assigns a subspace of T_pP for each point $p \in \pi^{-1}(x)$. This overdescribes the Yang-Mills field, since the connection at one point on a fibre uniquely determines the connection elsewhere on the same fibre via the condition that the connection is compatible with the action of G on P (i.e. $R^{g*}(H_pP) = H_{pg}P$).

We can remove this excess structure via a construction called the *bundle of connections*:¹⁴

Definition (Bundle of Connections). Let (P, π, M) be a principal bundle with structure group G ; TP is its tangent bundle. Let TP/G denote the quotient of TP by G . Then $(TP/G, d\pi, TM)$ is a fibre bundle over TM (the tangent bundle of the spacetime manifold M), called the *bundle of connections*. There is a one-to-one correspondence between connections on P and linear sections Γ of the bundle of connections.

The bundle of connections involves two innovations. Firstly, the bundle projects onto the tangent space TM of the manifold M . Instead of a horizontal subspace for each point, the bundle of connections assigns to each vector $v \in TM$ a subspace of TP such that each vector in this subspace points in the same direction as (the lift of) v . This overcomes the first redundancy. Secondly, in the definition of the bundle of connections we have quotiented by the action of G , such that each vector in TM is attached to just one copy of the tangent bundle T_pP at some point p . This overcomes the

to spacetime. On the former approach, it is a brute fact that all matter fields have the same symmetries. But if we assume that matter surveys the structure of spacetime, then the latter can explain this coincidence of symmetries.

¹⁴ For more on this construction, see Kobayashi (1957, Ch. 4). Gomes (2021) calls this object the Atiyah-Lie manifold.

second redundancy. The result is a construction which assigns a vector space $T_p P/G$ to each element of TM (or, alternatively, to each pair of infinitesimally close spacetime points). The Yang-Mills field is then represented by a section of the bundle of connections, which assigns to each tangent vector of the spacetime manifold a unique horizontal vector of the principal bundle in the same direction. This might be taken to vindicate Wallace’s claim that the ontology of gauge theories is dualistic: matter fields inhere in points, whereas gauge fields can be understood as relations between points.¹⁵

The advantage of the inflationary view is that the Yang-Mills field is an independent entity which we can use to explain the fact that distinct matter fields behave in the same way under parallel translation. The explanation is simply that both matter fields survey the same connection, here understood as a section of the bundle of connections.

Finally, we can apply anti-quidditism to the Yang-Mills field, just as we did for the matter field. In this case, a gauge transformation acts directly on a section of the bundle of connections. But once more, such transformations preserve all structural patterns. Therefore, the image of ω after a gauge transformation is really the same connection differently represented on the bundle of connections. The conclusion is the same: gauge-related models represent the same physical possibility.

7 Locality, Separability, Holism

The sophisticated interpretation of the fibre bundle formalism faces no underdetermination. But an equally puzzling feature of the Aharonov-Bohm effect is its apparent non-locality. Is this also the case for fibre bundle realism? In this section, I will argue that fibre bundle realism is both local and separable. But there is another—less problematic—sense in which the fibre bundle account is *holistic*, which explains the non-local nature of the Aharonov-Bohm effect.

7.1 Local Action

According to literal-minded A -field realism, the four-potential field acts locally on the matter field, shifting its phase as the field propagates across regions within which A_μ is non-trivial. But this account assumes that field values are comparable across points, such that A_μ has an unequivocal in-

¹⁵ But note that since TM is itself a bundle over M , we can always re-interpret the bundle of connections as another bundle over M .

fluence on ϕ . This is no longer true on a fibre bundle account, since each spacetime point now has its own set of field values. The idea here then is that the connection determines the evolution of the matter field. Locally, the connection defines a notion of ‘sameness’ across bundles: if q is the parallel translation of p over a path γ , then p and q lie in the same horizontal plane. The dynamics of classical field theories tell us that matter fields propagate along this horizontal direction. The account is thus fully local: it is only the connection at the location of the matter field which determines the evolution on the fibre bundle over time.

In the Aharonov-Bohm effect, the connection is such that the parallel translation of the field along a path through the left slit and one along a path through the right shift differ. When both halves of the matter field meet at Q , they find themselves at different points on the fibre above Q : this is the phase difference that causes the shift in the interference pattern. The degree to which parallel translation around closed curves is curved is called the *curvature*, characterised by $F_{\mu\nu}$. This is indeed a global feature of the bundle, but on our account it supervenes on local values of the connection.

7.2 Separability

Is fibre bundle realism also *separable*? Dewar (2019, fn. 56) alleges that the connection on the principal bundle is not; see also Lyre (2004) and Martens and Read (2020). On the contrary, I believe that the connection *is* separable. Of course, the connection connects infinitesimally close points, so in that sense it is not completely local. But such violations of separability are not particularly worrisome, and this is clearly not the sense of non-separability that Dewar has in mind. There is little reason to believe that even classical theories are truly local in this sense, as Butterfield (2006) has argued. Instead, consider *regions* rather than points. Specifically, consider distinct regions U and V which partially overlap. The connection on U determines the horizontal lift of all paths on U , and the connection on V determines the horizontal lifts of all paths on V . But since paths on $U \cup V$ are composed of paths on U and paths on V , their horizontal lifts are now fully determined as well. It follows that the fibre bundle picture *is* separable.

In personal correspondence, Dewar has clarified his claim as follows. Consider regions U and V which individually do not surround the solenoid, but whose union $U \cup V$ does. The connections on U and V individually are locally isomorphic to a connection which is zero everywhere within these regions, even though the connection on $U \cup V$ is not. Therefore, the connections on U and V *up to gauge transformations* are insufficient to determine

the connection on $U \cup V$. But this shouldn't come as a surprise, since the gauge-invariant content of a connection over a region U just consists of the holonomies, which we already saw are non-separable. The fibre bundle realist can resist this conclusion.¹⁶

The problem with Dewar's argument is that it assumes that sophistication considers symmetry-related states of *subsystems* of the universe as equivalent. But sophistication (as I have presented it) is only interested in universe states. After all, the issue of indeterminism only arises when we consider universe symmetries. In the case of gauge theories, this is obscured by the fact that subsystem symmetries which vanish to the identity *are* universe symmetries. But a discussion of sophistication in the context of classical mechanics clarifies this issue. Consider a pair of (dynamically isolated) classical systems; for instance, Rovelli's (2014) fleets of spaceships. Both of these systems are invariant under static shifts, hence the symmetry-invariant content of each system consists only of distance relations. However, the distances between ships of the first fleet and the distances between ships of the second fleet fail to fully determine the joint state of both fleets. After all, the latter also includes the distance between both fleets. From this reasoning, it would seem that even classical mechanics is non-separable—a conclusion which Dewar *et al.* are unlikely to defend. But the mistake in this argument is that sophistication is not committed to a symmetry-invariant ontology of distance relations. On the contrary, sophistication embraces realism about non-invariant quantities such as positions on a manifold. In the same vein, Gomes (2019) argues that 'forgetting' symmetry-variant structure of subsystems is fine when we consider such systems in themselves, but causes trouble when we consider the 'gluing' of subsystems. Instead, Gomes advocates 'external sophistication and internal reduction': in order to glue subsystems together, we need to pay attention to their symmetry-variant features.

7.3 Holism

Yet it remains the case that the Aharonov-Bohm effect is distinctly non-local in character. This is clearly brought out when we ask *where* the electron picks up a phase: where does the Aharonov-Bohm effect *come about*? As Healey points out, for any local representation of the connection and any region U which does not enclose the solenoid, there exists a gauge such

¹⁶ The same point is made by Dougherty (2017), who shows that only a 'truncated' version of the fibre bundle formalism leads to non-separability. I conjecture that Lyre and Dewar have this truncated version of the formalism in mind.

that A_μ vanishes in U . Therefore, Healey concludes, there is no region in which the matter field non-trivially interacts with the Yang-Mills field—and hence no interaction at all! But Healey’s argument is invalid. For while it is true that for any region there exists a gauge in which A_μ vanishes, there is no gauge such that A_μ vanishes everywhere. Therefore, it is a gauge-invariant fact that the (local representative of the) connection is non-trivial *somewhere*.

But there is no meaningful answer to the question *where* the matter field picks up a phase because there are no meaningful cross-point comparisons of phase, since each point of the manifold is associated with its own fibre of field values. In particular, there is no canonical map from local fibres into $U(1)$. There simply *is* no picking up of a phase. The matter field moves across the fibre as dictated by the connection. We can only compare field values at the same point, when the two branches of the matter field meet at the point Q on the screen. The holism of the Aharonov-Bohm effect consists of the fact that there is not enough physical structure to express phase differences along open paths, but only around closed ones.

The following analogy is helpful.¹⁷ Consider the Twin Paradox: one twin remains in her inertial rest frame on earth, while the other travels makes a return trip to Alpha Centauri. Since the latter measures less proper time, she is younger on return than her twin. Just as we are interested in where the phase difference occurs in the Aharonov-Bohm effect, we may wonder when the age difference between the twins comes into existence. This is easy enough to answer with respect to certain planes of simultaneity. From the perspective of the earth-twin, the rocket-twin’s clock runs slow at a constant rate. But it is often¹⁸ thought that simultaneity is *conventional*. If that is indeed the case, different simultaneity choices yield different results for the differential ageing of the twins (Debs and Redhead, 1996). For example, one can choose a convention such that the rocket-twin ages at the same rate as her earthbound-twin on the outbound leg of her trip, but then ages much more slowly on the return leg. This result is analogous to the fact that one can choose a gauge such that A_μ is zero over any open path. Indeed, Rynasiewicz (2012) has shown that choice of simultaneity convention is formally equivalent to a choice of ‘gauge’ in General Relativity. Just as there are no cross-point comparisons of field values, so there are no objective cross-point simultaneity relations.

This implies that effects such as the Twin Paradox and the Aharonov-

¹⁷ For another helpful analogy (with currency exchange rates), see Maldacena (2014).

¹⁸ But not universally—see Malament (1977).

Bohm effect are holistic in this sense: although the total effect size (the age difference or interference shift) is measurable, there is no fact of the matter as to how this effect comes about as the result of small local differences. The final age difference between the twins is not the result of many small age differences that accrue locally. Likewise, the final phase difference in the Aharonov-Bohm effect is not the sum of the phase differences over infinitesimal paths. Although puzzling, such holism is simply a consequence of the fact that value spaces are localised.¹⁹ This is not a defect in our theories, but a consequence of the novel metaphysics of fibre bundles.

8 Conclusion

I have defended a sophisticated interpretation of gauge theories in the fibre bundle formalism. The interpretation is local, separable and deterministic. I have explained the puzzling nature of the Aharonov-Bohm effect in terms of a certain form of holism which is a consequence of the fact that each point comes attached with its own set of field values. The account easily generalises to non-Abelian gauge theories, since nothing in the above depends on the fact that the $U(1)$ structure group of electrodynamics is Abelian. This makes sophistication preferable over Wallace's deflationary account.

There is also a broader lesson here, namely that symmetries are an important guide to the structure of physical quantities. This is especially clear in the fibre bundle framework, since it is essentially the local $U(1)$ symmetry-group of electrodynamics which determines the structure of the principal bundle, and hence of the associated bundle which represents matter fields. Far from a redundancy in our description of the world, symmetries are carriers of physical information. Instead of eradicating this structure from our theories, sophistication does justice to its significance.

References

- Aharonov, Y. and D. Bohm. 1959. Significance of Electromagnetic Potentials in the Quantum Theory. *Physical Review* 115(3): 485–491.
- Aharonov, Y. and D. Bohm. 1962. Remarks on the Possibility of Quantum Electrodynamics without Potentials. *Physical Review* 125(6): 2192–2193.

¹⁹ Note: I do not claim that localised value spaces are a *necessary* condition for such holism to occur. For example, one finds a similar sort of holism in the so-called 'topological' Aharonov-Bohm effect, which does not require a fibre bundle account.

- Arntzenius, F. 2012. *Space, Time, and Stuff*. Oxford, New York: Oxford University Press.
- Baez, J. and J. P. Muniain. 1994. *Gauge Fields, Knots and Gravity*. Singapore ; River Edge, N.J: World Scientific Publishing Company.
- Barrett, J. W. 1991. Holonomy and path structures in general relativity and Yang-Mills theory. *International Journal of Theoretical Physics* 30(9): 1171–1215.
- Belot, G. 1998. Understanding Electromagnetism. *The British Journal for the Philosophy of Science* 49(4): 531–555.
- Black, R. 2000. Against quidditism. *Australasian Journal of Philosophy* 78(1): 87–104.
- Brown, H. 2016. The (magnetic) Aharonov-Bohm effect [lecture notes].
- Butterfield, J. 2006. Against Pointillisme about Mechanics. *The British Journal for the Philosophy of Science* 57(4): 709–753.
- Debs, T. A. and M. L. G. Redhead. 1996. The twin “paradox” and the conventionality of simultaneity. *American Journal of Physics* 64(4): 384–392.
- Dewar, N. 2019. Sophistication about Symmetries. *The British Journal for the Philosophy of Science* 70(2): 485–521.
- Dewar, N. 2020. General-Relativistic Covariance. *Foundations of Physics* 50(4): 294–318.
- DeWitt, B. S. 1962. Quantum Theory without Electromagnetic Potentials. *Physical Review* 125(6): 2189–2191.
- Dougherty, J. 2017. Sameness and Separability in Gauge Theories. *Philosophy of Science* 84(5): 1189–1201.
- Feynman, R. P., R. B. Leighton, and M. Sands. 1964. *The Feynman Lectures on Physics*, Volume III. Reading: Addison-Wesley.
- Gomes, H. 2019. Holism as the significance of gauge symmetries. <http://philsci-archive.pitt.edu/16499/>.
- Gomes, H. 2021. Same-diff? Part I: Conceptual similarities (and one difference) between gauge transformations and diffeomorphisms. *arXiv:2110.07203 [gr-qc, physics:hep-th, physics:physics]*.

- Healey, R. 1997. Nonlocality and the Aharonov-Bohm Effect. *Philosophy of Science* 64(1): 18–41.
- Healey, R. 2007. *Gauging What’s Real: The Conceptual Foundations of Contemporary Gauge Theories*. Oxford, New York: Oxford University Press.
- Isham, C. J. 1999. *Modern Differential Geometry for Physicists: Second Edition* (2nd Revised edition edition ed.). Singapore ; River Edge, N.J: WSPC.
- Jacobs, C. 2021. Invariance or Equivalence: A Tale of Two Principles. *Synthese*.
- Kobayashi, S. 1957. Theory of connections. *Annali di Matematica Pura ed Applicata* 43(1): 119–194.
- Leeds, S. 1999. Gauges: Aharonov, Bohm, Yang, Healey. *Philosophy of Science* 66(4): 606–627.
- Lyre, H. 2004. Holism and Structuralism in U(1) Gauge Theory. <http://philsci-archive.pitt.edu/1831/>.
- Malament, D. 1977. Causal Theories of Time and the Conventionality of Simultaneity. *Noûs* 11(3): 293–300.
- Maldacena, J. 2014. The symmetry and simplicity of the laws of physics and the Higgs boson. *arXiv:1410.6753 [hep-ex, physics:hep-lat, physics:hep-ph, physics:hep-th, physics:physics]*.
- Martens, N. C. M. and J. Read. 2020. Sophistry about symmetries? *Synthese*.
- 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.
- Maudlin, T. 1998. Healey on the Aharonov-Bohm Effect. *Philosophy of Science* 65(2): 361–368.
- Maudlin, T. 2007. Suggestions from Physics for Deep Metaphysics. In *The Metaphysics Within Physics*. Oxford University Press.
- Mulder, R. A. 2021. Gauge-Underdetermination and Shades of Locality in the Aharonov–Bohm Effect. *Foundations of Physics* 51(2): 1–26.

- Myrvold, W. C. 2011. Nonseparability, Classical, and Quantum. *The British Journal for the Philosophy of Science* 62(2): 417–432.
- Nounou, A. M. 2003. A Fourth Way to the Aharonov-Bohm Effect. In *Symmetries In Physics: Philosophical Reflections*, eds. K. Brading and E. Castellani. Cambridge University Press.
- Pooley, O. 2006. Points, Particles, and Structural Realism. In *The Structural Foundations of Quantum Gravity*, eds. D. Rickles, S. French, and J. T. Saatsi, 83–120. Oxford University Press.
- Rosenstock, S. and J. O. Weatherall. 2016. A categorical equivalence between generalized holonomy maps on a connected manifold and principal connections on bundles over that manifold. *Journal of Mathematical Physics* 57(10): 102902.
- Rosenstock, S. and J. O. Weatherall. 2018. Erratum: “A categorical equivalence between generalized holonomy maps on a connected manifold and principal connections on bundles over that manifold” [J. Math. Phys. 57, 102902 (2016)]. *Journal of Mathematical Physics* 59(2): 029901.
- Rovelli, C. 2014. Why Gauge? *Foundations of Physics* 44(1): 91–104.
- Rynasiewicz, R. 2012. Simultaneity, convention, and gauge freedom. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 43(2): 90–94.
- Wallace, D. 2014a. Deflating the Aharonov-Bohm Effect. *arXiv:1407.5073 [quant-ph]*.
- Wallace, D. 2014b. Review of Interpreting Quantum Theories, by Laura Ruetsche. *The British Journal for the Philosophy of Science* 65(2): 425–428.
- Wallace, D. 2015. Fields as Bodies: A unified presentation of spacetime and internal gauge symmetry. *arXiv:1502.06539 [gr-qc, physics:hep-th]*.
- Weatherall, J. O. 2016a. Fiber bundles, Yang–Mills theory, and general relativity. *Synthese* 193(8): 2389–2425.
- Weatherall, J. O. 2016b. Understanding Gauge. *Philosophy of Science* 83(5): 1039–1049.
- Wolff, J. E. 2020. *The Metaphysics of Quantities*. Oxford, New York: Oxford University Press.