

# Interpreting Gauge Theories

Review of *Gauging What's Real*, Richard Healey

Chris Smeenk

The Standard Model of particle physics describes three of the four fundamental interactions as Yang-Mills gauge theories. What does the empirical success of these theories reveal about the world? On Healey's view, general "empiricist principles" and the structure of the theories themselves are sufficient to single out one plausible interpretation for gauge theories – namely, the holonomy interpretation for classical theories, and an extension of it to quantum theories. The most striking feature of gauge theories on this interpretation is that they ascribe properties non-locally to loops rather than to points or their neighborhoods. Healey regards the conflict between this interpretation and principles such as Humean supervenience and separability as an opportunity for philosophical progress, another example of physical theory forcing us to consign cherished metaphysical principles to the flames.

Below I will spell out my reservations regarding Healey's formulation of the interpretative project and assess his case for the holonomy interpretation. But before doing so, I should emphasize that Healey's contribution is not limited to the articulation and defense of the holonomy interpretation. This book provides a concise overview of foundational topics in this area that would be an excellent starting point for philosophers eager to heed Earman's call to take up "gauge matters." He stakes out positions on a number of important issues that I will not have the space to discuss. The book also includes streamlined technical summaries of a variety of topics, such as the fiber bundle formalism, constrained Hamiltonian formalism, quantization techniques, and loop representations, to name a few. He does not limit consideration to the (potentially misleading) simplest Yang-Mills theory, QED, instead carefully noting complications that arise in the non-Abelian case, and general relativity is also often employed as a useful contrast with Yang-Mills gauge theories. Overall, Healey's exposition of these diverse technical topics is reliable and clear, although I doubt it is entirely accessible to the broad audience he hopes to reach. The downside of this synthetic approach is that Healey devotes relatively little effort to clarifying the connections between the pieces of technical machinery introduced along the way, and there is an opportunity for valuable foundational work exploring the subtle technical and conceptual differences among the various analyses of gauge freedom (see, especially, Belot 2003, Earman 2003, and Wallace 2003).

## 1 Fiber Bundles and Holonomies

The central technical result underwriting Healey's proposal is a reconstruction theorem showing how to express the content of a gauge theory in terms of holonomies rather than in the more geometrical language of the fiber bundle formalism. Given this result, the gauge interpreter can offer various reasons to prefer one or the other formal framework as a clearer guide for gauge metaphysics. But before turning to the interpretative arguments it will be useful to review briefly

the relation between fiber bundles and holonomies.<sup>1</sup>

Suppose that the phase space  $E$  used to describe a physical system fails to faithfully represent the physical states of the system, in that distinct points  $\in E$  map onto the same physical state. The natural response to this mathematical redundancy is to construct a smaller space  $M$  by identifying points in  $E$  that represent the same state. In a gauge theory, this redundancy has the following character: the collection of points that map into the same physical state are given by the orbit of a group  $\mathcal{G}$  of gauge transformations acting freely on  $E$ .<sup>2</sup> What mathematical structure to require of a gauge theory — whether, e.g., gauge theories are coextensive with or a superset of Yang-Mills theories — is a contentious issue, albeit one that is tangential to Healey’s concerns since he is focused on Yang-Mills theories.

Now we turn to the formulation of a Yang-Mills theory that includes “gauge redundancy.”<sup>3</sup> Consider a free field theory consisting of a set of field equations governing a field that takes values in some vector space  $V$  at each point in  $\mathbb{R}^4$  (e.g.,  $\phi : \mathbb{R}^4 \rightarrow \mathbb{C}$  for a complex scalar field). Yang and Mills demanded symmetry under “variable phase transformations”: the requirement that for a solution  $\phi(x)$ ,  $\phi'(x) = g(x)\phi(x)$  is also a solution — where  $g(x)$  assigns elements of a gauge group  $G$  to the points of  $\mathbb{R}^4$ . This was meant to strip the value of  $\phi(x)$  at a particular point of physical significance. Fulfilling this demand leads naturally to the fiber bundle formalism. A fiber bundle is a space  $E$  that has the structure of “fibers” attached to a base space  $M$ , defined via a projection map  $\pi : E \rightarrow M$  such that the fibers are given by  $\pi^{-1}(p)$  for  $p \in M$ . The fiber at each point is a copy of a vector space or a group, but there is not a way to identify the “same element” across distinct copies. Just as an  $n$ -dimensional differentiable manifold “locally looks like”  $\mathbb{R}^n$ , as required by the existence of coordinate charts, the total space of the fiber bundle “locally looks like” a product  $M \times G$ .<sup>4</sup> In the case at hand, the field  $\phi$  is treated as a section of a vector bundle, namely a map that assigns an element of the fiber  $V$  to each point in the base space  $M$ . The field equations are formulated in terms of sections, but we require further structure in order to define derivatives of sections — needed for the field equations — compatible with variable phase symmetry. A connection provides this further structure, in the form of a path-dependent comparison between different fibers over  $M$ . The connection is defined on a principal bundle, in which the fibers are isomorphic to the gauge group  $G$ , naturally associated to the vector bundle. Using the projection map  $\pi$  one can define a “vertical subspace” of the tangent space at any point  $p \in E$ , which corresponds to tangent vectors “lying in the same direction as the fibers”; the connection further assigns a horizontal subspace to each point. This makes it possible to define the “horizontal lifts”  $\{\tilde{\gamma}\}$  of any curve  $\gamma$  in  $M$ , a class of curves in  $E$  whose tangent vectors in  $T_p E$  lie in the horizontal subspace at each point. The connection naturally induces a covariant derivative in the vector bundle, and we have all the structures needed to formulate field equations with a variable phase symmetry. We can also introduce the curvature associated with a connection. A curve  $\tilde{\gamma}$  will return to the fiber over the base point of  $\gamma$ , but not necessarily to the same point along the fiber; the failure of  $\tilde{\gamma}$  to close is a measure of the curvature.<sup>5</sup>

---

<sup>1</sup>This is a quick tour of topics Healey covers in detail in Chapters 1, 3 and Appendix B. I would further recommend Baez and Munian (1994) as a fine pedagogically-oriented introduction to these topics.

<sup>2</sup>As a result  $E$  has the structure of a principal fiber bundle with the gauge group  $\mathcal{G}$ , but the base space is *not* the spacetime manifold. This bundle is distinct from that discussed below, where the structure group  $G$  is a compact group such as  $U(1)$  or  $SU(3)$ , and the base space is the spacetime manifold.

<sup>3</sup>See Martin (2003) for illuminating discussions of the conceptual foundations and historical development of Yang-Mills theory, including a detailed assessment of the “gauge argument.”

<sup>4</sup>Note that the fiber bundle is not quite a local product, because there is no canonical isomorphism from the fiber to the group  $G$  — the fiber does not have an intrinsically distinguished identity element.

<sup>5</sup>The curvature is defined in terms of the failure of “infinitesimal” loops to close, since a finite curve can fail to

How do we describe the physical states in such a theory? The choice of a section maps (locally) the connection onto the gauge potential  $A_\mu^a(x)$  and the curvature onto the gauge field strength  $F_{\mu\nu}$ . But due to the gauge redundancy, there are no observable differences among an equivalence class of fields and gauge potentials related by gauge transformations (elements of  $\mathcal{G}$ ). Within the fiber bundle formalism, gauge transformations are typically defined as smooth maps of the fiber bundle to itself that preserve both the fibers over points in the base space and the group action on the fibers (vertical bundle automorphisms). These maps transform the connection on the bundle, and induce the familiar gauge transformations for  $A_\mu^a(x)$  and  $\phi(x)$ . The observationally distinguishable possibilities then correspond to sets of gauge-related states  $\{A_\mu^a(x), \phi(x)\}$ , i.e. equivalence classes of states under the action of  $\mathcal{G}$ , also called gauge orbits.

Healey’s preferred interpretation is based on a way of directly characterizing the gauge-invariant structure in terms of holonomies. The horizontal lift  $\tilde{\gamma}$  of a closed curve  $\gamma$  need not return to the same point in the fiber; whatever point it does return to will be related to the initial point by some element of the gauge group  $G$ . Thus the horizontal lift  $\tilde{\gamma}$ , determined by the connection, generates a linear map from a fiber into itself that depends on the curve  $\gamma$ , called the holonomy  $H(\gamma)$ . One can define a group structure on closed curves in the base space  $M$ , what Healey calls the “group of hoops,” where the group product is given by composition of hoops (equivalence classes of closed curves differing only by curves that enclose no area). The holonomies can be treated as a convenient way of characterizing features of a fiber bundle, with the connection still regarded as more fundamental – but this choice is not dictated by the mathematics. Barrett (1991) showed that given a smooth homomorphism  $H$  from the group of hoops to  $G$  satisfying certain constraints, then there is a principal fiber bundle, unique in an appropriate sense, with group  $G$  and a connection such that  $H$  is the holonomy map for the bundle. Unlike gauge potentials, holonomies do not exhibit the redundancies associated with gauge freedom. In the Abelian case, the holonomies are gauge independent and independent of the choice of a base point for the group of hoops. The non-Abelian case is more subtle (pp. 106-110), in that the holonomies themselves are neither gauge-invariant nor independent of the choice of base point. But in both cases the gauge invariant states can be characterized without appealing to equivalence classes of gauge-dependent quantities. However, the cost of banishing gauge redundancy is a form of nonlocality or nonseparability (more on this below). Since gauge transformations vary spatially, gauge-invariant states cannot be built up out of equivalence classes of localized states. The holonomy interpretation displays this quite clearly: the states are defined with respect to closed curves in the manifold, not with respect to spacetime points.

## 2 The Ontological Debate

Healey sets up the interpretative debate as a contest between three ways of understanding the ontology of gauge theories. The first possibility is never treated as a serious contender. This account treats gauge potentials much like the vector potential in classical electromagnetism prior to the discovery of the Aharonov-Bohm effect — that is, as a useful calculational device that does not represent new physical properties. Healey argues that such an interpretation implies action-at-a-distance (§2.4), and furthermore that there are cases of physically distinct states that differ only with respect to gauge potentials (§4.1). The main attraction is the fight between the two remaining views. In one corner, we have the connectionist, who interprets gauge potentials realistically, and takes them to represent local properties; in the other corner, Healey defends the

---

close even though the curvature vanishes if the base space is not simply connected.

claim that the gauge theories instead represent non-localized properties via holonomies. The fight focuses on the implications of gauge freedom.

One might expect Healey to attempt a knock-out blow against the connectionist with the hole-argument combination, used here in the case of gauge theories rather than in the fight regarding spacetime ontology. According to the connectionist, a particular connection  $\omega$  (or, locally, a particular gauge potential  $A_\mu^a(x)$ ) represents the physical state of a system. Applying a gauge transformation yields  $\tilde{\omega}$ , which the realism of the connectionist apparently commits her to treating as a distinct physical state. But treating  $\omega$  and  $\tilde{\omega}$  as distinct physical states leads to a failure of determinism, because the dynamical evolution of the theory only fixes unique solutions up to gauge equivalence. Determinism should not be allowed to fail for “metaphysical” reasons, such as a poor choice regarding individuation of physical states. So gauge-variant quantities like  $\omega$  should be dropped in favor of gauge-invariant quantities when it comes to representing physical states.

Healey takes a closely related argument to leave the connectionist sprawled on the mat, despite denying that the blow has any force in the fight over spacetime ontology. Like many more recent “sophisticated substantialists,” Healey has argued that a spacetime substantialist can accept Leibniz equivalence (Healey 1995) (that is, take two spacetime models related by a diffeomorphism, the analogs of  $\omega$  and  $\tilde{\omega}$ , to represent the *same* state of affairs) and thus avoid the hole argument. But he draws a distinction between the case of general relativity and Yang-Mills theories. In the fiber bundle formulation of general relativity, the connection  $\omega$  naturally breaks down into the “solder-form”  $\Theta$  plus the connection  $\nabla$ .<sup>6</sup> A gauge transformation  $\omega \rightarrow \tilde{\omega}$  induces a transformation of both  $\Theta$  and  $\nabla$ , such that the geometrical properties originally assigned to a manifold point  $m \in M$  are mapped onto some other point  $f(m)$ . But since  $m$  and  $f(m)$  agree on all geometrical properties, Healey takes them to represent the *same* spacetime point  $p$  — and hence the two models related by a symmetry represent the same state of affairs. However, in the case of gauge transformations in a Yang-Mills theory the “internal properties” assigned at spacetime points are changed. As a consequence, according to Healey (pp. 97-98), one cannot similarly claim in a Yang-Mills theory that  $\omega$  and  $\tilde{\omega}$  are different representations of the same state of affairs, as they assign *different* gauge potential properties to the same spacetime point (the connection is not “soldered” to the manifold). Thus the connectionist cannot accept an analog of Leibniz equivalence. I called this a variation on the hole argument because Healey takes the force of the punch to be the multiple realizability of the theory (pp. 89-99): the theory cannot uniquely specify the distribution of gauge potential properties represented by a given  $A_\mu^a(x)$ . This is closely related to hole-argument indeterminism, and I am skeptical how much more illuminating it is to reframe the debate in these terms — the status of Leibniz equivalence still seems to be the crucial point. Whether the hole argument combination is effective depends on whether the technical contrast Healey emphasizes really forces the connectionist to deny Leibniz equivalence.

The connectionist has two counter-punches. First, how does the holonomy interpretation explain the constraints that the holonomies must satisfy? (The connectionist learned this move from the spacetime substantialist, who similarly objects that the relationalist cannot explain why distance relations between objects satisfy the axioms of Euclidean geometry.) For two curves  $\gamma_1, \gamma_2$ , the holonomy of the curve given by composing them ( $\gamma_1 \circ \gamma_2$ ) is just  $H(\gamma_1)H(\gamma_2) = H(\gamma_1 \circ \gamma_2)$  (pp. 120-121). The connectionist can treat these as derived from facts regarding gauge potentials,

---

<sup>6</sup>As Healey reviews in §3.2, “soldering” refers to the fact the frame bundle  $F(M)$  is more closely tied to the base space  $M$  than other principal bundles because the elements of the frame bundle correspond to frames, i.e. sets of basis vectors for the tangent space  $T_pM$ . This extra structure in the frame bundle is exhibited by the solder-form  $\Theta$ :  $\Theta$  maps tangent vectors in  $T_pF(M)$  into their components in the corresponding frame, and it can also be regarded as a one-form on  $T_pM$  mapping vectors into their components.

but claims that Healey must take such constraints as brute facts since the holonomies are treated as primitive. Second, the gauge-invariant state favored by the holonomy interpretation is a holistic, non-local mess. As noted above, holonomies are defined over closed loops, extending to arbitrarily large distances, rather than over spacetime points and their neighborhoods.<sup>7</sup> Furthermore, for a theory including matter fields the gauge transformations act on the connection as well as the field in the associated vector bundle, and as a result the gauge-invariant characterization of the state requires a combination of the connection and matter fields.

The first point does not win the fight for the connectionist, because what constitutes a good explanation is itself contentious: the connectionist “explains” the holonomy constraints as consequences of the local properties represented by gauge potentials, Healey objects that good explanations cannot appeal to such inaccessible quantities, and we are back to the issue of gauge freedom and its implications. Healey responds to the second point with a judo-like maneuver, using the attack to throw the connectionist into the position of a reactionary metaphysician. Formulating the precise sense in which the holonomy interpretation is “non-local” is a delicate matter. Healey argues that on this interpretation, gauge theories describe “strongly non-separable” processes, defined as such by their failure to supervene on “intrinsic physical properties” defined over spacetime points or their neighborhoods. Healey draws two further conclusions: first, that there is no troubling action-at-a-distance in this case and, second, that on the holonomy interpretation gauge theories violate Lewis’s Humean Supervenience. So much, says Healey, for Humean Supervenience — gauge theories reveal it to be “outdated” (p. 128).

I find an alternative approach to defining locality and separability more promising than the line Healey pursues, because I am unsure what to make of basic notions employed in his definitions (such as “all the causes of an event” and qualitative intrinsic properties). My preferred approach (exemplified by Earman 1987) would rely instead on the mathematical structure of the physical theories at hand, and appeal to notions such as the state of a system and the functional dependences that hold between states in different regions as a result of the equations of motion. Similarly, rather than formulating holism in terms of supervenience, I would prefer to define a theory to be holistic if (roughly speaking) it is impossible to infer the global properties of a system based on accessible information regarding its subsystems, shifting the focus to the way a theory defines the joint state for a composite system (following Seevinck 2004’s analysis of holism in QM). Although this line of analysis would lead to a number of differences of detail with Healey’s account, I do not expect that it would undermine his main conclusions. However, it does indicate that his conclusion regarding Humean supervenience might not be as powerful as it seems at first glance. The blow could be dodged by a reformulation of Humean Supervenience that retains the spirit of Lewis’s proposal without a commitment to a “vast mosaic of local matters of fact.” Earman and Roberts (2005) define the Humean base as the non-nomic facts that can be the output of a reliable measurement confined to a finite spacetime region; Humean supervenience is then the claim that the laws of nature supervene on the Humean base. (They were motivated in part by the non-separability of non-relativistic quantum mechanics.) What can be reliably measured depends on current theory, so Humean supervenience is no longer pinned down to classical physics as it was in Lewis’s original formulation. Although one might well object to their proposal on other grounds, I don’t see any new threats arising from gauge theories.

In sum, Healey argues that the holonomy interpretation ends the bout as the clear winner. While I do not think the case is quite as clear cut, a more important question is whether this is

---

<sup>7</sup>In special cases one can still recover a “local” state from holonomies; for example, in the formulation of classical electromagnetism as a gauge theory, if space is simply connected then the holonomies in a region fix the magnetic field within that region, and thus one can build up a global, gauge-invariant state from localized states.

the right fight to be having (cf. Belot 2003). The fight has been staged for *classical* Yang-Mills theories, but this should be only a warm-up for the real contest regarding the *quantized* Yang-Mills theories that have been successfully applied to the world.

### 3 From Classical to Quantum

How does the interpretation of classical gauge theories relate to that of their quantized counterparts? My main disagreements with Healey’s interpretative project are related to his answer to this question, insofar as I have correctly understood it. The question is pressing in part because non-Abelian gauge theories have not been empirically successful in *any* domain; they contribute to empirically successful predictions of the Standard Model only indirectly as the basis for our (limited) understanding of quantized Yang-Mills theories. Healey acknowledges this point (pp. 82, 130), but then insists that we should “bracket” the question of empirical adequacy in interpreting the classical theories. When he turns to quantized Yang-Mills theories in the latter half of the book, the extension of the holonomy interpretation takes the form of a preference for an intrinsic, gauge-free formulation of the theory, via the loop representation. Healey apparently takes the arguments above in the classical case to give us reason to prefer a gauge-free formulation of a theory, a conclusion that is further supported by his classification of local gauge symmetry as a theoretical rather than empirical symmetry whose elimination entails no loss of empirical content.<sup>8</sup> Healey does not claim to fully spell out an interpretation based on the loop representation, limiting himself instead to suggesting how it relates to several general lines of attack for interpreting quantum field theories. But he does argue that the non-separability of the holonomy interpretation probably carries over to quantized theories.

My objection is that Healey does not sufficiently acknowledge the subtle interplay between interpretation of gauge theories and quantization techniques. Healey’s position is based in part on extending a line of reasoning from the classical to the quantum case. But this is especially precarious if crucial aspects of the classical theory are introduced for pragmatic reasons. The gauge potentials have a role within the theory dictated by the need to generate a quantum theory given limits imposed by the mathematics of quantization, and not directly by the demand to suitably represent properties. Guay (2008) has argued that gauge potentials (in conjunction with BRST symmetry) allow for the formulation of a local field theory that can be quantized.<sup>9</sup> Healey is certainly aware of these issues and emphasizes that there are sundry quantization techniques that yield different candidates for the quantized version of a given classical theory. If gauge reduction and quantization fail to commute, different interpretative stances regarding the classical theory would lead to empirically inequivalent quantum theories. In this situation empirical considerations would then give us guidance for interpreting the quantum theory and its classical progenitor. (The failure to lead to equivalent theories might instead lead to refinement of the quantization techniques.) However, Healey explicitly does not make the case on these grounds: he assumes that there is a so-called “loop transform” relating any loop representation (admitting a Fock

---

<sup>8</sup>This interesting line of argument is developed in Chapter 6. It is not clear why Healey does not introduce the distinction between empirical and theoretical / formal symmetries earlier and apply it also to the classical case.

<sup>9</sup>Guay(2008) emphasizes the important role of BRST symmetries: first the phase space is extended to include extra, non-physical “ghost” degrees of freedom to enable quantization, but these are subsequently eliminated by imposing an additional constraint on the “physical” Hilbert space. Guay further discusses DeWitt’s argument that gauge and BRST symmetries are needed to insure the equivalence of the quantization of a gauge theory and that of a gauge-reduced theory.

representation) to a connection representation, extrapolating from results showing that such a transform exists for some specific cases (p. 199).

Regarding loop representations, I am also inclined to emphasize more pragmatic reasons for their importance. Wilson loops were introduced in the study of quark confinement in QCD (namely, the physical states are invariant under the gauge group  $SU(3)$  even though quarks are not), a striking feature of the quantized theory that is absent from classical Yang-Mills theory and inaccessible to perturbation theory. The main appeal of getting rid of gauge freedom was to calculate quantities directly (in lattice gauge theory) without using perturbative expansions and gauge fixing. Loop representations for the continuum case will presumably provide insights into various non-perturbative aspects of QFT, and this seems to be the main reason for their appeal. Healey emphasizes instead their representational transparency due to the elimination of gauge freedom, achieved by defining the Hilbert space in terms of wave-functionals over the domain of loops rather than gauge potentials. He acknowledges an important limitation faced by existing loop representations, namely an inability to incorporate interactions with fermion fields (pp. 195-197). There is another problem best brought out by a contrast with loop quantum gravity. Recall that the holonomies around different loops are not independent of each other; instead they satisfy various constraints, reflecting the fact that they are over-complete as coordinates for the reduced configuration space. The resulting difficulty in finding a set of linearly independent basis loop states leads to various problems in quantization, in particular the kinematical Hilbert space produced by canonical quantization is non-separable. This is a desirable result in the case of loop quantum gravity, because the imposition of the diffeomorphism constraint eliminates the surplus states and yields a separable Hilbert space. But it is a significant obstacle to using loop representations in continuum Yang-Mills theories.

Loop representations are surely an important part of the quest for a firmer understanding of the rich mathematical structure of quantum Yang-Mills theories. I agree with Healey that it is interesting to explore their implications for foundational debates in QFT. Healey takes up this task in Chapter 8, surveying the status of holonomy properties in quantized Yang-Mills theories within the context of particle, Bohmian, Copenhagen, Everettian, and modal interpretations, in each case assessing whether holonomy properties are more representationally transparent than gauge potentials by the lights of the interpretation. In each case, Healey's conclusions are qualified due to open technical questions (e.g., whether there is a Fock representation for a particular Weyl algebra of loops, p. 208, or whether the loop representation is disjoint from usual Fock representation, p. 214) or more basic problems regarding the interpretation itself. While Healey broaches a number of interesting issues here, I take the implications of loop representations to be not so directly tied to Healey's brief for non-separable holonomy properties and instead related to their utility in addressing non-perturbative problems. Addressing these issues — whether via loop representations or other approaches — will certainly shed more illumination on the mathematical structure of Yang-Mills theories, and perhaps it will eventually be possible to make a more direct case for the existence of non-separable processes in quantized Yang-Mills theories.

## References

- Baez, J. and Muniain, J. (1994). *Gauge fields, knots and gravity*. World Scientific River Edge, NJ.
- Barrett, J. (1991). Holonomy and path structures in general relativity and Yang-Mills theory. *International Journal of Theoretical Physics*, 30(9):1171–1215.

- Belot, G. (2003). Symmetry and gauge freedom. *Studies in History and Philosophy of Modern Physics*, 34(2):189–225.
- Earman, J. (1987). Locality, nonlocality, and action at a distance: A skeptical review of some philosophical dogmas. In Achinstein, P. and Kargon, R., editors, *Kelvin's Baltimore Lectures and modern theoretical physics: historical and philosophical perspectives*, pp. 449–490. MIT Press, Cambridge.
- Earman, J. (2003). Tracking down gauge: an ode to the constrained Hamiltonian formalism. In Brading, K. and Castellani, E., editors, *Symmetries in Physics: Philosophical Reflections*, pp. 140–162. Cambridge University Press, Cambridge.
- Earman, J. and Roberts, J. (2005). Contact with the Nomic: A Challenge for Deniers of Humean Supervenience about Laws of Nature Part I: Humean Supervenience. *Philosophy and Phenomenological Research*, 71(1):1–22.
- Guay, A. (2008). A partial elucidation of the gauge principle. *Studies in History and Philosophy of Modern Physics*, 39(2):346–363.
- Healey, R. (1995). Substance, modality and spacetime. *Erkenntnis*, 42(3):287–316.
- Martin, C. (2003). On continuous symmetries and the foundations of modern physics. In Brading, K. and Castellani, E., editors, *Symmetries in Physics: Philosophical Reflections*, pp. 29–60. Cambridge University Press, Cambridge.
- Seevinck, M. (2004) Holism, physical theories and quantum mechanics. *Studies in the History and Philosophy of Modern Physics*, 35(4):693–712 (2004)
- Wallace, D. (2003). Time-dependent symmetries: the link between gauge symmetries and indeterminism. In Brading, K. and Castellani, E., editors, *Symmetries in Physics: Philosophical Reflections*, pp. 163–173. Cambridge University Press, Cambridge.