Chapter 8

Case Study II: Gauge Quantities

8.1 Introduction

The symmetries discussed in the previous chapters — Leibniz shifts, uniform mass scalings, and so on — are all examples of *global* symmetries. In this chapter I turn my attention to *local* symmetries. Global symmetries are often defined as symmetries whose action is constant, whereas local symmetries vary across spacetime. But this definition is problematic: for example, boosts of the universe are paradigmatic global symmetries, but their action varies over time. Therefore, I will adopt a more precise definition according to which local symmetries depend on an infinite number of parameters, whereas global symmetries only depend on finitely many parameters (for this reason, Gomes (2019) calls global symmetries 'rigid' and local symmetries 'malleable').[1] In particular, I will focus on local symmetries of *internal* (i.e. non-spatiotemporal) degrees of freedom, also called *gauge symmetries*.[2] The main example I will use in this chapter is the gauge freedom of scalar electrodynamics.

Gauge symmetries raise many of the same issues as the symmetries dis-

[1] For a slightly less formal characterisation, see Wallace (2002).

[2] The terminology here is particularly confusing: among philosophers, it is not uncommon to use the term 'gauge symmetry' for any symmetry that relates physically equivalent states, whereas physicists reserve the term for local symmetries; see Weatherall (2016b). Furthermore, there is some debate over whether external symmetries, such as the diffeomorphism invariance of General Relativity, count as gauge; see, for example, Wallace (2015) and Dewar (2020). I will use the term 'gauge symmetries' in the technical sense to denote local internal symmetries.

cussed before, such as the presence of undetectable quantities. The argument here is exactly analogous: since symmetry-transformations of the universe are unobservable, the existence of such symmetries seems to imply the empirical underdetermination of our models. As a consequence, the values that vary under those symmetries are unmeasurable. Furthermore, the local character of gauge symmetries entails a particularly harmful form of indeterminism, analogous to the Hole Argument in General Realtivity. Since gauge symmetries are local, there exist symmetries whose action is trivial up to some time t, but non-trivial thereafter. This means that the physical facts at and before t do not determine the facts after t: a failure of determinism. This failure seems particularly problematic because of the fact that the divergent futures are observationally equivalent.

As before, there are three broad 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, namely one that satisfies a certain gauge condition. Such a claim in effect elevates a particular gauge fixing condition to become an additional law.[3] But such a strategy has various issues — most importantly, that the particular choice of 'gauge law' is essentially arbitrary — so I will not further consider literalism here.

Instead, I focus on the choice between *reduction* and *sophistication.*[4] Recall that reduction aims at a novel *reduced* theory formulated in terms of invariant quantities, such that each model of the reduced theory uniquely correspond to an equivalence class of SRMs of the old theory. Sophistication, on the other hand, aims at a *restructured* theory such that the theory's SRMs are *isomorphic.* Since isomorphic models are structurally equivalent, the latter method allows one to interpret SRMs *anti-haecceitistically* or *anti-quidditistically* as physically equivalent. I will defend sophistication over reduction as the correct interpretation of gauge theories. Specifically, I will argue that sophistication makes most sense of physics' use of the *fibre bundle formalism* in modern formulations of gauge theories, which for a reductionist this formalism seems to possess excess structure.[5]

[3] This is similar to Maudlin's (1998) 'ONE TRUE GAUGE' approach, although Maudlin does not go so far as to consider the gauge fixing condition as a law.

[4] As far as I am aware, Dewar (2019) was the first to explicitly draw this distinction; see also Martens and Read (2020).

[5] I thus agree with Dewar's (2019) third example of sophistication. I extend on his example in this chapter by offering both a more detailed technical account of the example,

The plan is as follows. In §2, I use the Aharonov-Bohm effect to illustrate how gauge symmetries pose a problem for the interpretation of gauge theories. In §3, I survey and criticise various reductionist responses to this problem, including Healey's (2007) holonomy realism and Wallace's (2014a) deflationary account of the Aharonov-Bohm effect. In §4, I introduce the fibre bundle framework, which, I argue, is the appropriate setting for a sophisticated account of gauge symmetries. In §5, I consider gauge symmetries from the perspective of the fibre bundle picture. 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 about physical explanations. §8 concludes.

8.2 The Aharonov-Bohm Effect

The Aharonov-Bohm effect is essentially a modified double-slit experiment. Between the plate and the screen, a solenoid is placed. We assume that the solenoid is impenetrable. The Aharonov-Bohm effect refers to the fact that the interference pattern *changes* when we let a current run through the solenoid. Specifically, the positions of the interference peaks are shifted within the interference envelope. This is the case *despite* the fact that the electromagnetic field vanishes outside the solenoid, where the matter field is non-trivial. Hence, we cannot simply understand the effect as the result of that 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,[6] recall that the electromagnetic field tensor $F_{\mu\nu}$ can be expressed in terms of the electromagnetic four-potential A_μ:

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu \tag{8.1}$$

In terms of $F_{\mu\nu}$ and A_μ, the Lagrangian of scalar electrodynamics is

$$\mathcal{L} = (D_\mu \phi)(D^\mu \phi)^* - m^2|\phi|^2 - \frac{1}{4}F_{\mu\nu}F^{\mu\nu} \tag{8.2}$$

and by filling out the accompanying metaphysical picture.

[6] The exposition draws from Healey (2007) and especially Brown (2016).

where $D_\mu := \partial_\mu + iqA_\mu$. Here, ϕ denotes a classical matter field. The Lagrangian is invariant under the following *gauge transformation*, where e is the charge of the field:

$$\begin{aligned} A_\mu &\to A_\mu + \partial_\mu \alpha(x) \\ \phi &\to e^{iq\alpha(x)}\phi \end{aligned} \tag{8.3}$$

where $\alpha(x)$ is a function of the spacetime coordinates x and q is a scalar quantity that denotes the field's charge.

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

$$\psi(x) = \psi_I(x) + \psi_{II}(x), \tag{8.4}$$

where ψ_I and ψ_{II} are the components of the field that pass the left and right slit respectively. Let P denote the source of the field, and Q an arbitrary point on the screen. We can always choose a gauge such that $\alpha(P) = 0$. For a matter field in a region in which $F_{\mu\nu} = 0$, then, we can always choose a gauge such that $A_\mu = \partial_\mu \alpha$ (as long as our spacetime is simply connected) and hence write:

$$\alpha(x) = \alpha(x) - \alpha(P) = \int_P^x d\alpha(x) = \int_P^x \partial_\mu(\alpha) dx^\mu = \int_P^x A_\mu dx^\mu \tag{8.5}$$

With this result, we can re-express the transformation of ψ in (8.3) as follows:

$$\phi(x) \to e^{ie\int_P^x A_\mu dx^\mu}\phi(x) \tag{8.6}$$

When applied to the particular field of interest in (8.4), this results in:

$$\begin{aligned} \psi'(Q) &= e^{iq\int_I A^\mu dx_\mu}\psi_I(Q) + e^{iq\int_{II} A^\mu dx_\mu}\psi_{II}(Q) && (8.7) \\ &= e^{iq\int_I A^\mu dx_\mu}\big(\psi_I(Q) + e^{iq\oint A^\mu dx_\mu}\psi_{II}(Q)\big) && (8.8) \\ &= c\big(\psi_I(Q) + e^{iq\oint A^\mu dx_\mu}\psi_{II}(Q)\big) && (8.9) \end{aligned}$$

Since $c := e^{iq\int_I A^\mu dx_\mu}$ is a unit complex number, we can safely ignore it. From an application of Stokes' theorem, it follows that

$$\oint_{\partial D} A^\mu dx_\mu = \int_D F_{\mu\nu} dS \tag{8.10}$$

where D is the surface bounded by paths I and II through the left and right slit respectively. The expression in (8.8) is equal to the flux Φ through D. Therefore, our final expression for the gauge transformation of ψ is:

$$\psi(Q) \to \psi'(Q) = \psi_I(Q) + e^{iq\Phi}\psi_{II}(Q) \tag{8.11}$$

Importantly, this expression is gauge-invariant. This is puzzling, however, since Ψ depends on $F_{\mu\nu}$, which vanishes outside the solenoid. In other words, the phase shift of the matter field seems to depend on a field which with it does not locally interact. For this reason, physicists often consider the four-potential A_μ (which does not vanish outside the solenoid) as the physically real field that interacts with the wavefunction. Aharonov and Bohm, for example, wrote that

> [I]n a theory involving only local interactions [...] the potentials must, in certain cases, be considered as physically effective, even when there are no fields acting on the charged particles. (Aharonov and Bohm, 1959, 490)

Similarly, Feynman simply posits local action as a necessary condition for a field to be real:

> In our sense then, the A-field is "real." You may say: "But there *was* a magnetic field." There was, but remember our original idea—that a field is "real" if it is what must be specified *at the position* of the particle in order to get the motion. The B-field in the whisker acts at a distance. If we want to describe its influence not as action-at-a-distance, we must use the vector potential. (Feynman et al., 1964)

However, the fact that A_μ, unlike $F_{\mu\nu}$, is not gauge-invariant is problematic for now-familiar reasons. Solutions to (8.2) related by the transformations (8.3) are observationally equivalent. Therefore, reifying the four-potential implies underdetermination of the theory's models. Furthermore, there exist gauge transformations that act as the identity before some time t but non-trivially thereafter (i.e. $\alpha(x) \neq 1$ iff $x_0 > t$). 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* hold at the level of observables (unlike the indeterminism of QM). Therefore, a literalist approach to A_μ according to which gauge-related models represent physically distinct states of affairs has several undesirable features.

8.3 Problems with Reduction

There are various alternatives to A-field realism that aim to avoid this indeterminism. The three main proposals are *F-field realism*, *holonomy realism* and *field monism*. I will discuss each of these in turn.

As I will show, these accounts are all instances of *reduction*. Recall that reduction aims to avoid the underdetermination of SRMs by constructing a new theory — the reduced theory — in terms of invariant quantities of the old theory, such that there is a unique correspondence between models of the new theory and equivalence classes of symmetry-related models of the old theory. I will argue that these instances of reduction are unsatisfactory. Specifically, the No Cosmic Coincidences argument against reduction of Chapter 3 also applies to the first two approaches discussed below. The third approach, meanwhile, solves the issue for scalar electrodynamics, but does not extend to more complex gauge theories.

8.3.1 F-field realism

F-field monism is the view that the electromagnetic field-tensor $F_{\mu\nu}$ is fundamental. Historically, of course, F-field realism wasn't obtained as the reduct of A-field realism. Rather, F-field realism was the *de facto* account of electrodynamics that the Aharonov-Bohm effect put pressure on. Nevertheless, F-field 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).

There are two main issues with B-field realism. The first is the well-known fact that an explanation of the Aharonov-Bohm effect in terms of B implies a violation of the principle of Local Action. Since there is no overlap between the electromagnetic field and the wavefunction, the former can only act on the latter at a distance. This is universally seen as sufficient reason to reject B-field realism, and I concur with this view.

The second issue, less well-known but more closely related to the debate between reduction and sophistication, is that B-field realism implies a 'cosmic coincidence' (Dewar, 2019, 498). This follows from the fact that $F_{\mu\nu}$ is in some sense a *relational* quantity: it depends on the values of A at infinitesimally close points. I argued in Chapter 3 that any invariant relational quantity is involved in cosmic coincidences. In this case, the coincindence is the Gauss-Faraday law $\partial_{[\mu}F_{\nu\rho]} = 0$: expressed in terms of A_μ, this is a mathematical theorem. But if A_μ is merely a mathematic abstraction, then it is mysterious that $F_{\mu\nu}$ behaves *as if* it is the exterior derivative of the four-potential. Therefore, B-field realism incurs an *explanatory loss* in addition to its violation of Local Action.

8.3.2 Holonomy Realism

While F-field 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:

$$H(l) = \exp\left(-i \oint_l A_\mu dx^\mu\right) \tag{8.12}$$

On this picture, there is a fundamental *non-localised* property associated to every closed curve in spacetime: holonomies are not composed of field-values at individual space-time points, but attach to these curves as a whole. When a matter field interacts with these holonomies, it does so 'at once' around the loop. In this sense, interactions for holonomy realism are local: the holonomies spatiotemporally overlap with the matter field.

Like B-field realism, holonomy realism is an instance of reduction: holonomies are gauge-invariant quantities such that an assignment of holonomy values to spacetime loops uniquely corresponds to an equivalence class of gauge-related A-fields (Barrett, 1991; Rosenstock and Weatherall, 2016). Furthermore, holonomies are *relational quantities*. Specifically, the holonomy around a closed path l is a function of the path integrals of A_μ over any pair of paths l_1, l_2 that compose l. In other words,

$$H(l) = e^{-i \oint_l A_\mu dx^\mu} = e^{-i \int_{l_1} A_\mu dx^\mu} e^{-i \int_{l_2} A_\mu dx^\mu} \tag{8.13}$$

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

There are two issues with holonomy realism. The first is that an ontology of holonomies is *non-separable*: the intrinsic facts about a spacetime region $X \cup Y$ do not supervene on the intrinsic facts about X on its own and the intrinsic facts about Y on its own. We can easily see this when we consider two partially overlapping regions X and Y close to the solenoid, as in the following diagram:

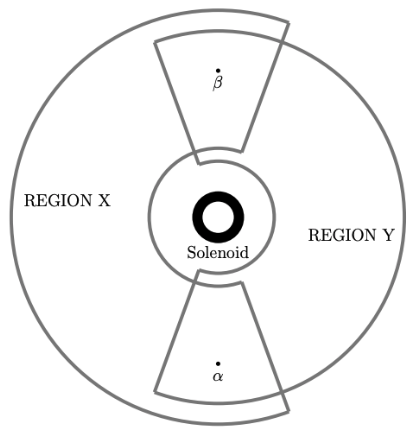

Figure 8.1: The holonomies in X and Y are zero, but there are non-trivial holonomies around the region $X \cup Y$ which encloses the solenoid. Reproduced from Wallace (2014a).

Since X does not enclose the solenoid the flux through its surface is zero, and similarly for Y. But now consider the union $X \cup Y$. This region *does* enclose the solenoid, and so has a non-zero holonomy value. Therefore, the intrinsic facts about X and Y on their own don't suffice to determine the intrinsic facts about $X \cup Y$: separability fails.

Now, it is questionable whether this is an issue. Perhaps the world just *is* non-separable — quantum entanglement already gives us some reason to think this is the case (although, as Maudlin (1998) argues, the entanglement here is of a different nature). Therefore, I find the second issue associated with holonomy realism more serious. This issue is that certain relations between holonomies are postulated as brute facts. This 'cosmic coincidence' is analogues to that faced by F-field realism.[7] Specifically, the holonomies of two distinct loops l_1 and l_2 must satisfy the following relation:

$$H(l_1 \circ l_2) = H(l_1)\, H(l_2) \tag{8.14}$$

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 (see Fig. 5.3). Call this feature *composite loop multiplication* (CLM).

CLM guarantees that we can represent holonomies as exponents of loop integrals of the gauge field, and hence express $\mathbf{B}$ in terms of $\mathbf{A}$. As Arntze-

[7] In Chapter 3, I offer a unified account of such coincidences.

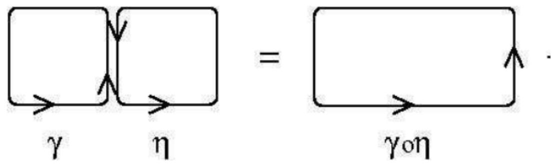

Figure 8.2: The loop $\gamma \circ \eta$ is the concatenation of γ and η. Reproduced from Asselmeyer-Maluga (2016).

nius (2012) 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 other words, if we define $H(l) = \exp\left[-i/\hbar \oint_l \mathbf{A} \cdot d\mathbf{x}\right]$, then

$$\begin{aligned} H(l_1)\, H(l_2) &= \exp\left[-i \oint_{l_1} A_\mu dx^\mu\right] \cdot \exp\left[-i \oint_{l_2} A_\mu dx^\mu\right] \\ &= \exp\left[-i \oint_{l_1 \circ l_2} A_\mu dx^\mu\right] \\ &= H(l_1 \circ l_2) \end{aligned} \tag{8.15}$$

Therefore, we can easily explain why CLM holds if we posit the existence of a local four-potential.

On the holonomy interpretation, on the other hand, there is apparently nothing that guarantees that CLM holds. 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. If the former is the case CLM could easily have failed, and it seems a conspiracy that the holonomies just happen to be related in such a way that we can represent them as loop integrals of local field values. In the second case, on the other hand, the value of $H(l_1)$ and $H(l)$ are counterfactually connected, despite the fact that l_1 and l only partially overlap. This counterfactual action-at-a-distance is at least as puzzling as non-separability, if not more so.

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 those smaller loops, which are themselves composed of even smaller ones. There is no end to this, and hence no smallest fundamental unit. The question then becomes: can loop supervenience explain the fact that CLM holds?

I believe that it cannot. Note that the sense in which smaller loops 'compose' larger loops is unusual. It is not the case that smaller loops constitue composite loops in the same way that the mass of two parts determines the mass of the whole. In the latter case, we are simply talking about mereological composition. But composite loops are not the mereological composite of smaller loops, since the small loops contain parts that don't overlap with any part of the composite loop (see Fig. 6.3). Furthermore, composition does not have the right formal properties to count as fusion. For example, mereological parthood is anti-symmetric: distinct objects cannot be proper parts of each other. On the other hand, we can prove that loop composition *is* symmetric. Let l_1^{-1} stand for the loop which 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 implies 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 means that loop supervenience is symmetric. If the supervenience of loops has no metaphysical basis in mereology, then Healey's response seems no more than a question-begging restatement of the very conspiratorial relation that we are trying to explain.

8.3.3 Field Monism

Wallace (2014b) introduces an interpretation of electrodynamics that is local in both senses mentioned above. Since Wallace's account implies that "the electromagnetic and scalar fields cannot be thought of as separate entities", but joinly "[represent] aspects of a single entity", I will call it 'field monism'. Instead of a complex matter field, Wallace's fundamental fields are the real scalar field $\rho = |\phi|$ and the covariant derivative field $D\theta = \partial_\mu \theta - A_\mu$. Field monism is another example of reduction: both ρ and $D\theta$ are invariant quantities, and joint distributions of these fields uniquely correspond to equivalence classes of symmetry-related models of electrodynamics. However, Wallace's account is not relational in the sense relevant here, since both the real scalar field ρ and the covariant derivative $D\theta$ are defined in terms of the values of ϕ and A_μ at unique points.[8] Therefore, Wallace's account does not entail any conspiracies of the sort discussed above.

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

[8] Compare this with F-field realism above. $F_{\mu\nu}$ is defined in terms of the values of A_μ at distinct (albeit infinitesimally close) points, whereas $D\theta_\mu$ is defined in terms of ϕ and A_μ evaluated at a single point.

of local gauge-invariant quantities in the non-Abelian case that can serve as a gauge-invariant representation". This suggest 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. Consider, for example, two complex-valued fields ψ and χ of different masses, both of which are coupled to A_μ. The Lagrangian of this system is:

$$\mathcal{L} = \mathcal{L}_1 + \mathcal{L}_2 - \frac{1}{4}F^{\mu\nu}F_{\mu\nu} \tag{8.16}$$

where

$$\mathcal{L}_1 = (\partial^\mu + ie_1 A^\mu)\psi^\star(\partial_\mu - ie_1 A_\mu)\psi \tag{8.17}$$

$$\mathcal{L}_2 = (\partial^\mu + ie_2 A^\mu)\chi^\star(\partial_\mu - ie_2 A_\mu)\chi \tag{8.18}$$

$\mathcal{L}_1$ 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 $D\theta_1^\mu = \partial^\mu\theta_1 - A^\mu$, where θ_1 is the phase of ψ. But since we have 'replaced' A^μ with $D\theta_1^\mu$ in $\mathcal{L}_1$, we have to make the same substitution in $\mathcal{L}_2$. This gives us an ontology consisting of a real-valued field ρ_1 with charge e_1, a complex-valued field χ with charge e_2, and a single connection $D\theta_1^\mu$. This ontology is indeed gauge-invariant.

However, we could also have started with $\mathcal{L}_2$ and written *that* Lagrangian in terms of a real-valued field and a connection. This results in an ontology consisting of a complex-valued field ψ with charge e_1, a real-valued field ρ_2 with charge e_2, and a connection $D\theta_2^\mu$. These are different ontologies! Although both consist of 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 exactly the other way around. Therefore, in this more complex case Wallace's approach implies a form of theoretical underdetermination: the choice between the two ontologies is arbitrary. The sophisticated account that I present below, on the other hand, is meant to extend to both non-Abelian theories and theories that couple distinct matter fields to the same gauge field.

8.4 Fibre Bundle Accounts

Instead of reduction, then, we may try *sophistication* as an approach to gauge symmetries.[9] Recall that the aim of sophistication is to *restructure* a theory's models such that symmetry transformations are isomorphisms. This allows for an anti-quidditist interpretation of gauge quantities, 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 framework. As my focus is on the physical interpretation of this formalism I do not aim for a comprehensive treatment of the mathematics of fibre bundles; for more details, see Baez and Muniain (1994); Isham (1999); Healey (2007); Weatherall (2016a).

I am not the first to suggest that the fibre bundle framework can aid our interpretation of electrodynamics: Leeds (1999); Nounou (2003); Maudlin (2007); Weatherall (2016a) all appeal to them one way or another. But my account differs from these on a few significant issues. Leeds, Nounou and Maudlin all aim to draw metaphysical conclusions from the fibre bundle formalism. Nounou, for example, argues that it is the topology of the fibres that explains 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 associated with gauge symmetries. Indeed, Leeds explicitly acknowledges that his picture "traffics heavily in non-measurable properties and quantities"; Nounou similarly admits that "we also part with determinism in the sense that [...] there are infinitely many gauge fields corresponding to one electromagnetic field". The aim of a sophisticated account of gauge theories, on the other hand, is to rid ourselves of underdetermination by interpreting SRMs as physically equivalent. The proposal is therefore closer Weatherall, who notes that gauge transformations in electromagnetism are similar to those in General Relativity. In the latter case, we already have a deflationary interpretation of gauge symmetries, and Weatherall suggests applying the same interpretation to the symmetries of electrodynamics. However, Weatherall remains silent on the *metaphysical* picture this implies: what are the fundamental entities of electromagnetism, and how do they interact? I aim to present a perspicuous metaphysical picture that corresponds to this formalism.

In the remainder of this section I will set out the mathematical details of the fibre bundle formalism. For an intuitive idea, start with the concept of a 'value space' discussed in earlier chapters. The value space of a quantity

[9] Dewar (2019) was the first to explicitly suggests that this is possible.

such as mass has a certain structure, and objects are mapped into this space. We could consider the value space of the four-potential, $\mathcal{V}_A$, to have the structure of a four-dimensional vector space. According to A-field literalism, then, spacetime points are mapped into this vector space by a function $A_\mu(x) : \mathcal{M} \to \mathcal{V}_A$.

Recall that in Chapter 6, I defined (internal) symmetries as transformations on a theory's models induced by bijections of their value spaces.[10] For example, if $\langle \mathcal{M}, \mathcal{V}_A, A \rangle$ is a model of the theory then a bijection ϕ is (or induces) a symmetry iff $\langle \mathcal{M}, \mathcal{V}_A, \phi_* A \rangle$ is also a model of the theory. But this definition does not yield local symmetries. The reason is that $\mathcal{V}_A$ carries no information about *where* particular field-values are instantiated. Therefore, a transformation ϕ applies everywhere equally. For concreteness, suppose that *before* a local symmetry transformation is applied the value of A is the same at x_1 and x_2. There exist local symmetries that non-trivially act on $A(x_1)$, but leave $A(x_2)$ the same. However, there is no transformation on $\mathcal{V}_A$ that implements this symmetry. If there were such a transformation, then $\phi(A(x_1)) \neq \phi(A(x_2))$. But $A(x_1) = A(x_2)$, so $\phi(A(x_1))$ *is* equal to $\phi(A(x_2))$

This may seem like a flaw of my definition of symmetries, but I believe that it is a virtue. The fact that a universal value space for A_μ does not have enough structure to account for local symmetries explains the need for a fibre bundle account. The essential idea of a fibre bundle account is to assign a local 'copy' of $\mathcal{V}_A$ to each spacetime point. In other words, instead of a single value value $\mathcal{V}_A$, there is a different 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*. As we will see, local symmetries are (vertical) automorphisms of this bundle, which I will define below.

I will now canvass this picture in some more detail. I provide definitions of all mathematical terms, but my aim is to explain the essential concepts in relatively non-mathematical terms.[11] Let's start with the concept of a fibre bundle:

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

[10] This definition is a little more loose and general than Dewar's (2020), but the definitions coincide for fibre bundle theories.

[11] For more details, see Isham (1999).

So, a bundle consists of two manifolds and a projection π that defines which points on the bundle lie 'above' which points on the spacetime manifold. It's called a *fibre* bundle because locally E looks the product $U \times F$. In physical terms, we can think of F as some value space, such as $\mathcal{V}_A$ above. The fibre $\pi^{-1}(x)$ above each point x, then, 'looks like' $\mathcal{V}_A$. But note that there is no 'canonical' map from points on E to points on F: there usually are distinct structure-preserving maps, or *local trivialisations*, from $\pi^{-1}(x)$ into F.

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

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

In other words, 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 determines a field value at each point. Sections thus replace functions $f : M \to \mathcal{V}$ from the manifold into some universal value space shared by all $x \in M$.

We now come to discuss some more specific types of fibre bundles that are relevant to 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 fibres are homeomorphic to a Lie group G, such that there exists a smooth and free right action of G on P such that for any local trivialisation $\xi : U \times G \to \pi^{-1}(U)$, $\xi(p, g)g' = \xi(p, gg')$.

In other words, a principal fibre bundle is a fibre bundle whose typical fibre is a group G, called the *structure group* of P. Furthermore, we require that any local trivialisation preserves this group-structure. An intuitive way to think of this is that the action of G defines the 'difference' between two points: if $p, q \in \pi^{-1}(x)$ and $q = pg$, then g is the difference between p and q. The requirement that a local trivialisation preserves this structure then means that if two points on the typical fibre G are some 'distance' g away from each other, then so are the images of ξ of these points.

As with bundles in general, there is no privileged map from 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_pP denote the *tangent space* at a point $p \in P$. The *vertical subspace* V_p of T_p is defined as $V_pP = \{\tau \in T_pP : \pi_*\tau = 0\}$, where π_* is the pushforward of π. A *connection* ω on P then assigns a *horizontal subspace* H_pP of T_pP to each point $p \in P$ such that (1) $T_pP \sim V_pP \oplus H_pP$ and (2) $R_{g*}(H_pP) = H_{pg}P$ (where R_g is the action of G on P).[12]

In effect, a connection determines which vectors tangent to p count as horizontal, such that any vector in T_pP has a decomposition in terms of H_pP and V_pP, where the latter consists of all vectors that point 'along' the fibre. The connection is compatible with the action of G on P, such that horizontal vectors on one fibre point in the 'same' direction across fibres.

I will now connect some of these mathematical structures to the theory of scalar electrodynamics discussed above. First, the connection on the principal bundle represents the electromagnetic potential. Relative to a choice of (local) section, we can represent the connection ω as a vector field A_μ on M. 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.

But we don't yet have a representation of the matter fields on which A acts. These matter fields, such as ϕ, 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 $\pi_F([p, v]) = p$. If F is a vector-representation of G, then the associated bundle is a vector bundle.

Matter fields are represented as sections of the associated bundle. For example, the structure group of electrodynamics is U(1), and hence the associated bundle is a vector bundle with typical fibre $\mathbb{C}$. Locally, a section is then an assignment of an element of $\mathbb{C}$ to each point in M. But note once more that there is no canonical map from P_F onto $\mathbb{C}$, and hence 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 that defines a notion of *parallel translation*. First, define the *horizontal lift* of a curve on M to P:

[12] Alternatively, one can define a connection algebraically as a Lie algebra-valued one-form that satisfies analogous conditions.

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. $\gamma^{\uparrow}(t) \in H_pP$. For each point $p \in \pi^{-1}(\gamma(0))$, there is a unique horizontal lift $\gamma^{\uparrow}(t)$ such that $\gamma^{\uparrow}(0) = p$.

This defines the horizontal lift of a curve on M into the principal bundle P, but we are interested in parallel translation on the associated bundle P_F. We can use the above 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 $k_v(p) = [p, v]$. Then, for a curve $\gamma(t)$ on M and a point $[p, v] \in \pi_F^{-1}(\gamma(o))$, the horizontal lift of $\gamma(t)$ that passes through $[p, v]$ is the curve $\gamma_F^{\uparrow}(t) : k_v(\gamma^{\uparrow}(t)) = [\gamma^{\uparrow}(t), v]$.

This definition allows us to define a notion of parallel translation on the associated bundle. We can define parallel translation as a ternary relation $S[\gamma, p, q]$ between a path $\gamma(t)$ on M and a pair of points p, q on P_F such that $p \in \pi^{-1}(\gamma(a))$ and $p \in \pi^{-1}(\gamma(b))$ for some $a, b \in \gamma(t)$. Then q is the parallel translation of p along $\gamma(t)$ iff the horizontal lift $\gamma_F^{\uparrow}(t)$ such that $\gamma_F^{\uparrow}(a) = p$ passes through q, i.e. $\gamma_F^{\uparrow}(b) = q$. Intuitively, what this means is that when one starts at p and travels along γ, one ends up at q. However, the relation of parallel translation is *path-dependent*: it is possible that p and q are equivalent via one path, but not via another. Therefore, parallel translation does not offer a well-defined *universal* notion of sameness across fibres.

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

8.5 Gauge Symmetries

According to sophistication, one can interpret one's symmetry-related models as physically equivalent *when those models are isomorphic*. Therefore, the appropriate structure of a theory is one whose isomorphisms are dynamical symmetries. In this section I argue that fibre bundle theories are appropriately structured in this sense.

Recall that we can extend Earman's (1989) symmetry principles to cover non-spatiotemporal symmetries (see Ch. 1). The idea is to define *value space symmetries* as the automorphisms of a value space $\mathcal{V}$, and *internal symmetries* as solutionhood-preserving transformations on a theory's models induced by bijections of $\mathcal{V}$. In detail, let a model $\langle M, \mathcal{V}, s\rangle$ consists of a manifold M, a value space $\mathcal{V}$ and a function $s : M \to \mathcal{V}$ that denotes some physical quantity. Furthermore, let $\mathcal{V} = \langle D, \mathcal{R}\rangle$, that is, the value space consists of a domain D and a set of relations $\mathcal{R}$ defined on D. The value space symmetries then are automorphisms ϕ of D such that $\phi(R) = R$ for all $R \in \mathcal{R}$. The dynamical symmetries, meanwhile, are transformations g of the form

$$\langle M, \mathcal{V}, s\rangle \xrightarrow{g} \langle M, \mathcal{V}, \phi_* s\rangle$$

such that $\mathcal{M}$ is a solution iff $g\mathcal{M}$ is.

The generalisations of Earman's principles then state that:

SP1 If ϕ induces an internal symmetry, it is a value space symmetry;

SP2 If ϕ is a value space symmetry, it induces an internal symmetry.

The importance of these principles is that when they are satisfied, models are symmetry-related iff they are isomorphic. For consider some relation R on $\mathcal{V}$, and a pair of field values $s(x)$ and $s(y)$. Let ϕ induce an internal symmetry. From SP1 it follows that $\phi(R) = R$, and hence $\langle s(x), s(y)\rangle \in R$ iff $\langle \phi_* s(x), \phi_* s(y)\rangle \in R$. Since this is the case for all relations $R \in \{R\}$, the SRMs are isomorphic. Conversely, suppose that some models $\langle M, \mathcal{V}, s\rangle$ and $\langle M, \mathcal{V}, \phi_* s\rangle$ are isomorphic, i.e. there exists some ϕ such that $\langle s(x), s(y)\rangle \in R$ iff $\langle \phi_* s(x), \phi_* s(y)\rangle \in R$ for all $R \in \{R\}$. This means that ϕ is a value space symmetry, and so from SP2 it follows that ϕ induces an internal symmetry. This reasoning straightforwardly generalises to polyadic relations.

Here is how all this applies to the fibre bundles. Instead of a universal value space $\mathcal{V}$ we now have a principal bundle P and an associated bundle P_F (where I now use P and P_F to denote the full bundle structure). The dynamical quantities are the section $s(x)$ that represents the matter field, and the connection ω that represents the Yang-Mills field. Therefore, our models are of the form:

$$\langle M, P, P_F, \pi, \pi_F, s(x), \omega\rangle$$

The equivalent of SP1 now is that for any model $\langle M, P, P_F, \pi, \pi_F, s, \omega\rangle$, the symmetry-related model $\langle M, P, P_F, \pi, \pi_F, \phi_* s, \phi_* \omega\rangle$ is a solution iff $\phi(P) = P$ and $\phi(P_F) = P_F$, that is, if ϕ is an automorphism of both fibre bundles. Likewise, SP2 says that if ϕ is an automorphism of P and P_F, then $\langle M, P, P_F, \pi, \pi_F, s, \omega\rangle$ is a solution of the theory iff $\langle M, P, P_F, \pi, \pi_F, \phi_* s, \phi_* \omega\rangle$ is.

In order to see whether fibre bundle theories satisfy SP1 and SP2, we need to know their (gauge) symmetries. Now, 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 this fixes a gauge, and that a gauge transformation simply is the choice of a different section with which to represent the Yang-Mills field. However, this is essentially a *passive* transformation: a different choice of section amounts to a different coordinatisation of A_μ. 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 $u : P \to P$ such that $u(pg) = u(p)g$. A principal bundle automorphism is *vertical* iff $\pi(u(p)) = \pi(p)$.

Vertical principal bundle automorphisms are symmetries of (P, π, M): such automorphisms preserve both the action of G on P and the map π from P to M. But the same transformations are also symmetries of the associated bundle (P_F, π_F, M)! This is so because any principal bundle automorphism induces an associated bundle automorphism $h_F : [p, v] \to [u(p), v]$. Therefore, gauge transformations are indeed local value-space symmetries, and hence SP1 an SP2 are satisfied.

We can also see this when we consider how gauge symmetries act on the physical quantities $s(x)$ 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 $u^*\omega$. Now, by definition u preserves all 'vertical' structure of both s and ω. This leaves us with their 'horizonal' structure, which is captured in their parallel translation. Specifically, consider a path γ and the value 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. This is easiest to see algebraically: because ω transforms via the pull-back of u, it is the case that $\omega(s(x)) = u^*\omega(u_*s(x))$, and hence the unique horizontal lift of γ through $s(x)$ is mapped onto the horizontal lift of γ through $u_*s(x)$. Therefore, the relation of parallel translation is preserved by gauge transformations. Since gauge transformations preserve all vertical and horizontal structure, it follows that gauge-related models are isomorphic.

8.6 The Metaphysics of Fibre Bundles

The fact that Yang-Mills theories have non-trivial gauge transformations is usually seen as a defect. Consider, for example, the matter field as represented by a section $s(x)$. Generally, $s(x) \neq u_*s(x)$, where u is an vertical bundle automorphism. However, since these fields are symmetry-related, there are no experiments that can measure the empirical difference between these sections. Since the symmetries of electrodynamics are local, this also seems to imply that electrodynamics is indeterministic in a way analogous to the Hole Argument.

These arguments are based on the implicit assumption that $s(x)$ and $u_* s(x)$ represent *distinct* fields. This *seems* to follow from fibre bundle realism: if every point $p \in P_F$ represents a field value, then an assignment of different elements of P_F to points on M must represent a distinct field. However, this claim rests on the assumption that any point p represents the *same* field value across models. Elsewhere, I call this assumption the Value-Magnitude Link. However, sophistication rejects this assumption. According to sophistication, field values are *qualitatively identified*. What this means is that *which* field value is instantiated at some point x depends on its structural relations to other field values at distinct locations.

I will elaborate on this view below, discussing matter fields and the Yang-Mills field in turn. The result is that, since gauge-related models are isomorphic, they represent physically equivalent states of affairs. This avoids the underdetermination and indeterminism usually associated with gauge theories.

8.6.1 Matter Fields

I propose *associated bundle Platonism*: elements of the associated bundle P_F represents physically real entities, namely the field values of the relevant matter field. In the case of semi-classical electrodynamics, these are field values. Sections of the associated bundle 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 distinct points are numerically distinct, and so distinct points simply *cannot* possess the same field value. The structure of the bundle represents the relations between field values, both vertical and horizontal; part of the horizontal structure is the ternary relation of parallel translation. In other words, there is a physical fact of the matter as to whether field values at distinct points are related to each other by parallel translation over some path γ.[13]

The above claims are meant literally: field values *exist* and bear certain

[13] In this way, the present view resembles Leeds (1999).

relations to each other. The view I propose is thus a form of Platonism. We can contrast this view with those of Maudlin (2007) and Arntzenius (2012). For Maudlin, field values are neither universals nor tropes; indeed, Maudlin rejects the existence of field properties wholesale. The reason is that properties, Maudlin argues, ought to induce a notion of *similarity* between objects. For example, tomatoes and strawberries are similar to the extent 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 sense of similarity for gauge fields. Maudlin concludes that the category of properties is superceded by that of fibre bundles. The view that I propose is dismissed in a footnote: "One could suggest that there are still color properties, but that every point in the space-time 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."

It seems to me that Maudlin's focuses too much on *identity* here. Sure, if distinct objects possess the *same* property, that's 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 (in other words, both objects instantiate different *determinates* of the same *determinable*). Furthermore, it is clear that both objects are more similar to each other than to some third object whose mass is 100 kg. This latter claim follows from the fact that mass value space has a certain structure for which mass ratios are well-defined: the closer the mass ratio of two objects is to one, the more similar they are with respect to their masses. But it is clear that the associated bundle that represents the matter field *also* has a highly non-trivial structure. The notion of parallel transport we defined above is just one instance of this. So, while the fibre bundle picture does not allow us to say whether distinct points possess the same field value — or, rather, it implies that distinct points *cannot* possess the same field value — this does not mean that we cannot say anything of interest about field properties at distinct points. Maudlin's objection to fibre bundle Platonism is unsuccessful.

Arntzenius (2012) has a different view of fibre bundles, which Wolff (2020) calls (fibre bundle) *locationism*. The idea of this view is that field values are not Platonic universals, but are metaphysically the same kind of entity as spacetime points are. In addition to their usual spatiotemporal location, then, objects also have a location in 'field value space'. The advantage of this view, or so Arntzenius argues, is that it satisfies Occam's razor insofar as posits fewer *kinds* of entities. 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, rather than particle-like objects. But while we can make sense of the idea that some discrete object has a location in both a spatiotemporal space and in some other value space, it does not make much sense to say that spacetime points are located in a value space. It seems that spacetime points simply are not the kind of things that can be located; after all, such points *are* locations. I therefore reject Arntzenius' fibre bundle locationism, and focus on the Platonist view in the below.

Fibre bundle Platonism is one half of sophistication. As seen in the previous section, the fibre bundle picture guarantees that the gauge symmetries of the theory are isomorphisms of its models. The other half is *anti-quidditism* about field-values.[14] According to anti-quidditism, there are no distinct possible worlds that agree on all structural facts, but disagree on *which* field values are instantiated. In other words, field values are *qualitatively individuated*; they are places in a relational structure. Consider two gauge-related sections s and s'. On the literalist view described above, these sections represent distinct field configurations. However, as I showed in the previous sections, there are no *structural* differences between s and s'. For any point $x \in M$, $s(x)$ and $s'(x)$ are qualitatively identical. Specifically, the relation of parallel translation is preserved by gauge transformations. According to anti-quidditism, then, $s(x)$ and $s'(x)$ represent the *same* field values, namely one with a certain qualitative profile.

Therefore, according to anti-quidditism there are no worlds in which *different* values of the matter field are instantiated that nevertheless occupy the same qualitative roles. It follows that gauge-related models represent physically equivalent states of affairs. There is thus no indeterminism in electrodynamics: the distinct futures compatible with the present state are merely different representations of the same future.

8.6.2 Yang-Mills Fields

Since the Yang-Mills field is represented as a connection on the principal bundle (which induces a notion of parallel translation on the associated bundle), it is tempting to adopt *principal bundle Platonism* in analogy with associated bundle Platonism. But there is an important disanalogy between

[14] As mentioned in a footnote above, the suggestion to apply anti-quidditism to fibre bundles originates from Dewar (2019); see also Martens and Read (2020). This move is analogous to the adoption of anti-haecceitism about spacetime points in response to the Hole Argument. For more on the latter, see Pooley (2005) and references therein.

the matter field and the Yang-Mills field: while the former are represented by *sections* of the associated bundle, the latter is represented by a *connection* on the principal bundle. And 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 in the bundle, but what does the principal bundle itself represent?[15]

I will describe two responses to this question. The first is a *deflationary approach*: neither the principal bundle nor the connection represent anything physical. Rather, it is the induced connection on the associated bundle that represents the Yang-Mills field. However, this approach has difficulties in accounting for distinct matter fields that couple to the same Yang-Mills field. The *inflationary approach*, on the other hand, reifies not the principal bundle but its tangent bundle. Yang-Mills fields then are sections of this 'bundle of connections'. The inflationary approach can explain how distinct matter fields couple to the same Yang-Mills field, and is for this reason preferable.

Deflationary Bundle Realism

Recall from §4 that the connection on P induces a connection on P_F. According to the deflationary approach, the connection on the latter represents the Yang-Mills field, which is thus not really a *field*. Rather, it specifies relations *between* field values. The Yang-Mills field is thus more like velocity, in the sense that velocities supervene on the relations between nearby spacetime points of a curve. As Wallace (2015) points out, this yields an essentially dualistic ontology of local field values on the one hand and infinitesimal field relations on the other.

The principal bundle, on this picture, is a mathematical abstraction. This is the view that Weatherall (2016a) defends. 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 is to coordinatise spacetime. The deflationary view leads us to wonder what the function of the principal bundle *is*. As Weatherall points out, 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

[15] Dewar (2019, fn. 42) also mentions this puzzle.

into the fact that both vector bundles are associated to the same principal bundle.[16]

The problem with this approach is that it is a brute fact *that* the two fields survey the same connection.[17] In other words, there really are two connections here: one defined over one associated bundle, and one defined over the other. These connections are the same in the sense that both are represented by the same connection on the corresponding principal bundle. But it is not the case that there is an independently existing Yang-Mills field that the associated bundle connections supervene on. This makes it seem somewhat mysterious that these connections *are* the same. This coincidence begs for a 'common cause' in the form of an independently existing Yang-Mills field. This objection is similar to the cosmic conspiracies of the reductionist approach considered in §8.3. I will therefore consider an alternative view which can explain the coincidence of distinct connections in terms of a physically real Yang-Mills field that is common to both.

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 what direction counts are horizontal. This strongly suggest that the 'values' of the Yang-Mills field are not elements of P, but of T_pP, the tangent bundle to P. Specifically, at each point p the 'value' of the Yang-Mills field consists of a subspace H_pP of T_pP that counts as horizontal.

As of yet, this is still too much structure, since the connection determines a horizontal subspace at *each* point of the principal bundle, and hence still lives on the bundle rather than on M. However, note that a specification of ω at all points of P overdescribes the Yang-Mills field, since the connection at one point on a fibre determines the connection elsewhere on the same fibre via the condition that $R^{g*}(H_pP) = H_{pg}P$. In other words, the connection is compatible with the action of G on P.

We can remove this superfluous structure by taking the *quotient* of TP by the action of G. This gives us the so-called *bundle of connections* (for

[16] The reason that Wallace's view struggled with such theories (§8.3.3) is now clear: to make sense of two matter fields that couple to the same gauge field, one represents the latter as a connection on distinct associated bundles of each matter field. But on Wallace's view, at least one of the two matter fields lives on a universal real vector space, rather than a fibre bundle.

[17] For a criticism of the technical aspects of Weatherall's deflationary view, see Menon (2018).

more on this construction, see Kobayaschi (1957, Ch. 4)):

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.

According to inflationary bundle realism, then, the bundle of connections is physically real, and its elements represent values of the Yang-Mills field. However, note that TP/G is defined over the tangent bundle to M, called TM. In other words, a section Γ of the bundle of connections assigns elements of that bundle to vectors rather than points. Or, equivalently, values of the Yang-Mills field are instantiated by *pairs* of infinitesimally close spacetime points. This, of course, vindicates Wallace's claim that the ontology of gauge theories is dualistic: matter fields inhere in points, whereas gauge fields are relations between points. The advantage of this inflationary view is that the Yang-Mills field is an independently existing entity which we can use to explain the fact that distinct matter fields are parallel translated in the same way.

Furthermore, we can apply anti-quidditism to the Yang-Mills field, just as we did with the matter field. For a gauge transformations maps pairs of infinitesimally close points to different vectors in the bundle of connections. But these vectors retain their qualitative characteristics. Specifically, the *curvature* is preserved by gauge transformations. Therefore, the image of ω under a gauge transformation is really the *same* connection differently represented on the principal bundle.

8.7 Locality, Separability, Holism

The sophisticated version of fibre bundle realism faces no underdetermination, since gauge-related models are physically equivalent. However, another puzzling feature of the Aharonov-Bohm effect is that it seems to imply that physics is, in some sense, non-local. Is this also the case for the fibre bundle formalism? In this section, I will argue that fibre bundle realism satisfies Local Action as well as separability. However, there is another — less problematic — sense in which the fibre bundle account is *holistic*. It is this holism that explains the non-local nature of the Aharonov-Bohm effect.

8.7.1 Local Action

In order to see whether fibre bundle realism violates Local Action, it is helpful to return to the Aharonov-Bohm effect. According to literal-minded A-field realism, there is a four-potential field that acts locally on the matter field, shifting its phase as the field propagates across regions in which A_μ is non-trivial. This account assumes that field values are comparable across points, such that A has an unequivocal influence on ϕ. This, of course, is no longer true on the fibre bundle account, since each spacetime point has its own set of field values. How, then, do we model *interaction* on the fibre account?

The idea here 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. This means that the matter field propagates *along* horizontal paths on the bundle. The matter field 'surveys' the connection at infinitesimal distances. The account is thus fully local: it is only the connection at (or infinitesimally close to) the location of the matter field that determines its evolution on the fibre bundle.

In the Aharonov-Bohm effect, the connection is such that the parallel translation of the field along a path to the left of the solenoid and the parallel translation along a path to the right of the solenoid differ. When both halves of the matter field 'meet' at Q, then, they find themselves at different points on the fibre above Q: this is just the phase difference that causes the shift in the interference pattern. The degree to which to which parallel translation around closed curves is closed is called the *curvature*, characterised by $F_{\mu\nu}$. This is indeed a global feature of the bundle[18], but it supervenes on local values of the connection.

8.7.2 Separability

Is fibre bundle realism also *separable*? Dewar (2019, fn. 56) alleges that the connection on the principal bundle is not, and Martens and Read (2020) concur. But I believe that on a sophisticated account the connection *is* separable. The connection on the principal bundle P connects infinitesimally close points. In that sense, it is not *completely* local. But such violations of separability don't seem particularly worrisome, and this is clearly not the sense of non-separability that Dewar intends. Indeed, there is little reason to believe that even classical theories are truly local in this sense,

[18] Or, on Nounou's (2003) account, a *topological* feature.

as Butterfield (2006) argues. Instead, then, consider *regions* rather than points. Specifically, consider two regions U and V that 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, surely *their* horizontal lift is now fully determined?

In personal correspondence, Dewar has clarified his claim as follows. Consider regions U and V that individually don't surround the solenoid, but their union $U \cup V$ does. The connections on U and V are *locallly* isomorphic to a connection that is zero everywhere within those regions (in some local representation), even though the connection on U and V is not. What this means is that *up to gauge transformations*, the connections on U and V *don't* determine the connection on $U \cup V$. This is the sense in which Dewar claims connections are non-separable. This shouldn't come as a surprise, since the gauge-invariant content of the bundle over a region U just consists of its holonomies, which we already saw are non-separable.[19]

The problem with Dewar's argument is that it assumes that sophistication considers states of *subsystems* of the universe as equivalent up to symmetries. But sophistication is chiefly interested in states of the universe as a whole. After all, the issue of indeterminism only arises when we consider global symmetries. In the case of gauge theories, this is obscured by the fact that local symmetries that vanish to the identity *are* global symmetries (Gomes, 2019). But a comparison with sophistication with classical mechanics clarifies this. Consider, for example, two (dynamically isolated) classical mechanical systems (such as fleets of spaceships, as in Rovelli (2014)). Both systems are invariant under static shifts, and hence the symmetry-invariant content of each system consists of their distance relations. However, the distances between ships of the first fleet and those between ships of the second fleet fail to fully determine the state of both fleets. After all, the latter also includes the distance between the two fleets, which is not contained in the invariant facts of either system. So it seems as if classical mechanics is also non-separable on a sophisticated account. But sophistication is *not* committed to a symmetry-invariant ontology of either distance relations or loop holonomies. On the contrary, sophistication embraces realism about non-invariant quantities such as fibre bundle connections or locations on a manifold. In the same vein, Gomes (2019) argues

[19] Note that, as Wallace (2014b) shows, this sort of non-separability of gauge-invariant content only occurs when the matter field vanishes — an assumption that is unphysical in light of quantum mechanics. Therefore, any form of non-separability discussed here is essentially fictional.

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.[20]

Indeed, the commitment to anti-quidditism implies that although universal symmetry-related models are physically equivalent, this is not the case for models of local subsystems. According to anti-quidditism, quantity-values are individuated by their relations. Crucially, these relations include relations to other subsystems. Therefore, sophistication opposes forgetting the symmetry-variant structure of subsystems, since this amounts to forgetting the 'identity' of qualitatively individuated entities. Since sophistication rejects Dewar's claim that only the symmetry-invariant content of subsystems is physically real, it does not imply non-separability.

8.7.3 Holism

Yet for all 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, along its trajectory, does the Aharonov-Bohm effect *come about*? As Healey points out, for any local representation of the connection and any subsection γ' of the electron's path, there is a gauge such that A_μ is flat over γ'. Therefore, Healey concludes, there is no path over which the electron non-trivially interacts with the Yang-Mills field — and hence no interaction at all!

Now, Healey's argument is clearly invalid, for while it is true that for any subsection there is a gauge in which A_μ vanishes, there is no gauge such that A_μ vanishes everywhere. In other words, it is a gauge-invariant fact that the local representation of the connection is non-trivial *somewhere*. But the argument is still instructive in highlighting the *holistic* nature of the Aharonov-Bohm effect.

The reason that there no meaningful answer to the question *where* the electron picks up a phase is that there are no meaningful cross-point comparisons of phase values, since each point is associated with its own set of field values. As we have seen, there is no canonical map from these local fibres into $U(1)$. Therefore, there *is* no picking up of phase; the electron simply moves across the fibre as dictated by the connection. We can only

[20] This is not to deny that isomorphic models of subsystems may have the same representational *capacities*; but this simply does not imply that such models are, for all purposes, physically equivalent. On this point, see also Roberts (2020).

compare field values at the same point, as when we compare the phase difference between the two electron-paths *at* Q. In other words, the holism of the Aharonov-Bohm effect consists of the fact that there is not enough structure to consider phase differences along open paths, but only around closed paths.

Perhaps the following analogy is helpful.[21] Consider the Twin Paradox: one twin remains in her inertial rest frame on earth, while the other travels to Alpha Centauri and returns. Since the latter measures less proper time, she is younger on return than the earth-twin. Just as we are interested in where the phase difference occurs in the Aharonov-Bohm effect, we may wonder when and at which rate the age difference between the twins comes into existence. Now, this is easy to do with respect to certain planes of simultaneity. For example, from the perspective of the earth-twin, the rocket-twin's clock runs slow at a constant rate. However, it is often[22] thought that simultaneity is *conventional*. And, as it turns out, different simultaneity choices yield different results for the differential aging of both twins (Debs and Redhead, 1996). For example, one can choose a convention such that the rocket-twin does not age slower at all on the first leg of the trip, but then ages *extra* slow on the return leg. This is rather similar to the fact that one can choose a gauge such that A is zero over any open path. (Indeed, Rynasiewicz shows that formally, a choice of simultaneity convention is equivalent to a choice of 'gauge' in GR). Just as there are no invariant cross-point comparisons of field values, so there are no objective cross-point simultaneity relations.

This implies that effects such as the Twin Paradox or the Aharonov-Bohm effect are holistic, in this sense: although the total effect size (e.g. the age difference) is measurable, there simply is no fact of the matter as to how this effect comes about as the result of small local differences. For example, the final age difference in the Twin Paradox is not the result of many small age differences that accrue locally; similarly, the final phase difference in the Aharonov-Bohm effect is not the sum of all the phase differences over infinitesimal paths. We can now see that such holism, although puzzling, is simply a consequence of the fact that value spaces are localised. This is not a defect in our theories, but a consequence of the novel metaphysics of fibre bundles.

[21] For another helpful analogy with currency exchange rates, see Maldacena (2014).

[22] But not universally — see Malament. I will not discuss this debate here.

8.8 Conclusion

I have defended a sophisticated interpretation of gauge theories in the fibre bundle formalism. This interpretation is local, separable and deterministic; it also clarifies the metaphysics of fibre bundles. I have explained the puzzling nature of the Aharonov-Bohm effect in terms of holism, which is a consequence of the fact that there are no cross-point comparisons between fibres. The account easily generates to non-Abelian gauge theories, since nothing in the above depended on the fact that the structure group $U(1)$ of electrodynamics is Abelian. Therefore, sophistication is preferrable over Wallace's gauge-invariant account.

There is also a broader lesson about symmetries 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 that determines the structure of its principal bundle, and hence of the associated bundle that represents matter fields. Far from a redundancy in our descriptions, symmetries are carriers of physical information. Instead of removing this structure from our theories, sophistication does justice to its significance.

Bibliography

Aharonov, Y. and D. Bohm. 1959. Significance of Electromagnetic Potentials in the Quantum Theory. *Physical Review 115*(3): 485–491.

Arntzenius, F. 2012. *Space, Time, and Stuff*. Oxford, New York: Oxford University Press.

Asselmeyer-Maluga, T. 2016. Smooth quantum gravity: Exotic smoothness and Quantum gravity. *arXiv:1601.06436 [gr-qc, physics:hep-th, physics:math-ph]*.

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.

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.

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

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/.

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.

Kobayaschi, 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/.

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*.

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.

Menon, T. V. 2018. *The Dynamical Foundations of Physical Geometry*. http://purl.org/dc/dcmitype/Text, University of Oxford.

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. 2005. 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.

Roberts, B. W. 2020. Regarding 'Leibniz Equivalence'. *Foundations of Physics 50*(4): 250–269.

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.

Rovelli, C. 2014. Why Gauge? *Foundations of Physics 44*(1): 91–104.

Wallace, D. 2002. Time-Dependent Symmetries: The Link Between Gauge Symmetries and Indeterminism. In *Symmetries in Physics: Philosophical Reflections*, eds. K. Brading and E. Castellani, 163–173. Cambridge University Press.

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.