

Taking particle physics seriously: a critique of the algebraic approach to quantum field theory

David Wallace*

October 1, 2010

Abstract

I argue against the currently-prevalent view that algebraic quantum field theory (AQFT) is the correct framework for philosophy of quantum field theory and that “conventional” quantum field theory (CQFT), of the sort used in mainstream particle physics, is not suitable for foundational study. In doing so, I defend that position that AQFT and CQFT should be understood as rival programs to resolve the mathematical and physical pathologies of renormalization theory, and that CQFT has succeeded in this task and AQFT has failed. I also defend CQFT from recent criticisms made by Doreen Fraser.

1 Prelude: a fable

Once upon a time there was a community of physicists. This community believed, and had good reason to believe, that quantum mechanics, not classical mechanics, was the right framework in which to do physics. They also had a good understanding, at the classical level, of the dynamics of solid bodies (vibrations in crystals, for instance): they knew, for example, that some such bodies could be analysed using Lagrangians like

$$\mathcal{L} = \frac{1}{2}\dot{\phi}^2 - \frac{1}{2}(\nabla\phi)^2 + (\text{higher terms}) \quad (1)$$

where $\phi(x)$ is the displacement of the part of the crystal which is at position x at equilibrium.

But the physicists were sad, because they knew nothing at all about the microscopic structure of matter, and so they did not have a good *quantum* theory of vibrations in crystals or of other solid-matter dynamics.

So one day, they set out to quantize their classical theories of solid matter. At first, they tried to do it naively, by putting the classical theory into

*Balliol College, Oxford; and Philosophy Faculty, University of Oxford. Email: david.wallace@balliol.ox.ac.uk

Hamiltonian form and replacing classical observables with self-adjoint operators. This worked quite well until the higher-order terms in (1) were included. But when the physicists tried to include those higher order terms, the theory became mathematically very badly behaved — all the calculations contained integrals that diverged to infinity.

Soon the physicists discovered that they could extract working calculational results if they just assumed that displacements couldn't vary on arbitrarily short lengthscales. This amounted to “cutting off” the range of integration in the divergent integrals, so that they got a finite result. When they did their calculations this way, the answers agreed very well with experiment.

But the physicists were still sad. “It's ad hoc”, they said. “It's inelegant”, they lamented. “It conflicts with the Euclidean symmetries of solid matter”, they cried.

So they went back to basics, and looked for an axiomatised, fully rigorous quantum theory, with displacements definable on arbitrarily short lengthscales and with exact Euclidean symmetries.

And to this day, they are still looking.

2 Introduction

Quantum field theory (QFT) certainly seems like an appropriate subject for philosophy of physics. It deserves to be called *the* great success story of post-war theoretical physics, with applications ranging far beyond its original home in particle physics and encompassing an increasingly large part of condensed-matter physics and statistical mechanics; if gravitational phenomena (and, perhaps, the quantum measurement problem) are set aside, it appears to provide a satisfactory underpinning for the whole of physics; its empirical successes include predictions which are confirmed to double-figure numbers of significant figures.

And indeed, QFT has had a significant, and increasing, degree of attention from philosophers of physics in recent years. But in much of this work¹, it is not really clear that what these philosophers are studying is the same quantum field theory which has achieved all these theoretical triumphs. For the dominant framework in which “philosophy of quantum field theory” is conducted is so-called *algebraic quantum field theory* (AQFT), the result of a program to put QFT on a rigorous footing which began with the axiomatic QFT developed by Wightman and co-workers in the 1960s, and which has been extensively studied by mathematical physicists in the interim. (For a reasonably up-to-date account of AQFT, see Haag (1996).)

Philosophers who work on AQFT are, naturally, quite aware of the difference between AQFT and “what the physicists do”. Halvorson (2007) notes explicitly, in the introduction to a monograph on AQFT, that

¹Examples include Baker (2009), Baker and Halvorson (2009), Clifton and Halvorson (2001), Earman (2004), Earman and Fraser (2006), Fraser (2008), Halvorson (2007), Ruetsche (2002), and Ruetsche (2003).

From the title of this Chapter [Algebraic Quantum Field Theory], one might suspect that the subject is some idiosyncratic approach to quantum field theory (QFT). The approach is indeed idiosyncratic in the sense of demographics: only a small proportion of those [physicists] who work on QFT work on algebraic QFT (AQFT). However, there are particular reasons why philosophers, and others interested in foundational issues, will want to study the algebraic approach.

What are these “particular reasons”? Basically: that AQFT is on a mathematically very firm footing, and that “standard” QFT (let me call it ‘conventional quantum field theory, or CQFT’²) is mathematically far too ill-defined for foundational study. As Halvorson goes on to say,

philosophers of physics have taken their object of study to be theories, where theories correspond to mathematical objects (perhaps sets of models). But it is not so clear where “quantum field theory” can be located in the mathematical universe. In the absence of some sort of mathematically intelligible description of QFT, the philosopher of physics has two options: either find a new way to understand the task of interpretation, or remain silent about the interpretation of quantum field theory.

It is for this reason that AQFT is of particular interest for the foundations of quantum field theory. In short, AQFT is our best story about where QFT lives in the mathematical universe, and so is a natural starting point for foundational inquiries.

In a previous paper (Wallace 2006), I argued that in fact CQFT can be put on a mathematical footing secure enough to do perfectly adequate foundational work. In this paper, I want to review that argument, and then to make explicit the too-tacit implication of that paper: that *given that* CQFT is foundationally adequate, there is really no reason to regard AQFT as a ‘natural starting point’ — or as any starting point at all — for foundational enquiries into QFT.

The structure of the paper is as follows. In section 3, I make the case that CQFT and AQFT should be seen as rival research programs, attempting in different ways to solve the problem of infinities in QFT. In sections 4 and 5 I explain how CQFT and AQFT, respectively, set out to solve that problem; in sections 6–7 I compare their respective successes and conclude (a) that CQFT is decisively more successful, and (b) its success undermines reasons for being interested in AQFT. In sections 8–9, I consider possible objections to this conclusion — in particular (in section 8), I critically assess a recent paper by Doreen Fraser (Fraser 2009) where she defends the strategy of using exclusively AQFT for foundational work on quantum field theory. Section 10 is the conclusion.

A disclaimer before I begin. I’m *obviously* not arguing that quantum mechanics can’t be given an algebraic formulation: that’s a straightforward mathematical fact. Nor am I arguing that algebraic *methods* do not sometimes —

²In Wallace (2006) I called it ‘Lagrangian quantum field theory’; I now feel that name was a little misleading.

do not frequently — give insight into the behaviour of quantum systems with infinitely many degrees of freedom: of course they do. By “algebraic quantum field theory”, I don’t simply mean “quantum field theory, studied by algebraic methods”.

So what do I mean by “AQFT”? The short answer is “quantum field theory, formulated via something like the axioms in Streater and Wightman (1964), Haag (1996) or Halvorson (2007), *crucially* including the assumption that quantum fields are operator- (or algebra-element-)valued distributions, so that field operators (or algebra elements) can be associated with arbitrarily small open subsets of spacetime.” But to get a more informative answer, it will be useful to review just why the algebraic approach developed in the first place.

3 AQFT and CQFT as rival programs

As is widely known,³ the original attempts — by Dirac, Jordan, Heisenberg and others — to develop quantum field theory foundered on the problem of infinities: attempts to calculate physical quantities just led to divergent integrals. Feynman, Tomonaga, Schwinger and Dyson sidestepped the problem by inventing a mathematical method — so-called *renormalization* — to remove the infinities. In doing so, they made quantum electrodynamics into a predictively powerful framework whose accuracy is still unsurpassed — but renormalization made little mathematical *or physical* sense. In this situation — and given the difficulties of extending quantum-field-theoretic methods beyond quantum electrodynamics — it was entirely reasonable for a group of mathematical physicists to try to put the subject on solid conceptual footing by writing down sensible axioms for a quantum field theory, and then looking for models for those axioms. As Streater and Wightman say in the introduction to their classic (1964) text,

[T]he quantum theory of fields never reached a stage where one could say with confidence that it was free from internal contradictions — nor the converse. In fact, the Main Problem of quantum field theory turned out to be to kill it or cure it: either to show that the idealizations involved in the fundamental notions of the theory (relativistic invariance, quantum mechanics, local fields, etc.) are incompatible in some physical sense, or to recast the theory in such a form that it provides a practical language for the description of elementary particle dynamics.

As of 1964, then, it would not have been unreasonable to see quantum field theory as being in the same foundational state as quantum mechanics in the late 1920s (that is, in the form proposed by Dirac (1930), with his extensive

³In this section, I greatly simplify the actual history of quantum field theory, which is Byzantine in its complexity. My purpose is more pedagogical than historical, and only the broadest brush-strokes will be significant for my purposes. For a much more historically careful account, see Cao (1997).

use of position eigenstates, delta functions, and the like) — predictively powerful, but resting on mathematically ill-defined notions — and to see Streater, Wightman *et al* as attempting to play von Neumann (or perhaps Gelfand)⁴ to Feynman/Tomonaga/Schwinger/Dyson’s Dirac. And as of 2009, philosophers of physics seem to regard the situation as essentially unchanged: conventional quantum field theory remains conceptually ill-formed; the axiomatic approach (and its algebraic successor) remains the natural way to solve this problem; if progress has been slow, this just reflects the intrinsic difficulty of the subject.

But forty-five years have passed. And they have seen theoretical physicists (notably Kenneth Wilson and John Kogut, in the early 1970s) approach the problem of renormalization from a very different direction. In doing so, these physicists made assumptions which directly contradicted some of the basic assumptions of the AQFT program: crucially, they assumed that instead of being definable on arbitrarily small spacetime regions, quantum field theory would break down at some short lengthscales. They worked out the general structure of a quantum theory with such a short-distance breakdown, and in doing so they not only put renormalization on a sound theoretical footing, but produced a general explanatory framework which yielded rich results both in particle physics and in condensed matter physics. (Because the short-distance breakdown of QFT used in this model is often called the ‘cutoff’, CQFT can be taken to stand for *cutoff quantum field theory* as readily as *conventional quantum field theory*.)

Given these more recent developments, it is no longer appropriate — if ever it was — to see AQFT as the proposed mathematically-rigorous version of CQFT. The two are better understood as *rival research programs*, trying in different ways to resolve the problem of renormalization. They make incompatible, and in principle empirically testable, physical assumptions. And — as I will argue in the next three sections — judged as research programs, one of those programs has basically been unsuccessful, and one has been fantastically successful. No prizes for guessing which is which.

4 Understanding renormalization in CQFT

One of the revolutions in post-1960s quantum field theory has been the application of QFT methods to condensed-matter physics, and the flow of ideas in both directions between condensed-matter and particle physics. In particular, ideas from condensed-matter physics strongly influenced Wilson and Kogut in their analysis of renormalization, and the condensed-matter approach to renormalization will be a useful starting point for our purposes.

In (real, not fabled!) condensed-matter physics,⁵ solid bodies (metals, crystals and the like) look like continuous systems — i. e. , fields — on large scales,

⁴von Neumann (1955) reformulated quantum mechanics in a way which made no use of position eigenstates or delta functions; Gelfand *et al* (1964) gave a rigorous mathematical reformulation of both.

⁵Here I mostly follow Binney *et al* (1992).

and so formally it is possible to apply QFT methods to them in just the same way as for “real” field theories. Just as in particle physics, though (and just as in the fable), this leads to integrals which diverge at short distances. In solid state physics, though, there is no particular reason to be surprised by this: the system being studied obviously cannot be treated as a continuum at arbitrarily short lengthscales, since on atomic scales matter has a discrete structure, so no wonder that the theory produces pathological results if we assume that it can be so treated.

This suggests that the divergent integrals should be evaluated not up to infinite momenta (that is, down to zero lengthscales) but should instead be ‘cut off’ at a lengthscale of $\sim 10^{-10}\text{m}$. But:

- Any (remotely simple) way of implementing this cutoff will inevitably be very crude and inexact: the way in which the approximation of treating the system as a continuum breaks down will realistically be much more complicated than any simple cutoff can capture.
- This ought to *matter*, because the short-distance interactions are much stronger than the longer-distance ones, and so we are crudely oversimplifying the dynamics exactly in the place where we would expect them to be most important.

This is where modern renormalization theory comes in. It establishes that the details do not after all matter for many purposes — because if we choose to implement the short-distance cutoff in two different ways, the only effects significant at lengthscales much longer than the cutoff lengthscales are changes in the coefficients of *finitely many* additional interaction terms, known as the *renormalizable* interactions. (Other interactions, the “nonrenormalizable” interactions, may be present but will become arbitrarily weak on lengthscales sufficiently long with respect to the cutoff lengthscale.)

For instance, in scalar condensed-matter systems (that is, systems where we can treat the relevant degree of freedom as a scalar field on long distance scales), these terms are

$$\phi^2 \quad \dot{\phi}^2 \quad (\nabla\phi)^2 \quad \phi^4. \tag{2}$$

So — as long as we are interested in the behaviour of the system on lengthscales long compared with the cutoff length — we can absorb all our ignorance of how the cutoff is implemented into the values of *finitely many* coefficients which can be measured empirically. To be sure, if we want to calculate those coefficients from first principles (from atomic quantum mechanics, say) then we will need to engage with all the messy details of the cutoff — indeed, in practice we will have to stop using quantum-field-theoretic methods altogether. But as long as we are happy to treat them as empirical parameters, the details of the cutoff process are irrelevant.

To reiterate: the moral of the story is that condensed-matter QFT, on lengthscales $\gg 10^{-10}\text{m}$, is totally insensitive to the details of the short-distance physics, except insofar as it determines the numerical values of the coefficients.

All we need is *some kind* of freezing out of the short distance degrees of freedom — exactly how that freezing out is implemented is irrelevant.

But if this works in condensed-matter QFT, nothing prevents us telling *exactly the same story* in particle physics, provided only that something freezes out the short-distance degrees of freedom on some lengthscale far below what current experimental physics can probe. This might be done, for instance, by

- Another field theory (though of course, that field theory in turn will need some kind of short-distance cutoff)
- A non-field theory, without its own infinities (string theory is expected by its proponents to be such a theory)
- A real lattice structure to space (seems inelegant, but who knows?)
- A discretisation of spacetime of the kind constructed in loop quantum gravity (as described by, e. g., Rovelli (2004)) — much more elegant than the real lattice route, and (it is hoped) relativistically covariant.
- Some as-yet-unimagined solution.

Renormalization theory itself tells us that if there *is* a short-distance cutoff, large-scale phenomenology will give us almost no information about its nature. As in condensed-matter physics, the only effect of changing the form of the cutoff is to rescale the parameters in the theory's Lagrangian.

We are, to be sure, more limited in particle physics than in condensed-matter quantum field theory. In the latter case, we have both the large-scale theory — quantum field theory — and the small-scale theory — the quantum mechanics of the atomic constituents. In the former case, we have only the large-scale theory. This means, in particular, that while the values of the renormalized parameters can be calculated (in principle) within condensed-matter physics, they can only be measured empirically within particle physics. But in particle physics, as in condensed-matter physics, the *only* significance at the large scale of the microphysical details — and, in particular, the only significance of the short-distance cutoff — is in setting the values of those parameters.

This, in essence, is how modern particle physics deals with the renormalization problem:⁶ it is taken to presage an ultimate failure of quantum field theory at some short lengthscale, and once the bare existence of that failure is appreciated, the whole of renormalization theory becomes unproblematic, and indeed predictively powerful in its own right.

Before considering that predictive power, though, we should look at the competition. How does AQFT deal with renormalization?

⁶I don't intend my argument in this paper to be sociological: it is no part of my thesis that *working physicists* (by which I mean, roughly: those who use QFT for concrete calculations) all think about renormalization this way. But for what it's worth, my anecdotal impression is that most do: or rather, most of those who worry about renormalization in the first place do. (As in discussions of the quantum measurement problem, a sizable fraction of physicists take a 'shut up and calculate' attitude.)

5 Understanding renormalization in AQFT

As I have already noted, algebraic quantum field theory deals with the paradoxes of renormalization by going back to first principles. Instead of writing down concrete examples of quantum field theories (using the admittedly deeply unattractive method of starting with a classical field and then “quantizing” it), which associated field operators to every spacetime point, and then trying to modify those theories to get rid of the infinities, the AQFT theorist writes down a formal mathematical framework which ought to be satisfied by any mathematically acceptable quantum field theory.

Crucially, that framework includes the assumption that quantum fields, though they cannot be defined *at points* without leading to mathematical pathologies, can be defined on arbitrarily small spacetime regions. (In the earlier axiomatic approach, this amounted to replacing an (operator-valued) quantum field with an *operator-valued distribution*, a map from smooth functions to operators, intuitively understandable as the integral of the operator field weighted by that smooth function. In more recent versions, the quantum field is replaced by an algebra net: a map from open subsets of spacetime to C^* -algebras. See Haag (1996) for the details.)

Once such a framework has been written down, it is possible to prove a great many fully rigorous theorems, which apply to any quantum theory satisfying those axioms. And many such theorems have indeed been proved; many have real mathematical beauty and importance. (Halvorson (2007) reviews some of this work.)

But none of this work means anything physical at all unless these axioms really do describe the quantum field theories that exist in nature. So in parallel with the search for general results, AQFT theorists have looked for concrete examples of quantum field theories which satisfy the axioms.

They made some initial success. It was recognised very early on that *non-interacting* quantum field theories — theories of free electrons, free photons, and the like — can readily be incorporated into the formal framework. Not much later, Glimm and Jaffe (1968) succeeded in constructing a two-spacetime-dimensional scalar quantum field theory with an interaction term ($\lambda\phi^4$) which was exactly definable without cutoffs. The obvious next step was to look for interacting theories in four spacetime dimensions: perhaps toy models to begin with, but ultimately, rigorous versions of the interacting field theories used in mainstream particle physics. Despite forty years of work, though, no such theory has ever been found: the only known physically realistic algebraic quantum field theories in four dimensions are free-field theories.

Another way to look at the situation may be helpful. As noted above, physicists classify interaction terms into *renormalizable* and *non-renormalizable*. This can be understood in physical terms via the renormalization theory of the previous section; it can also, though, be defined mathematically in terms of the behaviour of the short-distance divergences in the integrals that appear in the perturbative expansion of the interaction due to that term. Among

the renormalizable terms, though, are so-called *super-renormalizable* terms:⁷ those which lead to infinities in at most a finite number of elements in the perturbative expansion. The $\lambda\phi^4$ term in two-dimensional scalar field theory is super-renormalizable, and so it would be unsurprising to find that the theory can be reformulated, without a cutoff, just by modifying the Hamiltonian to remove precisely the divergent terms in the perturbative series. Glimm and Jaffe’s rigorous, cutoff-free field theory does precisely that. But the trick will not work for interactions which are not super-renormalizable. No-one has yet succeeded in constructing an interacting field theory whose interaction terms are not super-renormalizable. And in four spacetime dimensions, *no* interaction terms are super-renormalizable.

However we look at it, though, the result is the same, and cannot be stressed enough. **There is no known physically realistic interacting algebraic quantum field theory in four dimensions.**

6 Comparing research programs: CQFT versus AQFT

I argued in section 3 that CQFT and AQFT, because they take different *and conceptually incompatible* approaches to the problem of renormalization, should be understood as rival research programs and evaluated on that basis. We’ve now reviewed each program; how do they compare to one another?

To begin with the “conventional” or “cutoff” approach: CQFT makes a very large number of novel empirical predictions. Chief amongst these are the hundreds of cross-sections, decay rates, mass ratios, resonances, magnetic moments, and the like, calculable within the Standard Model and empirically confirmed: the only way we know to make mathematical and physical sense of the Standard Model is via CQFT, so it doesn’t seem unreasonable to count the quantitative predictions of the Standard Model as predictions of CQFT. (If AQFT could also reproduce those predictions, they would equally count as novel predictions of AQFT; it can’t.) Increasingly, the methods of lattice QCD allow us to extend the quantitative predictions of the theory beyond the region where scattering theory is applicable, and to calculate the mass spectrum of baryons and mesons with increasing (though still fairly low) accuracy.⁸

Over and above these specific predictions, CQFT provides an explanatory framework for the whole of particle-physics phenomenology, and one in which the details of renormalization theory play a central role. Amongst the explanatory successes of this framework are:⁹

⁷Many textbooks do not regard the super-renormalizable terms as renormalizable; this is just a matter of definition.

⁸See Kalloniatis, Leinweber, and Willians (2005) for a fairly recent review of progress in this field.

⁹See Peskin and Schroeder (1995), Cheng and Li (1984) or any other sufficiently advanced QFT textbook for details.

- “asymptotic freedom”: the feature of QCD whereby the effective coupling constant decreases in strength as we consider shorter- and shorter-distance interactions.
- The particular functional form of the Lagrangian of any interaction, given only the symmetry properties of the fields in that Lagrangian: whatever the short-distance form of the dynamics, the effective Lagrangian appropriate for physics at lengthscales long compared to the cutoff will contain all and only the renormalizable interaction terms.
- The reason that the weak interaction is weak: the weak interaction between e.g., electrons and neutrinos can be modelled as a four-fermion interaction, cut off at energy levels of $\sim 80\text{GeV}$ (the mass of the W boson) by the (spontaneously broken) electroweak interaction; a 4-fermion interaction term is nonrenormalizable, and so will be very weak at energy scales low compared to that cutoff level.
- Anomalous symmetry breaking: the process whereby classical symmetries cannot be preserved by any renormalization scheme, and therefore are violated by radiative-correction terms in the theory.
- The reason why internal symmetries of field theories are always represented irreducibly: renormalization will cause the coupling constants, including the masses, of different irreducible components in a reducible representation to shift relative to one another.¹⁰
- The reason for the close analogies between condensed-matter and particle-physics QFT.

So judged by the usual standards of science — wideness of applicability, predictive accuracy, explanatory and unificatory power — CQFT is doing pretty well ... to put it mildly.

What about AQFT? Being uncharitable, it makes no (non-falsified) empirical predictions whatsoever, because the world we live in manifestly has (a) interactions and (b) four (or more) spacetime dimensions, and we have **no interacting algebraic quantum field theories in four dimensions**. Even being charitable: the only empirical predictions of AQFT are general results (the spin-statistics theorem, the CPT theorem, etc.) which are also derivable (by the usual standards used in theoretical physics) in CQFT (perhaps only as extremely good approximations, depending on whether the world is Poincaré-covariant at the fundamental level.)

It is no part of my argument to object to the first-order project of continuing to look for interacting AQFTs. (It’s not the research project I’d put money on, but what do I know?) But in the absence of any such theories — and in the presence of a rival program which has been spectacularly successful — there is, at present, just no reason to expect that program to succeed. And if so,

¹⁰For more on this point, see Wallace (2009, section 7).

philosophers of physics should not be studying that research program in the expectation that they will learn about the structure of reality in doing so.

7 Another problem for AQFT: why trust claims about the arbitrarily small?

For the sake of argument, let us put aside the lack of *positive* reasons to expect that an empirically adequate AQFT could be constructed. Is there anything about the success of CQFT that actually rules them out?

One argument might run as follows:¹¹

1. We have a very well confirmed theory (the Standard Model, understood as a CQFT), one of whose central claims is that field degrees of freedom are frozen out at sufficiently short lengthscales.
2. As good scientific realists, we should tentatively accept that claim as approximately true.
3. Since AQFT denies that claim, its basic structure is wrong.
4. So no wonder that we can't construct empirically adequate theories within that structure!

This is a little too quick, though. CQFT, in itself, doesn't actually *require* the existence of a cutoff; it just tells us that *assuming* a cutoff suffices to make the theory well-defined (and that, having assumed it, the details of how it is imposed are irrelevant at long lengthscales). It is consistent with the CQFT framework that the theory's degrees of freedom after all remain defined on arbitrarily short lengthscales.

Now, *some* specific QFTs do (appear to) rule out this possibility.¹² In quantum electrodynamics, in particular, the strength of the interaction appears to increase without limit as we study the theory on shorter and shorter lengthscales, and in fact is believed to diverge below a certain lengthscale — for QED to be well-defined, then, it needs a cutoff which is above this lengthscale. But non-Abelian gauge theories — notably QCD and the electroweak theory — do not have this property: in fact, their interaction strength (as was briefly noted above) actually decreases with decreasing lengthscale. So there appears to be no lower limit to the value of the cutoff in QCD or electroweak theory, and therefore no mathematical objection (from this source, at any rate) to the existence of an algebraic version of either.

Now, as I have stressed, I am aware of no positive reason why such a theory *should* exist. But suppose, just for the sake of argument, that it does, and

¹¹I presented this argument in a talk at the Western Ontario conference on Philosophy of Quantum Field Theory, in spring 2009. I'm grateful to Nick Huggett for pointing out that it does not apply to theories with ultraviolet freedom.

¹²Again, see Peskin and Schroeder (1995), or any other modern QFT textbook, for details.

that we managed to find it. What would we gain from this, as physicists *or* philosophers?¹³

So far as I can see, basically nothing. Given any finite lengthscale Λ , we can reproduce the claims of this hypothesised theory about field expectation values on lengthscales above Λ , just using the CQFT version of the theory with a cutoff $\ll \Lambda$. And if we know all the field expectation values on lengthscales above Λ , we know *everything* that the theory says about the structure of the world on lengthscales above Λ .¹⁴

But what about the theory's claims about physics on *arbitrarily short lengthscales*? And in any case, never mind the *structure* of the expectation values, what about what the theory tells us about the *nature* of reality? Good questions¹⁵ — or rather, they would be if we had any reason at all to trust the theory on those lengthscales. But hardly anyone thinks we do: once we get down to Planckian lengthscales, the fiction that spacetime is nondynamical and that gravity can be ignored will become unsustainable. Whatever our sub-Planckian physics looks like (string theory? twistor theory? loop quantum gravity? non-commutative geometry? causal set theory? something as-yet-undreamed-of?) there are pretty powerful reasons *not* to expect it to look like quantum field theory on a classical background spacetime. As such, what QFT (of any variety) says about the nature of the world on lengthscales below $\sim 10^{-43}$ m (let alone below $10^{-10^{43}}$ m, which is a perfectly legitimate lengthscale according to AQFT) doesn't actually tell us anything about reality.

An analogy might help. Classical electromagnetism plus pointlike charged particles is known to be mathematically pathological: the field strength at the location of the particle is infinite, which makes it difficult at best to work out how the field affects the particle itself. It's a moderately interesting mathematical game to see how the theory could be made rigorous (see Rohrlich (2007) for a classic account), but it's *only* a game, because we know that classical electrodynamics just isn't a remotely accurate description of reality below certain lengthscales and above certain field strengths, and so if we really want to understand charged particles, we have to do quantum mechanics.¹⁶

To sum up: as long as we didn't have a satisfactory conceptual understanding

¹³As *mathematicians*, no doubt we would gain a lot from it.

¹⁴Obviously, this presupposes that we don't need to add hidden variables, or some kind of dynamical collapse, to QFT in order to extract empirical data from it. Those not persuaded by either the Everett interpretation of quantum mechanics (cf. Saunders, Barrett, Kent, and Wallace (2010) and references therein), or by operationalist or neo-Copenhagen approaches to quantum mechanics are invited to draw their own conclusions as to how far the account of this paper applies to their position; since we don't at present have a fully satisfactory dynamical-collapse or hidden-variable version of the Standard Model (though see Colin and Struyve (2007) and Struyve and Westman (2007) for interesting progress), let alone of some algebraic-QFT variant of the Standard Model, the question is probably premature in any case.

¹⁵Actually, I don't think questions about the nature (as opposed to the structure) of the world are good questions at all (cf. Ladyman and Ross (2007)), but never mind.

¹⁶Arguments like this seem under-appreciated in philosophy of physics, for reasons that perplex me; the continued attention by philosophers of statistical mechanics to the 'problem of measure zero', which is an artefact of classical physics with no direct quantum-mechanical equivalent, is another such pseudoproblem (see Wallace (2001) for further discussion).

of renormalization theory, it was possible to suppose that the only satisfactory way to understand it was via AQFT. But the success of CQFT in making conceptual sense of renormalization both undermines the reasons to expect that we will find interacting AQFTs, and removes the interest in so doing: for any such theory will say nothing *believable* about the world that could not already have been said via a CQFT.

8 Fraser’s defence of AQFT

By and large, philosophical explorations of AQFT have proceeded with little or no explanation as to why that framework is appropriate for the philosophy of quantum field theory. An honorable — and refreshing — exception is Fraser (2009), who explicitly argues that “an interpretation of QFT should be based on a rigorous axiomatic variant of QFT rather than any of the other variants.” Fraser engages directly with CQFT (notably in the form which I defended in Wallace (2006)) and tries to show why it is an inappropriate basis for foundational work; in this section, I will attempt to return the compliment. In section 8.1 I summarise the context of Fraser’s arguments; in sections 8.2 and 8.3, respectively, I respond to her criticisms of CQFT and her positive case for AQFT.

8.1 Underdetermination

Fraser begins her discussion by considering $\lambda\phi_2^4$ — scalar field theory in two spacetime dimensions, with a ϕ^4 interaction term. She considers three variants of this quantum field theory. The first is a ‘naive’ QFT of the kind used in the 1950s and 1960s, in which infinities are tamed by an unexplained algorithm. The second and third are, near enough, CQFT and AQFT versions of the theory, respectively.

(Three points should be noted here. Firstly (as was mentioned above) this theory really does have a rigorously defined AQFT version, developed by Glimm and Jaffe — thus far, at least, Fraser is discussing a concrete AQFT, not just a hoped-for one. Secondly, Fraser refers to the ‘naive’ QFT as a *theory*, whereas I would regard it as too conceptually ill-formed to deserve this name — but I doubt whether this is a difference of substance.

Thirdly, and more interestingly, Fraser’s cutoff QFT incorporates a long-distance as well as a short-distance cutoff, introduced to tame a different set of infinities resulting in the divergence of the integrals in the theory at *long* lengthscales. This is one place where my earlier disclaimer is relevant: this kind of *long*-distance (or ‘infra-red’ divergence) arises because the free-field vacuum and the interacting-field vacuum (that is, the ground states of scalar QFT with and without the ϕ^4 term) differ on arbitrarily long lengthscales and so are unitarily inequivalent. This really is a case where algebraic methods are required to understand what is going on. As I have stressed, the point at issue in this paper is not algebraic vs. non-algebraic methods, it is real cutoffs (represented

by CQFT) versus no cutoffs (represented by algebraic QFT with its requirement that fields are definable on arbitrarily small open sets).¹⁷

Fraser now argues that we have a real case of underdetermination here: the CQFT and AQFT versions of $\lambda\phi_2^4$ are empirically equivalent (in the sense that they predict the same scattering cross-sections) but make different theoretical claims (notably about whether space is discrete¹⁸). And she claims that in this case, there are clear theoretical grounds to prefer the AQFT version: the cutoff in the CQFT version is ad hoc and unmotivated.

So far, so good: I agree completely. Of course the cutoff is unmotivated: why suppose there is a cutoff when the theory demonstrably can be defined on all lengthscales? In the (fairly remote) counterfactual situation where I was an organism in a world which appeared to be accurately described by $\lambda\phi_2^4$, I'd follow Fraser in assuming, tentatively, that the AQFT version of the theory was true (while noting, from its empirical equivalence with the CQFT version, that nothing in the phenomenology thus far observed rules out the failure of the theory at some as-yet-unprobed lengthscale).

On the basis of this example, Fraser attempts to construe the whole CQFT-vs-AQFT debate as likewise a case of underdetermination. But I find this incomprehensible. The $\lambda\phi_2^4$ case is a genuine case of underdetermination, because we have two empirically equivalent versions of the theory — a CQFT version and an AQFT version. The case of QED, QCD or the standard model differs in a rather crucial respect, because we only have the CQFT version. Where is the underdetermination here?

Let me press the point. In the $\lambda\phi_2^4$ case, Fraser advocates that we reject CQFT because of the *actuality* of underdetermination, and instead accept an *actual* AQFT. But in the four-dimensional case, she advocates that we reject CQFT because of the mere *possibility* of underdetermination, and instead accept AQFT even though physically realistic AQFTs are currently *nonexistent*.

Although I don't think Fraser directly addresses this problem with her underdetermination thesis, she does discuss the absence of interacting AQFTs, which she acknowledges as a concern for her position. She writes, though, that “[t]his concern can be alleviated by recognising that the formal variant of QFT is best viewed as a program which has yet to be completed”. Clearly this is correct, but I'm unclear why this is supposed to help. Lots of research programs have been “yet to be completed” at various points in history. In some cases, this was because they would be completed at a later time. In other cases, it was because they were just wrong, and the problem they sought to solve would be solved by a different method. It's not at all clear why Fraser thinks AQFT is in the former category.

¹⁷I discuss this general issue further in I discuss this further in section 4 of Wallace (2006). In the concrete case of the infra-red divergences of quantum electrodynamics, we can explicitly see how the inequivalent representations arise: calculations done via an infrared cutoff (usually imposed as a low-mass limit on the photon) show that the ‘in’ and ‘out’ stage in QED scattering processes have an infinite expected number of photons, though only a finite expected energy (Chung (1965); see also Kulish and Faddeev (1970) and, for more recent perspectives, Horan, Lavelle, and McMullan (2000) and references therein).

¹⁸And finite, says Fraser; but see the above comment on infrared divergences.

Fraser does make some comments about the plausibility of the program’s ultimate success, but none of them really seem to help. She observes that “QFT is hardly unusual in requiring refinements over a period of time” — agreed, but that seems to presuppose that these refinements are the correct ones. And she claims that “the results that have been obtained so far do not give any reason to believe that this [an interacting AQFT] [is] an unattainable goal.” This is a curious attitude to the burden of proof on a scientific research program. If we want to take AQFT as *sufficiently secure that it should be our route to make foundational claims about the world* (and not merely as a promising line of research), isn’t more needed than this? It’s not at all clear to me that there were reasons internal to S-matrix theory, or to the old quantum theory, that gave us “any reason to believe” that they would not succeed. Rather, they were abandoned because something different came along that solved their problems more successfully.

I think Fraser’s position here really relies on her view that CQFT *does not* solve the problems of QFT successfully, so that AQFT is the only game in town. That is: her confidence that the AQFT program will succeed is based on a rejection of my claims, in this paper, that the CQFT program succeeds in resolving the problems of renormalization. If so, the situation isn’t really one of underdetermination, as Fraser suggests: it’s one where the only acceptable (and not just the best) way of making sense of QFT is via AQFT methods.

8.2 Fraser’s reasons to reject CQFT

Why does Fraser make this claim? Her main objection — so far as I can tell — to the existence of a real short-distance cutoff is that it is *ad hoc*: “it does not have an independent motivation” (p.552). She claims that “to the best of my knowledge — no-one defends the position that QFT offers evidence that space is discrete” (p.552). She contrasts this case with the case of quantum gravity: “arguments from quantum gravity that space is discrete are supported by deep theoretical considerations about how gravity is to be quantized” (p.552). But she denies that those very “deep theoretical considerations” can be used by the quantum field theorist to defend the cutoff:

the fact that quantum gravity indicates that space is discrete would not help to settle the question of how to interpret the cutoff variant of QFT because gravitational considerations are external to QFT. The point at issue is whether QFT dictates that space is necessarily discrete and finite in extent; that is, whether the discreteness and finitude of space is a foundational principle of QFT.” (p.552)

But the advocate of CQFT should reject this framework. If we’re serious about understanding our best extant physics and using it to learn about the world, the right question isn’t “what would the world be like in the hypothetical situation where there was no gravity and QFT was exactly true”? Addressing this question *obviously* can’t tell us whether to accept AQFT or CQFT, because we get different answers according to which theory we are taking to be exactly true!

Nor is the right question “what would it be rational to believe about QFT, if we lived in a world where there was no gravity but QFT was empirically very successful”? This question is fairly uninteresting, for the simple reason that we *don't* live in such a world. The scientifically and philosophically interesting question has to be “what is it rational to believe about QFT, given all the empirical data that we *in fact* have?”

From this point of view, it's perfectly reasonable (and not at all *ad hoc*) to be motivated to believe in a discrete structure to spacetime because it is independently motivated by various different strands in theoretical physics, *including* the need for a short-distance cutoff in QFT. And indeed (*contra* Fraser's claim) this position is pretty commonly taken by particle physicists. Peskin and Schroeder, for instance, write

In QED and other quantum field theories appropriate to elementary particle physics, the cutoff would have to be associated with some fundamental graininess of spacetime, perhaps a result of quantum fluctuations in gravity. (Peskin and Schroeder 1995, p.402)

To be fair to Fraser, her initial discussions of CQFT mostly seem to be concerned with a version which implements the cutoff via a literal discretisation of spacetime (a version which, so far as I know, basically no-one advocates). She distinguishes this from the version of CQFT that I present in this paper (and in Wallace 2006), in which the actual mechanism imposing the cutoff is taken to be unknown and the arguments of renormalization theory are deployed to explain why this does not matter as long as we are interested in the theory's claims on scales large compared to the cutoff.

Fraser characterises this strategy (accurately, I think) as changing the question of interpreting QFT from “if QFT were true, what would reality be like” to “given that QFT is approximately true, what is reality (approximately) like”, a move that (I think!) she is happy to permit. (Correctly so, surely, since we have excellent reasons to think that (a) QFT is not true, and (b) QFT is approximately true¹⁹.)

8.3 Fraser's two arguments for preferring AQFT

Even given this way of thinking about interpreting QFT, however, Fraser still argues for AQFT rather than CQFT as the correct approach. Her reasons are twofold:

1. AQFT and CQFT do not in fact agree even in the large-scale claims they make about the world.

¹⁹I merely copy Fraser's term here: I take it that she intends it as a placeholder for whatever notion of “approximately being right” the scientific realist should use to handle theory change. My own view is that the right analysis to use here is the sort of structural continuity advocated by structural realists (see Ladyman and Ross (2007) and references therein; see also section 1 of Wallace (2001).)

2. Pursuing AQFT rather than CQFT is a better way to find the true theory to which QFT is an approximation.

The first, so far as I can see, is straightforwardly false on technical grounds. The methods of renormalization theory tell us that any field theory — including one which is actually defined on all lengthscales — will look like a CQFT on sufficiently long lengthscales.²⁰ Here “look like” means not merely “is empirically equivalent to”, but rather “agrees on all long-distance field expectation values”. And since all structural claims about states in QFT — including highly empirically remote claims about, say, the interior structure of the proton or the Higgs mechanism — are expressible in terms of the expectation values of field operators and their derivatives, it follows that CQFT and AQFT must agree on all structural claims about structures defined on lengthscales large relative to the lengthscale at which QFT breaks down — that is, they must agree on all structural claims that we have reason to believe. (This is basically a reiteration of my argument in section 7.)

Fraser in fact gives two concrete examples of the large-scale inequivalence of AQFT and CQFT.²¹ The first is inequivalent representations: she claims (p.560) that

the cutoff variant does not have even approximately the same content as algebraic QFT because the cutoff variant has a finite number of degrees of freedom and therefore does not admit unitarily inequivalent representations; in contrast, algebraic QFT has an infinite number of degrees of freedom and therefore admits unitarily inequivalent representations. Spontaneous symmetry breaking is one case in which these unitarily inequivalent representations are put to use.

Here Fraser is again assuming that CQFT requires both a short- and a long-distance cutoff; but as I have noted, I (and I think most quantum field theorists) are happy to grant that *long-distance* divergences really should be tamed by algebraic methods, and that QFTs defined on spatially infinite manifolds really do have infinitely many degrees of freedom, and hence unitarily inequivalent representations. The issue is not really algebraic vs. non-algebraic methods: it is a short-distance cutoff vs. no cutoff. It is true, of course, that AQFT has unitarily inequivalent representations even on spatially finite regions, but this unitary inequivalence *only* manifests itself with respect to expectation values

²⁰Again, see (e.g.) Binney *et al* or Peskin and Schroeder (1995).

²¹In this part of her discussion, Fraser seems to set aside the fact that we don’t actually have any interacting algebraic quantum field theory, let alone ones which are empirically equivalent to the Standard Model. She writes, in fact (p.559) that “even it were granted that QFT should be regarded as an approximate guide to the ontology of relatively large distance scales, it does not follow that the cutoff variant of QFT succeeds in describing large-scale ontology.” — as if there were both cutoff and non-cutoff versions of the Standard Model. But for the sake of argument, in this part of my discussion I will grant Fraser the (in my view unmotivated) assumption that there is a quantum field theory in the AQFT framework which is empirically equivalent to the Standard Model.

on arbitrarily small spacetime regions, and these are exactly those expectation values which don't convey real information about the world.

Fraser's second example is the existence of quanta (by which she means, I think, the existence of a Fock-space representation of the theory). She has argued elsewhere (Fraser 2008) that an interacting AQFT (if it existed) would not admit such quanta, and states (p.560) that "a theory according to which quanta exist is not approximately equivalent to a theory according to which quanta do not exist".

Well, maybe not, but it is approximately equivalent to a theory according to which quanta do approximately exist! And if AQFT (more precisely, if this supposed interacting algebraic quantum field theory) does not admit quanta in at least some approximate sense, then so much the worse for it: the evidence for the electron is reasonably conclusive.

Moving on to Fraser's second argument for AQFT over CQFT in this context (which she takes (p.559) as a "more important line of objection"): her position, basically, is that we should expect that "theory X" (her term, borrowed from Wallace (2006), for whatever the theory is to which QFT is an approximation) will resemble QFT not just in its empirical claims but in its underlying structure: "theory change from QFT to theory X must involve approximate continuity at the theoretical level and not merely approximate continuity at the empirical level". I'm happy to grant this, if by "theoretical level" she means to include the theory's field expectation values and their dynamics on lengthscales long compared to the cutoff: we have plenty of reason, to be sure, to expect a large degree of structural continuity between QFT and X. But I think she means something more — something which would serve to differentiate the Standard Model from her (imagined) algebraic variant. And if so, I have no idea why she thinks this is the case.

She goes on to claim that

[U]ntil we actually find theory X and can determine the respects in which there is approximate continuity between theory X and QFT, we have no way of knowing how useful it would have been to work out the details of QFT [i. e. , develop AQFT]. The fact remains that — given our limited knowledge at the present time — continuing to pursue the development of QFT [i. e. , continuing to work on AQFT] is a good strategy.

Well, I suppose it might be. *Perhaps*, if we succeeded in finding an interacting QFT in the algebraic framework — or if we succeeded in proving that there is no such theory — it might tell us something salient to the search for quantum gravity. Who knows? But in any case, what has this to do with the *philosophy* of quantum field theory — that is, with the project of learning about the world by studying the conceptual claims in QFT? Fraser's argument is an argument for the first-order project of trying to find interacting algebraic quantum field theories, not for the philosophical project of asking what the world would be like if there was such a theory and it correctly described reality.

(I should note in passing that Fraser also claims that “[r]esting content with the cutoff variant of QFT because it is empirically adequate at large distance scales would be a strategic mistake because it would hinder the search for theory X”. This seems mistaken. For one thing, CQFT does give us some information about X (not much, to be sure, by the very nature of renormalization theory). The variation of the coupling constants in the Standard Model on different lengthscales, in particular, tells us something about ways in which the theory ought to break down or be modified at lengthscales far below our ability to probe; notably, one reason to take supersymmetry seriously is that, when we allow for it — but not otherwise — the strong and electroweak interactions are compatible with the hypothesis that they are aspects of some larger gauge group. (See chapter 22 of Peskin and Schroeder (1995), and references therein, for details.) More fundamentally, though, it is CQFT that forms the theoretical basis for our best current prospect of gaining *empirical* evidence of X, via the LHC.

Fraser deserves considerable credit for making an explicit case for AQFT to be the basis for philosophy of quantum field theory. But in doing so, I think, she really succeeds only in showing how surreal is the proposal that philosophy of QFT should deal, not with the — spectacularly successful — quantum field theory that we actually have, but with a research program which has yet to give us any evidence at all that it is correct.

9 Other objections to CQFT

Although I’m not aware of any defences of AQFT (as the basis for philosophy of QFT) *in print* other than Fraser’s, I’ve come across various other objections in conversation with colleagues; here I’ll try to set out, and respond to, the more common of them.

Objection: CQFT is not Poincaré-covariant.

Response: It’s true that those cutoff schemes we can actually concretely implement (notably, putting the theory on a lattice) violate Poincaré covariance (and, indeed, translation and rotation covariance). But provided that the real cutoff is actually imposed by some Poincaré-covariant theory (string theory and loop quantum gravity both have this property), this will just be an artefact of the particular cutoff schemes we’re using at the moment (and recall that the details of that scheme are irrelevant to the large-scale structure of the theory).

Having said this, if Poincaré covariance turns out to be only phenomenological, so what? (It is worth noting, in this context, that condensed matter physics manages fine with symmetry based reasoning — reasoning based on the rotational or translational covariance of many systems, for instance — even though those symmetries are only phenomenological and are broken at the atomic level.)

Objection: CQFT is useless for metaphysics given its underdetermination: precisely because it tells us nothing about the nature of the cutoff, it's not useful to do metaphysical work on it.

Reply: If what you need for whatever metaphysical project you're engaged in is structural information about the world on scales large compared to the Planck length, CQFT will provide that information just fine.

If you're after something else — structural information about the world on *all* lengthscales, say, or *non*-structural information about the world, tough. The current state of the art in physics doesn't tell us any such information; if you want it, quantize gravity (or go and do something else until other people have quantized gravity). Don't think that you'll get such information by studying AQFT, though: we have no reason to *believe* its claims about what the world is like on sub-Planckian lengthscales (indeed, we have excellent reasons *not* to believe them).

(I'll concede that some people like to ask “if theory T were exactly true, what would the world be like?”, even though they know that theory T *isn't* exactly true. Those people are welcome to take $T = AQFT$, provided they don't kid themselves that they're discovering what the world is actually like.)

Objection: CQFT is not mathematically rigorous, so we don't actually know the results that it purports to establish.

Reply: Although the mathematical *and physical* pathologies of pre-1970s QED made this a reasonable objection prior to the development of renormalization-group methods, things have now changed, and the level of mathematical rigor in CQFT is basically the same as elsewhere in theoretical physics. So applying this standard even-handedly would invalidate practically all of the theoretical physics of the twentieth century: atomic physics, non-relativistic scattering theory, astrophysics, nuclear physics, etc, etc.

Objection: CQFT is very ugly; AQFT is much more elegant.

Reply: This is a quite widespread view, but it seems to be held mostly by people who haven't really studied modern field theory in depth. I encourage them to do so; they may be favourably surprised by how elegant it is.

More seriously, theoretical elegance *might* be a criterion for theory selection between theories tied for first place with respect to the empirical evidence. Using it to reject a well-confirmed theory which has made abundant novel predictions in favour of a “theory” that currently doesn't exist does violence to the scientific method. (Feel free to hold out for more elegance in whatever physics lies beyond the effective-field regime, of course.)

10 Conclusion

It has not been my purpose in this paper to disparage the mathematical physicists who continue to search for physically realistic interacting algebraic quantum field theories. Leaving aside the mathematical interest of the task, it is impossible to know in advance just which mathematical highways and byways may prove to be a useful source of insights for the future progress of physics. Let a thousand flowers bloom — so long as it is understood that there is no requirement in quantum field theory itself that any such theory exists, and that the apparent paradox of ultraviolet divergences that originally motivated the axiomatic program in QFT has long since been solved by quite another route.

Nor has it been, and nor is it, my purpose to suggest that philosophy of quantum field theory done in the AQFT framework has nothing to teach us. On the contrary, much — perhaps most — of that work probably transfers across just fine to CQFT. And this includes philosophy of quantum field theory that makes essential use of algebraic *methods* — notably in the investigations of inequivalent representations which arise because quantum field theories, with or without cutoffs, have infinitely many degrees of freedom when defined on spatially infinite spacetimes.

But it is the case, I think, that the philosophy of quantum field theory is needlessly complicated, and often distorted, when it is carried out on the assumption that interesting QFTs must be formulated within the framework of cutoff-free algebraic quantum field theory. And that assumption, from the point of view of twenty-first century theoretical physics, is entirely without motivation. We possess a concrete example of an interacting quantum field theory — the Standard Model — and if it does not fit the framework of AQFT, then so much the worse for that framework. For it is foundationally satisfactory in its own right — despite, or perhaps because of, its explicit acknowledgement that it ceases to be defined below some lengthscale. And it is the most widely applicable, predictively powerful, and precisely tested scientific theory in human history. If the goal of philosophy of physics is to understand the deep structure of reality via our best extant physics, then to be lured away from the Standard Model by algebraic quantum field theory is sheer madness.

Acknowledgements

I am grateful for comments and constructive criticism from David Baker, Gordon Fleming, Hilary Greaves, Nick Huggett, Wayne Myrvold, Laura Ruetsche, Simon Saunders, Chris Timpson, and especially Jeremy Butterfield and Doreen Fraser.

References

- Baker, D. (2009). Against field interpretations of quantum field theory. *British Journal for the Philosophy of Science* 60, 585–609.

- Baker, D. and H. Halvorson (2009). Antimatter. *British Journal for the Philosophy of Science* 60, 1–29.
- Binney, J. J., N. J. Dowrick, A. J. Fisher, and M. E. J. Newman (1992). *The Theory of Critical Phenomena : an introduction to the renormalisation group*. Oxford: Oxford University Press.
- Cao, T. Y. (1997). *Conceptual Developments of 20th Century Field Theories*. Cambridge: Cambridge University Press.
- Cheng, T.-P. and L.-F. Li (1984). *Gauge Theory of Elementary Particle Physics*. Oxford, UK: Oxford University Press.
- Chung, V. (1965). Infrared divergence in quantum electrodynamics. *Physical Review* 140, B1110–B1122.
- Clifton, R. and H. Halvorson (2001). Are Rindler quanta real? inequivalent particle concepts in quantum field theory. *British Journal for the Philosophy of Science* 52, 417–470.
- Colin, S. and W. Struyve (2007). A Dirac sea pilot-wave model for quantum field theory. *Journal of Physics A* 40, 7309–7342.
- Dirac, P. (1930). *The Principles of Quantum Mechanics*. Oxford University Press.
- Earman, J. (2004). Curie’s principle and spontaneous symmetry breaking. *International Studies in the Philosophy of Science* 18, 173–198.
- Earman, J. and D. Fraser (2006). Haag’s theorem and its implications for the foundations of quantum field theory. *Erkenntnis* 64, 305–344.
- 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. Forthcoming in *Philosophy of Science*; available online at philsci-archival.pitt.edu .
- Gelfand, I. M., G. E. Shilov, M. I. Graev, and N. I. Vilenkin (1964). *Generalized Functions*. Academic Press.
- Glimm, J. and A. Jaffe (1968). A $\lambda\phi^4$ quantum field without cutoffs. I. *Physical Review (2)* 176, 1945–1951.
- Haag, R. (1996). *Local Quantum Theory: Fields, Particles, Algebras*. Berlin: Springer-Verlag.
- Halvorson, H. (2007). Algebraic quantum field theory. In J. Butterfield and J. Earman (Eds.), *Handbook of the Philosophy of Science: Philosophy of Physics, Part A*. Boston: Elsevier.
- Horan, R., M. Lavelle, and D. McMullan (2000). Asymptotic dynamics in quantum field theory. *Journal of Math J.Math.Phys.* 41 (2000) 4437-4451, 4437–4451.
- Kalloniatis, A., D. Leinweber, and A. Williams (Eds.) (2005). *Lattice Hadron Physics*. Berlin: Springer.

- Kulish, P. P. and L. D. Faddeev (1970). Asymptotic conditions and infrared divergences in quantum electrodynamics. *Theoretical and Mathematical Physics* 4, 745–757.
- Ladyman, J. and D. Ross (2007). *Every Thing Must Go: Metaphysics Naturalized*. Oxford: Oxford University Press.
- Peskin, M. E. and D. V. Schroeder (1995). *An introduction to Quantum Field Theory*. Reading, Massachusetts: Addison-Wesley.
- Rohrlich, F. (2007). *Classical Charged Particles* (3rd ed.). New Jersey: World Scientific.
- Rovelli, C. (2004). *Quantum Gravity*. Cambridge: Cambridge University Press.
- Ruetsche, L. (2002). Interpreting quantum field theory. *Philosophy of Science* 69, 348–378.
- Ruetsche, L. (2003). A matter of degree: Putting unitary inequivalence to work. *Philosophy of Science* 70, 1329–1342.
- Saunders, S., J. Barrett, A. Kent, and D. Wallace (Eds.) (2010). *Many Worlds? Everett, Quantum Theory, and Reality*, Oxford. Oxford University Press.
- Streater, R. F. and A. S. Wightman (1964). *PCT, Spin and Statistics, and All That*. New York: W. A. Benjamin.
- Struyve, W. and H. Westman (2007). A minimalist pilot-wave model for quantum electrodynamics. *Proceedings of the Royal Society of London A* 463, 3115–3129.
- von Neumann, J. (1955). *Mathematical Foundations of Quantum Mechanics*. Princeton: Princeton University Press.
- Wallace, D. (2001). Implications of Quantum Theory in the Foundations of Statistical Mechanics. Available online from <http://philsci-archive.pitt.edu>.
- Wallace, D. (2006). In defence of naiveté: The conceptual status of lagrangian quantum field theory. *Synthese* 151, 33–80.
- Wallace, D. (2009). QFT, antimatter, and symmetry. *Studies in the History and Philosophy of Modern Physics* 40, 209–222.