Received: 13 October 2016 | Revised: 5 January 2017 | Accepted: 25 January 2017
DOI: 10.1111/phc3.12414

WILEY

ARTICLE

A philosopher's guide to the foundations of quantum field theory

Noel Swanson

University of Delaware

Correspondence
Noel Swanson, Department of Philosophy, University of Delaware, 24 Kent Way, Newark, DE 19716, USA.
Email: nswanson@udel.edu

Abstract

A major obstacle facing interpreters of quantum field theory (QFT) is a proliferation of different theoretical frameworks. This article surveys three of the main available options—Lagrangian, Wightman, and algebraic QFT—and examines how they are related. Although each framework emphasizes different aspects of QFT, leading to distinct strengths and weaknesses, there is less tension between them than commonly assumed. Given the limitations of our current knowledge and the need for creative new ideas, I urge philosophers to explore puzzles, tools, and techniques from all three approaches.

1 | INTRODUCTION

Relativistic quantum field theory (QFT) unifies core ideas from quantum mechanics and relativity, providing a framework for the standard model of particle physics. Despite the extraordinary empirical success of the standard model, our current understanding of the mathematical and conceptual foundations of QFT is still nascent. One of the major difficulties facing interpreters of QFT is the problem of inequivalent representations. Because field systems have infinitely many degrees of freedom, standard quantization procedures for translating classical theories into quantum ones typically yield infinitely many physically inequivalent representations of the same basic physical quantities. How to make sense of this explosion of possibilities has been a topic of much debate. In a previous two-part article in this journal, Ruetsche (2012) examines both the interpretive challenges and numerous explanatory applications of inequivalent representations in QFT.

Here, we focus on a more fundamental difficulty: the proliferation of different theoretical frameworks for QFT. Unlike the situation in other corners of physics including classical field theory and non-relativistic quantum mechanics, we currently lack a canonical mathematical language for QFT that is both internally consistent and empirically adequate. The textbook Lagrangian approach to QFT, although astonishingly fruitful, is a grab bag of conflicting mathematical ideas. Axiomatic and constructive field theory offer greater rigor but have yet to reproduce the empirical and explanatory successes of the standard model.

This article aims to provide a map of this difficult, shifting terrain for the general philosophical audience. In Section 2, I give a brief survey of three of the main approaches to the foundations of QFT—*Lagrangian*, *Wightman*, and *algebraic QFT*. In Section 3, I chart the relationship between these approaches. Some philosophers have characterized them as inequivalent, rival research programs, arguing that it is important to determine which approach is the best before the task of interpreting QFT can begin. After reviewing the debate, I conclude that although each approach emphasizes different aspects of QFT, there is less tension between them than commonly assumed. Given the limitations of our current knowledge and the need for creative new ideas, I urge philosophers to take a more cosmopolitan stance that better aligns itself to scientific practice. In this spirit, Section 4 explores some recent developments in the rich trading zone between Lagrangian, Wightman, and algebraic QFT that should pique the interest of philosophers of physics.

2 | THREE VARIATIONS ON QFT

2.1 | Lagrangian QFT

Lagrangian QFT takes guiding ideas from classical field theory and adapts them to the quantum realm. Classically, a field is an extended object characterized by a suitably continuous assignment of mathematical variables (e.g., scalars, vectors, tensors, or spinors) to spacetime points.[1] Dynamically possible field configurations are encoded in a Lagrangian function via the principle of stationary action. Various quantization rules can then be employed to translate this classical structure into a QFT. The classical field variables become *field operators*, which act as (unbounded) linear transformations on a Hilbert space representing the possible states of the quantum theory, while the path integral formalism serves as an extension of the classical action principle. An assortment of perturbative approximation techniques codified by Feynman diagrams are then used to extract information from the model.

Although mathematically similar to the operators used in standard quantum mechanics to represent physical quantities, such as energy, spin, and charge, the field operators in Lagrangian QFT typically do not represent quantities directly. Instead, physical quantities (somewhat misleadingly called *observables*), are built out of algebraic combinations of the field operators. If there is a gauge symmetry present, only those algebraic combinations of the fields that are invariant under the gauge action are deemed to have direct physical significance. To put it metaphorically, the field operators are superfluous descriptive fluff, much like the vector potential in classical electromagnetism. There is still ongoing debate about whether or not this is the correct interpretation of gauge symmetries in QFT, but it is the consensus view amongst a healthy majority of working physicists and philosophers alike, and we will adopt this perspective here.[2] It should be further noted that gauge symmetries come in two varieties - *global* gauge symmetries, which act uniformly on all field operators, and *local* gauge symmetries, which vary continuously as a function of spacetime location. Local gauge symmetries are a crucial feature of the standard model and similar interacting QFTs.

In spite of a plethora of fruitful applications in particle physics and beyond, Lagrangian QFT is really a cluster of different computational techniques.[3] It does not give us a sharp mathematical characterization of what counts as a model of QFT, making it challenging to prove rigorous general theorems. Studying field theories in the Lagrangian approach usually proceeds on a model-by-model basis. In addition, most interacting models have unresolved internal problems. These can be divided into three broad classes.

First, there are problems with quantization rules. Standard canonical quantization treats nonlinear interactions as small perturbations of a free field system, an idealization that fails for strongly coupled theories such as quantum chromodynamics (QCD). More powerful path integral quantization methods avoid this problem, but these require choosing a measure over an infinite dimensional path space, a procedure that is only mathematically well-defined in certain special cases. In practice, these complications are typically ignored.[4] These issues aside, once we settle on a particular quantization method, there is still the vexing problem of inequivalent representations.

A second set of issues concern whether or not the output quantum models are mathematically consistent. To get physical predictions, we must compute functional integrals that include contributions from physical processes at both extremely short distances or high energies (ultraviolet regime) and extremely long distances or low energies (infrared regime). Both can cause the integrals to blow up, yielding senseless infinities. Even when the practical challenges posed by ultraviolet and infrared divergences can be tamed using cutoffs and renormalization group techniques, a number of conceptual problems continue to linger. For example, the standard picture of particle interactions models scattering processes using a unitary S-matrix that connects interacting fields to free incoming and outgoing fields. Haag's theorem shows that such a unitary operator does not exist unless the field is free at all times. Earman and Fraser (2006) and Fraser (2008) explore the foundational implications of Haag's theorem.

Finally, even in the weak-coupling limit where perturbation theory can be formally applied, there are questions about when it gives us an accurate approximation of the underlying physics. Dyson (1952) contains a famous heuristic argument concluding that generic renormalized perturbation series in quantum electrodynamics (QED) do not converge.[5] Even if the series diverge, however, they might still be asymptotic to some well-behaved limiting theory. The problem is that without a better theoretical understanding of the non-perturbative situation we cannot say very much about what this underlying physics looks like. This issue is highlighted by recent triviality results for an important interacting model called *ϕ^4-theory* (Fernandez, Fröhlich, & Sokal, 1992). In $d < 4$ dimensional Minkowski spacetime, the theory is nontrivial, but in $d > 4$, the renormalized perturbation series is asymptotic to a free field theory despite the fact that it appears to describe nontrivial interactions.[5] Similar triviality results exist for the borderline $d = 4$ case, although these rely on technical assumptions that require further scrutiny.[6] Results like these indicate that we still do not fully understand when we can trust the perturbative techniques employed by Lagrangian QFT.

2.2 | Wightman QFT

In order to address these problems and to provide a mathematically precise answer to the question "what is a model of QFT," various axiomatic approaches have been developed. The two most well-established frameworks are Wightman and algebraic QFT.[7] Both start from standard Lagrangian models formulated in Minkowski spacetime and attempt to frame precise, general principles that any physically reasonable QFT must obey. They focus on slightly different parts of the theory; however, and as a result, they differ in their respective strengths and weaknesses.

Wightman QFT hews closely to Lagrangian QFT, providing axioms describing covariant field operators acting on a fixed Hilbert space.[8] Because of ultraviolet divergences, field operators cannot be assigned to individual spacetime points. We can only coherently talk about the field configuration in extended regions of spacetime. Nonetheless, the axioms ensure that these regions can always be made arbitrarily small and compact. Formally, this requires representing the fields as tempered operator-valued distributions, $\phi(f)$, where f is a smooth test function of compact support on Minkowski spacetime.

A model of the Wightman axioms consists of a 4-tuple $(\{\phi(f)\}, \mathcal{H}, U, \Omega)$, where $\{\phi(f)\}$ is a family of field operators acting on a separable Hilbert space $\mathcal{H}$, U is a continuous group of unitary operators representing the Poincaré transformations, and Ω is a unique translation-invariant vector representing the vacuum state. The generators of the translation subgroup represent the energy–momentum observables; thus, the dynamics are implicitly contained in the representation U.

There are three central axioms. The first, *covariance*, requires that field operators transform in the same way that spacetime coordinates do under U, making it easy to formulate Lorentz-invariant laws. The second, *microcausality*, demands that spacelike separated field operators must either commute or anticommute. Although usually viewed as a sufficient condition to enforce relativistic no-signaling constraints, Ruetsche (2011, section 5.3.1) details a number of shortcomings of this interpretation, arguing instead that microcausality is best viewed as a mereological constraint on independent spacelike separated subsystems. (It should also be noted that microcausality is closely tied to the spin-statistics theorem: in order for a model to satisfy the remaining Wightman axioms, fields with integer spin must commute and fields with half-integer spin must anticommute.) The third central axiom, the *spectrum condition*, requires that the energy must be positive in all Lorentz frames, ensuring that the vacuum state is a stable ground state.[9] Streater and Wightman (1989) emphasize how these three axioms along with the compact localization properties of the field operators entail that many physically significant functions of the fields (e.g., the n-point vacuum correlation functions) can be analytically continued to holomorphic functions on the complexification of Minkowski spacetime, $M^{\mathbb{C}} = M \oplus iM$. Although the physical significance of this fact remains mysterious, the mathematical upshot is that holomorphic functions are equal to their own Taylor series; thus, their global behavior is determined by their local behavior in the neighborhood of any point. These extension techniques form the backbone of the first rigorous proofs of a number of important structural results including the parity-charge-time (PCT) and spin–statistics theorems (Jost, 1957). In addition, they give rise to powerful tools for constructing models of the Wightman axioms (Glimm & Jaffe, 1987).

Despite these strengths, the formalism does have certain weaknesses. First, the axioms concern gauge-dependent field operators. Because these do not directly represent physical quantities on the standard view, it makes the physical interpretation and justification of the axioms difficult. Second, many of the proofs in the Wightman framework, although mathematically rigorous, are physically opaque because the analytic continuation techniques they rely upon lack clear physical or explanatory content.[10] Third, a series of no-go results due to Strocchi indicate that the axioms cannot directly accommodate local gauge theories such as the standard model. Either one must employ a generalized Hilbert space with an indefinite inner product (a so-called *Krein space*) or work with non-local field operators that violate the microcausality and covariance axioms.[11]

2.3 | Algebraic QFT

Rather than assumptions about field operators, algebraic QFT supplies axioms directly for gauge-invariant observables.[12] On this picture, each region of spacetime is assigned a collection of observables $\mathcal{A}(O)$ corresponding to physical quantities that can be measured by experiments causally confined to that region. At the outset, we do not assume that these observables are concretely represented as linear transformations on a particular Hilbert space. Instead, the collection comes equipped with an intrinsic algebraic structure (usually that of a C^*-algebra), which can be used to directly describe relations between quantities.[13]

States are given by (positive and normalized) linear functionals on the net of local algebras that directly encode the expectation values of the observables. Given a state, the Gelfand–Naimark–Segal (GNS) construction determines a unique representation of the net as an algebra of bounded linear operators acting on a separable Hilbert space. Working in a concrete representation yields additional topological tools that can be used to define important global quantities including temperature, energy, charge, and particle number. Unlike Wightman or Lagrangian field theory, algebraic QFT typically makes use of multiple unitarily inequivalent representations, which are distinguished by different values of these global observables. It is this flexibility that has allowed the approach to provide deep insight into a range of phenomena including quantum phase transitions, spontaneous symmetry breaking, charge superselection rules, particle statistics, and the concept of antimatter (Halvorson & Müger, 2006; Ruetsche, 2011).

For example, the Doplicher–Haag–Roberts (DHR) analysis of charge superselection rules begins with the assumption that the only physically possible global states are ones that differ from the vacuum in some compact region (Doplicher, Haag, & Roberts, 1969a,b). (Such states are conjectured to be suitable characterizations of charged states in strongly interacting gauge theories such as QCD.) The family of corresponding GNS representations has the structure of a symmetric tensor *-category, which allows us to define matter and antimatter states lying in unitarily inequivalent, conjugate representations. This gives rigorous mathematical content to the idea that such states have opposite charge quantum numbers. As Baker and Halvorson (2010) argue, the DHR picture has the advantage of applying to interacting QFTs, which lack a clear particle interpretation. In contrast, the standard characterization of antimatter from Lagrangian QFT relies on the emergent particle picture supplied by non-interacting Fock-space theories (Wallace, 2009).

Putting these pieces together, we can view a model of algebraic QFT as a 4-tuple $(\{\mathcal{A}(O)\}, \{\rho\}, \alpha, \omega)$, where $\{\mathcal{A}(O)\}$ is a net of local C^*-algebras, $\{\rho\}$ is a collection of physically possible global states, α is a continuous group of net automorphisms representing the Poincaré transformations, and ω is a unique translation-invariant vacuum state.[14] The main Haag–Kastler axioms mirror the central Wightman axioms: the local observable algebras must transform covariantly under Poincaré symmetries (or at least the translation subgroup), spacelike separated observables are required to commute, and in the vacuum GNS representation (where the translations are implemented by unitary operators), the spectrum condition is required to hold. Algebraic covariance is a weaker, more natural assumption than Wightman covariance because it is possible to build a covariant net of observables out of noncovariant Wightman fields (Strocchi, 2013). Algebraic microcausality has a similar advantage over its Wightman counterpart. In $d < 4$ dimensions, it permits plektons, new particle types that obey generalized Bose–Fermi statistics called *braid group statistics* (Fredenhagen, Rheren, & Schroer, 1989). Moreover, the analysis of Buchholz and Fredenhagen (1982) shows that the Haag–Kastler axioms are consistent with solitons and other types of topological charges, which cannot be generated by compactly localized Wightman field operators.

Because the algebraic approach has more physically transparent initial assumptions, many theorems in QFT receive their sharpest, most perspicuous formulation within the setting of algebraic QFT. It is largely for this reason that algebraic QFT has received a fair amount of attention from philosophers of physics. At the same time, it turns out to be quite difficult to build algebraic models directly. The normal strategy is to build a Wightman model by exploiting the analyticity properties of local field operators and then show that it generates a net of observables satisfying the Haag–Kastler axioms. Summers (2012) contains an excellent survey of the problems and methods of constructive field theory, including recent advancements in purely algebraic constructive techniques.

3 | RIVAL PARADIGMS?

With three versions of QFT on the table, a natural foundational question arises: How are they all related? Because algebraic and Wightman QFT have precise axioms, we can ask well-posed mathematical questions about their relationship. It turns out that in their standard form, they are not equivalent. Given a Wightman model, we can always construct a net of associated observable algebras by taking bounded functions of the field operators. The resulting net automatically satisfies all of the Haag–Kastler axioms except the requirement that the local algebras are C^*-algebras. This can be ensured by imposing H-bound constraints (Borchers & Yngvason, 1992). Such bounds are satisfied in all known models of physical interest, and without them, it is possible to construct pathological models, where the local field strength can be arbitrarily high while the energy is arbitrarily low. (In low-dimensional massless theories, additional constraints are needed to rule out this kind of pathology.)

Thus, with some mild physicality assumptions in place, Wightman QFT becomes a subtheory of algebraic QFT. In the other direction, much less is known. The examples canvassed in Section 2.3 indicate that algebraic QFT is more general than standard Wightman field theory; however, the possibility of plektons and topological charges suggests natural extensions of the standard Wightman framework involving field operators obeying generalized commutation relations and localized in cones extending to spacelike infinity (Rehren, 1992). In certain restricted contexts, it is possible to reconstruct a model of Wightman QFT from an algebraic model. In theories with a compact global gauge symmetry, Doplicher and Roberts (1972) prove that it is possible to reconstruct an underlying field algebra and gauge group given the observable net and a family of states satisfying the DHR selection criteria. In certain cases it has been shown that the field algebra generates a family of Wightman fields in a suitable GNS representation, but in general the question remains open.[15] In a broader setting, Fredenhagen and Hertel (1981) explore the possibility of defining local field operators by taking intersections of energy-damped local observable algebras. Building on their work, Bostelmann (2005) shows that under the assumption that the theory has a well-behaved ultraviolet regime, this technique succeeds in defining a family of field operators that satisfy the main Wightman axioms. In general, however, these field operators do not generate the entire net of observables, so the original algebraic model contains more physical information than the reconstructed Wightman model.

Because Lagrangian QFT is more mathematically amorphous, its relationship to the two axiomatic approaches is harder to pin down. On the one hand, recent algebraic constructions of low-dimensional interacting theories have produced models with no known Lagrangian, suggesting that the algebraic approach is more general (Lechner, 2015). On the other, the Lagrangian approach accommodates local gauge theories, which have yet to be rigorously constructed in either the algebraic or Wightman framework. Prima facie, the Haag–Kastler axioms should be compatible with the kinds of non-local field operators required by one horn of the Strocchi no-go results mentioned at the end of Section 2.2; however, our current mathematical understanding of such objects is poor. So while we have no reason to suspect that local gauge theories are incompatible with algebraic QFT, it is too soon to say for sure. Strocchi (1978) takes the other viable path, proposing an extension of the Wightman axioms to field operators acting on a Krein space. Known as *Strocchi–Wightman QFT*, this approach can take advantage of the rich analyticity properties of local field operators. The disadvantage is that additional constraints are needed to identify the physical states of the theory, which lie in some Hilbert subspace of the Krein space. Nonetheless, it dovetails elegantly with Becchi-Rouet-Stora-Tyutin (BRST) quantization techniques employed in Lagrangian QFT and has helped shed light on the non-perturbative structure of the Higgs mechanism and the U(1) problem in QCD (Strocchi, 2013). Finally, an important series of papers culminating with Balaban (1989) employs tools from both Lagrangian and Wightman QFT to explore the renormalization group flow of strongly interacting SU(N) gauge theories. Starting with a Lagrangian QFT defined on a lattice with both a short and long distance cutoff, the papers succeed in giving a rigorous proof of the ultraviolet stability of the theory in the limit as the lattice spacing goes to zero. Although this suggests that there exists a well-defined continuum description of such theories, the calculation of the associated vacuum expectation values and a proof that they satisfy the constraints of axiomatic QFT remains an open problem.[16]

Although these results support a cautious optimism about the future convergence of Lagrangian, algebraic, and Wightman QFT, some philosophers of physics have taken a different stance. Both Fraser (2009, 2011) and Wallace (2006, 2011) argue that Lagrangian and axiomatic QFT represent rival theoretical alternatives. As Fraser puts it

> *QFT presents a genuine example of the underdetermination of theory by empirical evidence. There are variants of QFT—for example the standard textbook formulation and the rigorous axiomatic formulation —that are empirically indistinguishable yet support different interpretations. (2009, p. 536)*

Fraser champions algebraic QFT, citing the explanatory utility of unitarily inequivalent representations. As we saw in Section 2.3, algebraic QFT employs inequivalent representations to explain a range of important phenomena from spontaneous symmetry breaking to the existence of antimatter. Fraser argues that Lagrangian QFT needs to impose cutoffs to deal with ultraviolet and infrared divergences, rendering all representations finite-dimensional and thus unitarily equivalent by the Stone–von Neumann theorem. Consequently, even if two models of algebraic and Lagrangian QFT are empirically indistinguishable, the Lagrangian model cannot appeal to the explanatory power of inequivalent representations.

In response, Wallace (2011) notes that all of the explanatory examples Fraser discusses concern inequivalent representations associated with long-distance physics. Although Lagrangian QFT can address infrared divergences by imposing a long-distance cutoff, these are not physically realistic in an infinite spacetime setting. Instead, a better approach is to impose asymptotic boundary conditions by selecting values for certain global quantities (e.g., Kulish & Faddeev, 1970). But results from algebraic QFT like the DHR analysis suggest that this procedure effectively amounts to choosing a particular inequivalent representation to work in.[17] Thus, Wallace agrees with Fraser that techniques from algebraic QFT can help explain the long-distance structure of QFT, but in this context, the algebraic methods are actually compatible with and complimentary to Lagrangian methods.

The short-distance structure of QFT is another matter. Wallace (2006, 2011) argues that Lagrangian and axiomatic QFT represent fundamentally different approaches to dealing with ultraviolet divergences. Central to Wallace's argument is the concept of an *effective field theory*. Given a universal theory describing physics at all length scales, Wilsonian renormalization analysis provides a set of sophisticated coarse graining tools for extracting stable regularities in the long-distance limit. In the best cases, it turns out that these regularities are largely insensitive to the specific details of the underlying theory. What is more, renormalization gives us a kind of explanation for when and why this occurs.[18] So even when we do not know what the underlying universal theory looks like, we can proceed by guessing an effective theory with a floating cutoff and whose interaction terms are renormalizable (i.e., insensitive to the details on how the cutoff is taken). If the predictions of this effective field theory are experimentally confirmed, then we have good reasons to trust it as an approximation in the regime above the cutoff. Wallace contends that the standard model is best viewed as an effective QFT of this kind, whereas algebraic and Wightman QFT aim to capture fundamental theories valid at all length scales. As we have seen, Wightman QFT employs field operators associated with arbitrarily small compact regions. Likewise, the Haag–Kastler axioms usually include an assumption known as *weak additivity*, which entails that the local algebras of arbitrarily small regions are nontrivial. Such assumptions are incompatible with the effective field theory picture.

In one sense, Wallace is right. A major goal of axiomatic QFT is to provide a rigorous description of fundamental QFTs, valid at all length scales. In order to do this, it makes certain assumptions that conflict with an effective field theory viewpoint. But it also has the goal of characterizing in rigorous mathematical terms what all QFTs have in common, and insofar as effective field theories are QFTs, these should be part of the equation. To this end, it is important to distinguish between axioms that are more central to the conceptual framework of Wightman and algebraic QFT (such as covariance, microcausality, and the spectrum condition) and those that serve as auxiliary technical assumptions.

In algebraic QFT, weak additivity is typically viewed as the latter kind of assumption, serving primarily to establish the Reeh–Schlieder property concerning superentanglement in the vacuum state. From a physical standpoint, this property is the consequence of long-range correlations and continues to hold in effective field theories at length scales above the cutoff.[19] In fact, the vast majority of results from algebraic QFT do not rely on assumptions about the ultraviolet regime at all. (Notable exceptions include the field reconstruction results surveyed at the beginning of this section, as well as theorems on the Type III structure of double-cone algebras, e.g., Fredenhagen, 1985.) Furthermore, recent work on scaling algebras (Buchholz & Verch, 1995) and perturbative algebraic QFT (Brunetti, Dütsch, & Fredenhagen, 2009) aims to extend the tools of renormalization group analysis and perturbation theory to the algebraic framework, presenting a possible avenue to formalize the concept of an effective field theory within algebraic QFT.[20] Assumptions about the short distance properties of field operators are far more central in Wightman QFT where they underwrite heavily used analytic continuation methods. Nonetheless, even if one has to give up certain techniques for building models or proving theorems, there is nothing conceptually incoherent about an effective field theory built from Wightman field operators assigned to large regions. It is worth noting that Euclidean construction techniques, which are intimately linked to Wightman QFT via the Osterwalder–Schrader reconstruction theorem are frequently employed in the presence of a cutoff (Jaffe, 2000).[21]

Although the disagreement between Fraser and Wallace extends to other important questions concerning the role of mathematical rigor in physical theories and the explanatory significance of renormalization analysis, from our perspective, the issue at the heart of their debate is the difference between fundamental and effective QFTs, and this is largely orthogonal to the divisions between Lagrangian, Wightman, and algebraic QFT.[22] As we have seen in the course of our brief tour, each approach has different strengths and weaknesses. Lagrangian field theory supplies the raw data, a set of tremendously powerful, if somewhat piecemeal, predictive and explanatory schemas that are central to our understanding of nature at its most fundamental level (even if they are a bit rough around the edges). As interpreters of QFT, our principal goal is to understand these schemas better, and they put significant constraints on our theorizing. Axiomatic approaches give us more precise regimentations of Lagrangian QFT, but we are still in the process of figuring out how they work and if they are fully faithful to the original Lagrangian picture. Wightman field theory has numerous sophisticated tools for building concrete models of QFT. Its main drawbacks are its reliance on localized gauge-dependent field operators. Algebraic QFT stands to provide a more physically transparent, gauge-free description of QFT, but the construction of particular models is very challenging. Therefore, we need to look to the other, more concrete approaches for clues as to how to bring algebraic methods to bear on a wider range of problems, as well as for ideas about possible modifications or extensions of the Haag–Kastler axioms. Arguably, the best way to advance this goal is to approach the problem, from a number of different, complementary angles. As Haag (1996) puts it, “we need a synthesis of the knowledge gained in the different approaches” (p. 326). This is not to say that there is universal harmony amongst mathematicians and physicists working on these various research programs (quite the opposite), but it does seem that as a group, they are more open to the possibility of progress coming from different directions.

To sum things up, we have three complementary, yet partial pictures of QFT, pictures with significant overlap, differential advantages, and no deep incompatibilities. At this stage, it would be premature to cast any of them aside. The moral for philosophers of QFT is this: we should be more open to ideas, puzzles, and techniques from all three approaches. QFT is still evolving, and if philosophers wish to actively participate in its development, we must adopt a more cosmopolitan attitude towards the various conceptual schemas employed by working physicists.

4 | NEW DIRECTIONS

Our field guide concludes with two foundational puzzles located at the intersection of Lagrangian, Wightman, and algebraic QFT. Although the puzzles themselves are decades old, they are the subject of much recent, cutting-edge work in mathematical physics. Moreover, they remain largely unexplored by philosophers of physics and thus represent new avenues for research into the conceptual foundations of QFT.

4.1 | The Bisognano–Wichmann property

In a pair of seminal papers, Bisognano and Wichmann (1975, 1976) uncovered a deep structural property linking local observable algebras and spacetime symmetries. The property itself is rather technical, and its physical basis is not well understood. Nonetheless, it is known to hold in all rigorously constructed models of physical interest and has spurred several important advancements in algebraic and constructive field theory.

As a consequence of the Reeh–Schlieder property, every local observable algebra in the vacuum representation has a pair of special algebraic invariants, the modular conjugation operator, J, and the modular automorphism group, $\{\Delta^{it}\}$. Under the assumption that the observable algebras are generated by Wightman fields, Bisognano and Wichmann (1975, 1976) proved that the modular invariants associated with certain regions of spacetime generate Poincaré transformations. The regions in question are spacelike wedges, infinitely extended spacelike regions bounded above and below by two non-parallel lightlike hypersurfaces. They can be visualized easily in d = 3 Minkowski spacetime as a wedge-shaped region nestled into the side of a light cone (see Figure 1). The modular automorphism group associated with the wedge algebra generates the unique subgroup of wedge-preserving Lorentz boosts, and the modular conjugation operator generates a partial PCT reflection that reverses the direction of time, conjugates all charges, and reflects one spatial direction perpendicular to the edge of the wedge (the x-direction in Figure 1).

Although the original theorem applies to Wightman QFT, it is widely thought that the Bisognano–Wichmann property is a more fundamental consequence of the algebraic structure of observables. Using only translation covariance of the observable net and the spectrum condition, Borchers (1992) proves that in the vacuum representation, the wedge modular invariants must act in the correct geometric fashion except possibly in the spatial direction along the edge of the wedge. Yngvason (1994) constructs several counterexamples with this feature, although none of them satisfy both full Poincaré covariance and the spectrum condition, preserving the possibility that an elementary algebraic proof exists. Haag (1996) conjectures that the Bisognano–Wichmann property follows from the Haag–Kastler axioms plus some natural bound on the local degrees of freedom in QFT such as the *split* or *modular nuclearity* properties.[23] Recent work by Morinelli (2016) supports Haag's conjecture, indicating that the Bisognano–Wichmann property is much weaker than the split property. It is still an open question whether or not the Bisognano–Wichmann property entails that the net of observable algebras is generated by Wightman fields, although the warped deformation constructions surveyed by Lechner (2015) appear to imply a negative answer. Thus, further exploration of the property's algebraic basis has the potential to shed light on the relationship between algebraic and Wightman QFT.

Beyond this, the Bisognano–Wichmann property has a host of important technical and physical applications in algebraic QFT. The interplay between modular structure and spacetime symmetries endows the modular invariants with analyticity properties that can be used in place of assumptions about gauge-dependent Wightman fields. These analyticity properties play a central role in algebraic proofs of the PCT and spin–statistics theorems (Borchers &

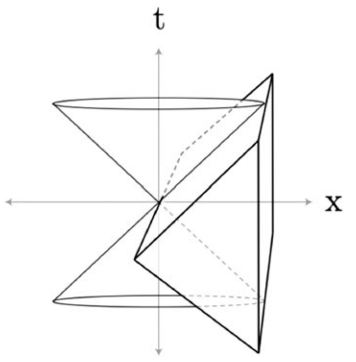

FIGURE 1 A spacelike wedge

Yngvason, 2000) and provide a physical basis for explaining the Unruh effect (Sewell, 1980). In addition, they underwrite new algebraic construction techniques that have produced some of the first examples of interacting four-dimensional QFTs (Summers, 2012). Geometric constraints on wedge modular invariants inspired by the Bisognano–Wichmann property have also been proposed as possible generalizations of the covariance and spectrum axioms applicable to QFTs in curved spacetime (Buchholz & Summers, 1993). (Because generic models of general relativity lack global isometries, the usual versions of these axioms no longer make sense.)

In each of these cases, the Bisognano–Wichmann property stands to provide both mathematical and physical insight into key ideas from Lagrangian and Wightman QFT. There is much in these explanatory applications for philosophers of physics to explore. In addition, there are more elementary questions about the link between modular invariants and spacetime symmetries that deserve attention. No-go results due to Trebels (1997) indicate that in general, we do not expect the modular invariants of other regions or of wedge regions in non-vacuum representations to act geometrically. What makes wedge regions and the vacuum state so special? Why are modular invariants linked to symmetries in the first place?

4.2 | The infrared problem

In QED, the presence of massless particles and long-range forces gives rise to troublesome infrared divergences. These arise when perturbative calculations fail to take into account contributions from low-energy (i.e., *soft*) photons. In QED, as in classical electromagnetism, any accelerating charge will emit electromagnetic radiation. In the quantum theory, this radiation takes the form of a cloud of photons, including an infinite number of soft photons. Because of the long-range nature of forces in QED, each configuration of charges and the associated photon cloud leads to distinct boundary conditions for the electromagnetic field at spatial infinity. In principle, the subtleties of such boundary conditions must be included in any non-perturbative calculation. In practice, however, they can be safely ignored. The standard Lagrangian solution to the infrared problem exploits the fact that any realistic particle detector will only be sensitive to photons with energy above some threshold; hence, measurements of the final scattering state cannot distinguish between different configurations of the soft-photon cloud. By summing over these possible configurations while taking into account quantum corrections associated with the production of virtual photons at all energies, it can be shown that infrared divergences cancel to all orders (see Weinberg, 2005, Chapter 13).

As Strocchi (2013, Chapter 6.4) notes, this practical solution glosses over a number of conceptual difficulties. By imposing a soft photon cutoff, the procedure explicitly breaks Lorentz invariance while making the perturbative calculation dependent on the nature of the relevant measuring device. In addition, it crucially relies on the following mysterious trick: our ignorance of the soft-photon cloud structure effectively amounts to a violation of the positivity of transition probabilities, allowing positive and negative probabilities to cancel out in the end. We, therefore, need both a better understanding of the non-perturbative boundary conditions and a more transparent explanation for the success of the practical Lagrangian solution.

Although progress in the former direction has been slow, the investigation has revealed several important foundational insights. Ultimately, what is needed is an analysis of the global superselection structure of QED akin to the DHR picture discussed in Section 2.3. According to an influential ansatz of Chung (1965), the superselected global states of QED should correspond to coherent states that spontaneously break Lorentz symmetry. This idea forms the basis of subsequent explorations of asymptotic boundary conditions in the Lagrangian and Wightman frameworks (see Kulish & Faddeev, 1970 and Frölich, 1973, respectively). Seeking to adapt this idea to algebraic QFT, Buchholz (1982) starts from the assumption that electrically charged states can be localized in spacelike cones. His analysis yields an uncountable set of superselection sectors labeled by global charge and an asymptotic flux parameter related to the spatial direction of the localization cone. While this represents a promising start, the superselection structure is much more complicated than the DHR picture, and a full account of sector composition and statistics still needs to be developed.[24]

A second obstacle concerns the existence of rigorous scattering theory for QED. Haag and Ruelle originally developed a version of scattering theory for massive particles that avoids the paradoxical implications of Haag's theorem. This was subsequently extended to cover interactions between massless particles by Buchholz (see Buchholz & Summers, 2006 for a summary of these techniques). More recently, Dybalski (2005) has formulated a version of Haag–Ruelle scattering theory for interactions between massive and massless particles. Central to all of these programs, however, is the assumption that particles correspond to irreducible representations of the Poincaré group labeled by mass and spin (an idea that goes back to Wigner). Because charges in QED are always surrounded by an infinite cloud of soft photons, this assumption fails. Particles in QED are so-called *infraparticles* rather than *Wigner particles*. Buchholz (1986) makes this idea precise in the context of Wightman QFT, proving that electrically charged states cannot be sharp eigenstates of the mass operator. (It should be noted that this result is intimately linked to the spontaneous breaking of Lorentz symmetry by coherent global states in QED.) Dybalski and Tanimoto (2011) examine infraparticle interactions in low-dimensional conformal theories; however, a fully general scattering theory for infraparticles remains to be developed.

In light of these difficulties, Buchholz and Roberts (2014) adopt a different strategy, seeking to clarify the physical basis of the practical Lagrangian solution to the infrared problem without giving a full non-perturbative account. Instead of energy threshold constraints, they consider a more general restriction—no matter what, experiments can only measure observables localized in some forward light cone. In QED, all charged particles have mass; hence, the global charge can in principle be determined by such a measurement (i.e., if two representations are unitarily equivalent on some forward light cone, they have the same global charge). Buchholz and Roberts use this idea to divide the unruly menagerie of superselection sectors into equivalence classes called *charge classes* labeled by global charge. They go on to develop an account of antimatter and particle statistics in this new setting with charge classes playing an analogous role to superselection sectors in the DHR analysis. The method works by wiping out certain physical details, namely, the long-range features of soft-photon clouds, which cannot be measured by light cone-localized observables, prompting Buchholz and Roberts to conjecture that their method is essentially equivalent to the usual Lagrangian strategy of summing over unknown soft-photon configurations. If correct, this would represent a major conceptual advancement. Without positing a soft-photon cutoff, the charge class explanation preserves Lorentz invariance and does not hinge on the details of a particular detection device. Moreover, it avoids the bizarre computational crutch of negative transition probabilities. A comparison of the charge class hypothesis with recent work on the infrared problem in Lagrangian QFT by Steinmann (2000) might help settle this conjecture.

As with the Bisognano–Wichmann property, current work on the infrared problem draws upon tools and techniques from all three formulations of QFT. There is a host of similar foundational problems situated at the nexus of these frameworks that mathematically oriented philosophers of physics are well situated to make progress on. A partial list includes anomalous symmetry breaking, the Wigner particle concept, rigorous quantization methods, renormalization, color confinement, statespace nuclearity properties, the spin–statistics connection, and the role of Gauss's law in local gauge theories. We have much to gain from a broadened worldview.

ACKNOWLEDGEMENTS

I would like to thank David Baker, Detlev Buchholz, Hans Halvorson, and an anonymous referee along with audiences at the University of Delaware, Carnegie Mellon, and the University of Minnesota for helpful comments and criticism on early versions of this paper.

ENDNOTES

[1] See Weinberg (2005) and Peskin and Schroeder (1995) for standard treatments. Duncan (2012) explores a number of foundational issues such as Haag's theorem that other authors working in the Lagrangian tradition frequently gloss over.

[2] See Rovelli (2014) for a dissenting opinion. Healey (2007) provides a comprehensive introduction to the philosophical problems surrounding gauge symmetries in QFT.

[3] Although this is the view of most mathematical physicists, Wallace (2006, 2011) argues that modern renormalization methods have put Lagrangian QFT on reasonably secure mathematical footing.

[4] More mathematically rigorous quantization procedures exist such as geometric quantization and deformation quantization, although these take their cue from the mathematics of constructive and algebraic QFT and are very difficult to implement in concrete models. A more thorough understanding of various quantization techniques, their domains of applicability, and how they are interrelated is a major open problem in the foundations of quantum theory. See Henneaux and Teitelboim (1992) and Landsman (1998) for discussions of more rigorous quantization methods.

[5] Let F be some physical quantity of interest written as a power series in terms of the electromagnetic charge e:

$$F(e^2) = a_0 + a_1 e^2 + a_2 e^4 \ldots$$

If this series converges for some positive value of e, then it is analytic at $e = 0$ so both $F(e^2)$ and $F(e^{-2})$ should converge for sufficiently small values of e. Dyson interprets $F(e^{-2})$ as describing the same quantity in a fictitious world, where like charges attract and opposites repel. In this world, by creating a large number of pairs of oppositely charged particles from the vacuum and separating them into clusters of like charge, it is possible to indefinitely lower the energy. Because of quantum tunneling effects, the vacuum is no longer a stable ground state. It is thermodynamically favored to decay into a roiling sea of particles and antiparticles with infinitely negative energy. This pathological behavior means that $F(e^{-2})$ cannot converge, so neither can $F(e^2)$.

[6] Extant triviality theorems employ lattice regularization techniques to control the ultraviolet limit. Because the asymptotic behavior of the renormalized perturbation series can be extremely sensitive to the choice of regularization, Rivasseau (1991, Chapter III.3.I) argues that these results lack full generality. Fernandez et al. (1992, Chapter 15.6.1) discusses the possibility that a finite-field renormalization scheme might evade the triviality theorems, although they conclude that this scenario, which has no connection to standard renormalization group theory, is unlikely. A third possibility is that a ϕ^4 theory coupled to a local gauge theory might somehow escape triviality (a tantalizing prospect given the role of the ϕ^4 interaction in the standard model Higgs mechanism).

[7] A third axiomatic approach, functorial QFT, also deserves mention. Functorial QFT aims to provide axioms characterizing the Feynman path integral as a functor from a category of spacetime cobordisms into a category of Hilbert spaces (Baez, 2006). So far mathematicians have only figured out how to implement this idea in the special cases of topological and conformal QFTs; however, there is interesting overlap between functorial QFT and recent proposals for extending algebraic QFT to curved spacetime (Fredenhagen & Rejzner, 2016).

[8] The Wightman axioms were first proposed by Wightman and Gårding in the 1950s. Streater and Wightman (1989) and Jost (1965) are classic sources.

[9] In addition, there are several technical assumptions concerning the domain of definition of the unbounded field operators and the cyclicity of the vacuum state. See Streater and Wightman (1989, Chapter 3.1) for a full list of axioms.

[10] Although the ubiquity of analytic continuation arguments poses a special interpretive challenge for Wightman QFT, it should be noted that similar methods are also widely employed (albeit with less mathematical rigor) in Lagrangian QFT. Though they are less common in algebraic QFT, they are crucial for establishing a number of central results related to modular invariants and the Bisognano–Wichmann property discussed in Section 4.1. Thus, the question of their physical significance cannot be completely sidestepped by moving to a different framework.

[11] These theorems emerged over a series of several papers. See Strocchi (2013, Chapter 7.3) for references. Horuzhy (1990) also contains a succinct overview of the Strocchi no-go theorems in the appendix.

[12] The locus classicus is Haag (1996). For a compact, technical introduction, see Brunetti and Fredenhagen (2004). Halvorson and Müger (2006) is a thorough mathematical treatment of the subject from a philosophical viewpoint. Ruetsche (2011) focuses on the problem of inequivalent representations and associated interpretive challenges.

[13] A C^*-algebra is isomorphic to a subalgebra of bounded linear operators acting on some Hilbert space. Such algebras are closed under addition, multiplication, and complex-scalar multiplication. (In quantum theories, the multiplication operation

is noncommutative.) Additionally, the algebra is equipped with a canonical involution mapping (the abstract analogue of the conjugate transpose operation for matrices) and a norm. The algebra must be closed in the associated norm topology. Von Neumann algebras, which also play a major role in algebraic QFT, are C^*-algebras with a complete set of projection operators.

14 As in Wightman QFT, the uniqueness assumption can be dropped in order to model theories with multiple vacuum states.

15 The Wightman field operators act on a separable Hilbert space $\mathcal{H}$. Under the action of the reconstructed gauge group, G, the Hilbert space can be decomposed into a direct sum of irreducible representations of G, $\mathcal{H}=\oplus\mathcal{H}_\sigma$. These representations are in 1–1 correspondence with the original algebraic superselection sectors, that is, unitary quivalence classes of representations of the net of observable algebras satisfying the DHR selection criteria. The Casimir invariants, σ, associated with G correspond to the global charge operators whose values label each superselection sector.

16 The main technical hurdle currently blocking progress in this direction is known as the Gribov problem. See Rivasseau (1991, Chapter III.5.F) for a discussion.

17 Kulish and Faddeev (1970) explore this idea in the context of quantum electrodynamics, where the connection to algebraic results has yet to be made precise. A full account of superselection theory in QED requires a solution to the infrared problem, which we discuss in Section 4.2.

18 Fraser (2011) objects to this point, arguing that many of these explanations rest on analogies with solid-state physics that may not carry over to the particle physics case. See Butterfield and Bouatta (2015) for a recent survey of philosophical issues connected to renormalization.

19 Wallace (2006) discusses the physical ramifications of the Reeh–Schlieder property from the perspective of effective Lagrangian QFT. From the algebraic perspective, weak additivity usually requires that (in the vacuum GNS representation) the global algebra can be generated by spacetime translations of the local algebra associated with *any region O*. It is only this last bit that conflicts with the effective field theory picture. We can easily replace it with the assumption that only translations of algebras associated with regions larger than some cutoff generate the global algebra. The standard proofs of the Reeh–Schlieder property can all be carried out with this modified assumption. See, for example, Horuzhy (1990) Theorems 1.3.1 and 1.3.2.

20 The exact relationship between perturbative and traditional algebraic QFT is still an open question. The main difference is that the local algebras in the perturbative approach are not C^*-algebras, but rather more general *-algebras whose elements are sequences of formal power series

$$\sum_{n=0}^{\infty} A_n t^n$$

with coefficients A_n in some *-algebra, $\mathcal{A}$. Consequently a state over these algebras generates a power series rather than a complex number. As in standard Lagrangian perturbation theory, we must take a partial sum of this series to get an expectation value.

21 QFTs with a built-in ultraviolet cutoff also manifestly violate the covariance and microcausality axioms; however, this is usually viewed as an artifact of choosing a particular cutoff scheme. Though it remains an open possibility that the underlying fundamental theory might violate these central axioms, they should remain approximately valid for the effective theory and a full explanation of the limiting relation between the two should explain why.

22 To be fair, by the end of their exchange, both Wallace and Fraser come to recognize this. Wallace (2011) specifically aims his sights at axiom systems “crucially including the assumption that quantum fields [...] can be associated with arbitrarily small open subsets of spacetime” (p. 117). But as we have seen, such assumptions are far less central to the conceptual backbone of algebraic and Wightman QFT than both Wallace and Fraser suggest.

23 These properties are related to the existence of product states across regions separated by some finite spacelike distance. Although they are not part of the standard Haag–Kastler axioms, they are widely viewed as a necessary constraint on any QFT with emergent particle structure and are playing an increasingly central role in current research in algebraic QFT. See Haag (1996, Chapter V.5) for a precise formulation and discussion of the physical consequences of these properties.

24 There is an important difference between this analysis and the generalized DHR picture developed by Buchholz and Fredenhagen (1982), which also begins with the assumption that charged states can be localized in spacelike cones. The Buchholz-Fredenhagen picture applies to purely massive theories with topological charges. In this case, representations with similar charges localized in different cones are all unitarily equivalent, i.e., the localization cone is “arbitrary.” This simplifies the superselection structure and enables a full extension of the DHR results. In theories with massless particles, like QED, the direction of the localization cone corresponds to a superselected observable quantity, greatly complicating matters. We need a better mathematical understanding of how charges localized in different cones are related before we can replicate the DHR picture in this setting.

It remains possible that charged states in QED might have better localization properties relative to some natural subalgebra of the full algebra of observables. In this case we could use standard DHR tools for a limited range of observations. This is true in classical electrodynamics and non-interacting versions of QED, where electric charges are localized in compact regions relative to the subalgebra generated by the electric current. Intriguingly, this turns out to be false for interacting QED. Using rigorous perturbative results established by Steinmann (2000), Buchholz, Doplicher, Morchio, Roberts, and Strocchi (2001) prove that electrically charged states in QED cannot be localized in this manner. The observed delocalization in the full theory must arise from a combination of quantum effects and the presence of interactions. The physical implications of this fact still need to be explored.

WORKS CITED

Baez, J. (2006). Quantum quandaries: A category-theoretic perspective. In D. Rickles, S. French, & J. Saatsi (Eds.), *The structural foundations of quantum gravity*. Oxford: Clarendon Press.

Baker, D. J., & Halvorson, H. (2010). Antimatter. *British Journal for the Philosophy of Science*, *61*, 93–121.

Balaban, T. (1989). Large field renormalization II. Localization, exponentiation, and bounds for the R operation. *Communications in Mathematical Physics*, *122*, 355–392.

Bisognano, J., & Wichmann, E. (1975). On the duality condition for a Hermitian scalar field. *Journal of Mathematical Physics*, *16*, 985–1007.

Bisognano, J., & Wichmann, E. (1976). On the duality condition for quantum fields. *Journal of Mathematical Physics*, *17*, 303–321.

Borchers, H. J., & Yngvason, J. (2000). On the PCT—theorem in the theory of local observables. http://arxiv.org/abs/math-ph/0012020

Borchers, H. J. (1992). The CPT-theorem in two-dimensional theories of local observables. *Communications in Mathematical Physics*, *143*, 315–332.

Borchers, H. J., & Yngvason, J. (1992). From quantum fields to local von Neumann algebras. *Reviews in Mathematical Physics*, *4*, 15–47.

Bostelmann, H. (2005). Phase space properties and the short distance structure in quantum field theory. *Journal of Mathematical Physics*, *46*, 52301–52321.

Brunetti, R., & Fredenhagen, K. (2004). The algebraic approach to quantum field theory. http://arxiv.org/abs/math-ph/0411072

Brunetti, R., Dütsch, M., & Fredenhagen, K. (2009). Perturbative algebraic quantum field theory and the renormalization group. *Advances in Theoretical and Mathematical Physics*, *13*, 1–56.

Buchholz, D. (1982). The physical statespace of quantum electrodynamics. *Communications in Mathematical Physics*, *85*, 49–71.

Buchholz, D. (1986). Gauss's law and the infraparticle problem. *Physics Letters B*, *174*, 331–334.

Buchholz, D., & Fredenhagen, K. (1982). Locality and the structure of particle states. *Communications in Mathematical Physics*, *84*, 1–54.

Buchholz, D., & Roberts, J. E. (2014). New light on infrared problems: Sectors, statistics, symmetries and spectrum. *Communications in Mathematical Physics*, *330*, 935–972.

Buchholz, D., & Summers, S. J. (1993). An algebraic characterization of vacuum states in Minkowski space. *Communications in Mathematical Physics*, *155*, 449–458.

Buchholz, D., & Summers, S. J. (2006). Scattering in relativistic quantum field theory: Fundamental concepts and tools. In *Enclyclopedia of mathematical physics*. Amsterdam: Elsevier.

Buchholz, D., & Verch, R. (1995). Scaling algebras and renormalization group in algebraic quantum field theory. *Reviews in Mathematical Physics*, *7*, 1195–1239.

Buchholz, D., Doplicher, S., Morchio, G., Roberts, J. E., & Strocchi, F. (2001). Quantum delocalization of the electric charge. *Annals of Physics*, *290*, 53–66.

Butterfield, J., & Bouatta, N. (2015). Renormalization for philosophers. In T. Bigaj, & C. Wütrich (Eds.), *Metaphysics in contemporary physics*. Leiden: Brill.

Chung, V. (1965). Infrared divergences in quantum electrodynamics. *Physical Review B*, *140*, 1110–1122.

Doplicher, S., & Roberts, J. E. (1972). Fields, statistics and non-Abelian gauge groups. *Communications in Mathematical Physics*, *28*, 331–348.

Doplicher, S., Haag, R., & Roberts, J. E. (1969a). Fields, observables and gauge transformations. I. *Communications in Mathematical Physics*, *13*, 1–23.

Doplicher, S., Haag, R., & Roberts, J. E. (1969b). Fields, observables and gauge transformations. II. *Communications in Mathematical Physics*, *15*, 173–200.

Duncan, A. (2012). *The conceptual framework of quantum field theory*. Oxford: Oxford University Press.

Dybalski, W. (2005). Haag–Ruelle scattering theory in presence of massless particles. *Letters in Mathematical Physics*, *72*, 27–38.

Dybalski, W., & Tanimoto, Y. (2011). Asymptotic completeness in a class of massless relativistic quantum field theories. *Communications in Mathematical Physics*, *305*, 427–440.

Dyson, F. (1952). Divergence of perturbation theory in quantum electrodynamics. *Physical Review*, *85*, 631–632.

Earman, J., & Fraser, D. (2006). Haag's theorem and its implications for the foundations of quantum field theory. *Erkenntnis*, *64*, 305–344.

Fernandez, R., Fröhlich, J., & Sokal, A. (1992). *Random walks, critical phenomena, and triviality in quantum field theory*. New York: Springer-Verlag.

Fraser, D. (2008). The fate of 'particles' in quantum field theories with interactions. *Studies in History and Philosophy of Modern Physics*, *39*, 841–859.

Fraser, D. (2009). Quantum field theory: Underdetermination, inconsistency and idealization. *Philosophy of Science*, *76*, 536–567.

Fraser, D. (2011). How to take particle physics seriously: A further defense of axiomatic quantum field theory. *Studies in History and Philosophy of Modern Physics*, *42*, 126–135.

Fredenhagen, K. (1985). On the modular structure of local algebras of observables. *Communications in Mathematical Physics*, *97*, 79–89.

Fredenhagen, K., & Hertel, J. (1981). Local algebras of observables and pointlike localized fields. *Communications in Mathematical Physics*, *80*, 555–561.

Fredenhagen, K., & Rejzner, K. (2016). Quantum field theory on curved spacetime: Axiomatic framework and examples. *Journal of Mathematical Physics*, *57*, 31101–31139.

Fredenhagen, K., Rheren, K.-H., & Schroer, B. (1989). Superselection sectors with braid group statistics and exchange algebras. *Communications in Mathematical Physics*, *125*, 201–226.

Frölich, J. (1973). On the infrared problem in a model of scalar electrons and massless, scalar bosons. *Annales de l'Institute Henri Poincare*, *A19*, 1–103.

Glimm, J., & Jaffe, A. (1987). *Quantum physics: A functional integral point of view*. New York: Springer-Verlag.

Haag, R. (1996). *Local quantum physics*. Berlin: Springer-Verlag.

Halvorson, H., & Müger, M. (2006). Algebraic quantum field theory. In J. Butterfield, & J. Earman (Eds.), *Philosophy of physics* (pp. 731–922). Amsterdam: Elsevier.

Healey, R. (2007). *Gauging what's real*. Oxford: Oxford University Press.

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

Horuzhy, S. S. (1990). *Introduction to algebraic quantum field theory*. Dordrecht: Kluwer Academic Publishers.

Jaffe, A. (2000). Constructive quantum field theory. In *Mathematical physics 2000*). Singapore: World Scientific.

Jost, R. (1957). Eine bemerkung zum CTP theorem. *Helvetica Physica Acta*, *30*, 409–416.

Jost, R. (1965). *The general theory of quantized fields*. Providence: American Mathematical Society.

Kulish, P., & Faddeev, L. (1970). Asymptotic conditions and infrared divergences in quantum electrodynamics. *Theoretical and Mathematical Physics*, *4*, 745–757.

Landsman, N. P. (1998). *Mathematical topics between classical and quantum mechanics*. New York: Springer-Verlag.

Lechner, G. (2015). Algebraic constructive quantum field theory: Integrable models and deformation techniques. In *Advances in algebraic quantum field theory Mathematical Physics Studies*. Cham: Springer.

Morinelli, V. (2016). An algebraic condition for the Bisognano–Wichmann property. https://arxiv.org/abs/1604.04750

Peskin, M., & Schroeder, D. (1995). *An introduction to quantum field theory*. Boulder: Westview Press.

Rehren, K.-H. (1992). Field operators for anyons and plektons. *Communications in Mathematical Physics*, *145*, 123–148.

Rivasseau, V. (1991). *From perturbative to constructive renormalization*. Princeton: Princeton University Press.

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

Ruetsche, L. (2011). *Interpreting quantum theories*. Oxford: Oxford University Press.

Ruetsche, L. (2012). Philosophical aspects of quantum field theory I and II. *Philosophy Compass*, *7*, 559–584.

Sewell, G. (1980). Relativity of temperature and the hawking effect. *Physics Letters A*, *79*, 23–24.

Steinmann, O. (2000). *Perturbative quantum electrodynamics and axiomatic field theory*. Berlin: Springer.

Streater, R. F., & Wightman, A. (1989). *PCT, spin and statistics, and all that*. New York: Addison-Wesley.

Strocchi, F. (1978). Local and covariant gauge quantum field theory. *Physical Review D*, *17*, 2010–2021.

Strocchi, F. (2013). *An introduction to non-perturbative foundations of quantum field theory*. Oxford: Oxford Science Publications.

Summers, S. J. (2012). A perspective on constructive quantum field theory. http://arxiv.org/abs/1203.3991

Trebels, S. (1997). Über die geometrische wirkung modularer automoprhismen. Ph.D. thesis, Göttingen.

Wallace, D. (2006). In defence of naiveté: On the conceptual status of Lagrangian quantum field theory. *Synthese*, *151*, 33–80.

Wallace, D. (2009). QFT, antimatter, and symmetry. *Studies in History and Philosophy of Modern Physics*, *40*, 209–242.

Wallace, D. (2011). Taking particle physics seriously: A critique of the algebraic approach to quantum field theory. *Studies in History and Philosophy of Modern Physics*, *42*, 116–125.

Weinberg, S. (2005). *The quantum theory of fields* (Vol. I). Cambridge: Cambridge University Press.

Yngvason, J. (1994). A note on essential duality. *Letters in Mathematical Physics*, *31*, 127–141.

AUTHOR BIOGRAPHY

Noel Swanson is an assistant professor of philosophy at the University of Delaware, where he works on foundational problems in physics, mathematics, and general philosophy of science. He earned his PhD from Princeton University in 2014 under the direction of Hans Halvorson.

How to cite this article: N Swanson. A philosopher's guide to the foundations of quantum field theory. *Philosophy Compass*. 2017;12:e12414. https://doi.org/10.1111/phc3.12414