

Towards a Realist View of Quantum Field Theory

Draft for volume *Scientific Realism and the Quantum*, OUP

James D. Fraser

1 Introduction

Quantum field theories (QFTs) seem to have all of the qualities that typically motivate scientific realism. Alongside general relativity, the standard model of particle physics, and its subsidiaries, like quantum electrodynamics (QED) and quantum chromodynamics (QCD), are our most fundamental physical theories. They have also produced some of the most accurate predictions in the history of science: QED famously gives a value for the anomalous magnetic moment of the electron that agrees with experiment at precisions better than one part in a trillion. When it comes to actually articulating a realist reading of these theories, however, we run into serious difficulties.

This chapter puts forward what I take to be the most promising strategy for developing a realist epistemology in this context. I set up the discussion by highlighting the difficulty of making sense of QFT in orthodox realist terms if we restrict our attention to perturbative and axiomatic treatments of the theory (section 2). I then introduce the renormalization group (section 3), and argue, drawing on previous work by Wallace (2006, 2011) and Williams (2017), that it points to a way of rescuing a realist reading of QFT (section 4). I close by considering some challenges facing this renormalization group based realism (section 5). Besides some brief remarks in this final section, I will mostly be bracketing the measurement problem and associated interpretive puzzles about the physical content of quantum theories.

2 Realism and Quantum Field Theory: A First Look

Scientific realists disagree about how their position should be formulated in detail. We can discern some general ideas that lie behind many of these statements of the realist creed, however. Two, in particular, will be my focus here. The first is an explanatory thesis: realists take the empirical success of the sciences to be explained by the fact that they are getting something right about the way the world is. This contrasts with the anti-realist, who either denies that the empiri-

cal success of science stands in need of explanation, or else adopts an alternative explanatory strategy that does not invoke the representational veracity of our theories.¹ The second idea concerns the epistemic achievements of current scientific theories. Realists typically take at least some extra-empirical beliefs about the world to be supported by the predictive success of the sciences; that is, they take science to furnish knowledge of the unobservable. Again, this contrasts with anti-realist skepticism towards scientific claims about unobservable entities and properties. The constructive empiricist, for instance, takes an agnostic attitude towards the extra-empirical content of successful theories.

While these two claims do not strictly imply one another, connections between them are typically posited. The dominant approach to developing the realist position in recent decades has been to identify the theoretical claims we ought to commit ourselves to using an explanatory criterion. Rather than taking successful theories as a whole to be the locus of realist commitment, this ‘selective’ form of realism advocates belief in a subset of their descriptive content; namely, those parts of the theory which essentially contribute to, and therefore explain the success of, its empirical predictions. This way of cashing out the epistemic achievement component of realism is explicitly deployed in Psillos’s (1999) and Kitcher’s (2001) influential discussions of historical theory change, but the basic selectivist intuition arguably animates many of the other variants of realism put forward in the recent literature, such as structural realism. On this line of thought rolling out a realist epistemology across the various branches of science ought to proceed as follows: start with a theory’s empirical predictions and trace them back to theoretical claims that underwrite their success; these are the parts of science that constitute knowledge of the unobservable.

How does this programme fare in the case of QFT? While I will ultimately claim that there is scope for harmonising QFT with this selective realist scheme, I’ll start with a more pessimistic reading of the situation. Ignoring advances in renormalization theory since the 1950s, and restricting our attention to the perturbative and axiomatic approaches to QFT, it is possible to paint a fairly bleak picture of realism’s prospects in high energy physics.

As has already been mentioned, QFTs are wildly empirically successful. The problems start when we look at how these predictions are obtained. The most important source is the perturbative approach to QFT developed by Feynman, Tomonaga, Schwinger and Dyson in the late 1940s.² This is where the famous calculation of the anomalous magnetic moment of the electron, and the bulk of the standard model predictions currently being tested at the Large Hadron Collider, come from. I’ll sketch how this formalism works with reference to the ϕ^4 theory,

¹Van Fraassen’s (1980, 40) analogy between theory selection and natural selection can perhaps be read in either of these two ways. See also Saatsi (2015) for a discussion of anti-realist style explanations of the empirical success of particular theories.

²This is less true today than it was some decades ago. Non-perturbative calculation methods, especially Monte Carlo simulations in lattice QFT, are becoming increasingly powerful sources of empirical predictions. A complete epistemological study of high energy physics would need to look at these methods in detail but I leave this as a project for another time.

a well-understood model that shares most of the relevant features of the gauge theories that make up the standard model. The classical action of this model is given by:

$$S = \int d^4x \left[\frac{1}{2} (\partial_\mu \phi)^2 - \frac{1}{2} m^2 \phi^2 - \frac{\lambda_4}{4!} \phi^4 \right], \quad (1)$$

where ϕ is a scalar field, m is its associated mass, λ_4 is a positive constant parameterising the strength of the interaction and the integral ranges over Minkowski space-time. As is the case with realistic QFTs, there is little hope of finding exact solutions of this theory. The perturbative approach aims to get around this problem by setting up a power series expansion for quantities in an interaction coupling, in this case, λ_4 . The coefficients of this series can, in principle, be calculated using what we know about the corresponding exactly solvable free theory (with $\lambda_4 = 0$), though increasingly difficult computations are required at each order. The hope is that, if λ_4 is small, the first few terms of the series will provide an accurate approximation of the behaviour of a weakly interacting model.

This approximation strategy is, of course, ubiquitous in applied mathematics, but QFT perturbation series have their own peculiarities. As is well known, naively applying the perturbative method to QFT yields ill-defined results, leading many to consider abandoning the theory in its early days. Computing the series coefficients, for both the ϕ^4 theory and more realistic models, requires one to evaluate integrals over momentum space that diverge in the high momentum, or ultraviolet, region.³ In order to get meaningful predictions out of QFT, the so-called renormalization procedure was devised. This works roughly as follows. First, the offending integrals are rendered convergent; in the simplest case, this is achieved by imposing a cutoff—a large but finite upper limit on the momentum integrals. The expansion parameter is then redefined so as to remove the perturbative coefficient’s dependence on the cutoff. It turns out to be possible to completely ‘absorb’ the divergent part of the coefficients at each order into a finite number of parameters in the action if the theory contains only interaction terms parameterised by couplings with a zero or positive mass dimension—so-called renormalizable interactions.⁴ The ϕ^4 interaction has this property, as do the electroweak and strong interactions of the standard model. Indeed, during the years the standard model was being developed renormalizability was often viewed as a necessary condition for taking a theory seriously and was demanded a priori. At the end of the renormalization process, the cutoff is removed by taking it to infinity. Following this recipe produces a series with finite coefficients, and it is truncations of these renormalized perturbation series which are compared to data gathered from collider experiments.

³QFT perturbation series can also contain infrared divergences—integrals which blow up in the region of very low momentum. More generally, the low energy/long distance structure of QFT models raises distinctive problems of its own, which will not be discussed here.

⁴A brief note on dimensional analysis in QFT for the uninitiated. High energy physicists typically work with so-called natural units, which set $\hbar = c = 1$. This has the effect of drastically simplifying dimensional analysis, such that the dimensionality of any quantity can be expressed as a power of energy/mass/momentum, known as its mass dimension.

Despite its extraordinary empirical success, this approach to QFT has long been seen as conceptually suspect. Some have suggested that the perturbative formalism is insufficiently mathematically rigorous to engage with philosophically, and even inconsistent owing to its apparent clash with Haag’s theorem. I have argued in previous work (James Fraser, 2017) that the situation is not quite this bad. Perturbative QFT is, at least, internally coherent. There is a problem here, however, which should be especially worrying for the would-be scientific realist. Many of the manipulations that go into perturbative calculations, and the renormalization procedure, in particular, seems to be completely ad hoc. What licenses us to remove the infinities that appear in naive QFT perturbation series? The best we can really say within the confines of the perturbative approach itself is that we have to do this for the series to be well behaved and yield sensible predictions. But this is a purely instrumental rationale. Furthermore, it doesn’t give us any reason to think that the resulting renormalized series ought to produce accurate predictions. The realist, remember, was committed to there being an explanation for the success of scientific predictions which ultimately rests on the way the world is. On the face of it at least, perturbative QFT violates this demand.

The prospects of extracting knowledge of the unobservable from the perturbative formalism also appear to be grim. The root of the perturbative formalism’s ad hoc character is arguably that it fails to provide a characterization of the mathematical structure, and descriptive content, of QFT models. I’ll push this point in the language of the functional, or path integral, approach to quantization here (in part because it lays the ground for the discussion of the renormalization group in the next section). The key quantity in this approach is the partition function Z , which (in the case of a single scalar field ϕ) is associated with the integral:

$$Z = \int \mathcal{D}\phi e^{iS}. \tag{2}$$

Informally, the measure $\mathcal{D}\phi$ indicates that a sum is being taken over all possible configurations of the field (and S is, again, the classical action). Once we have the partition function of a QFT model we can, in principle at least, derive all of its physical quantities. More importantly, we can explicitly construct it as a mathematical object.

It turns out to be very difficult to precisely define this integral for fields that live on continuous space-times, however. Owing to the infinite number of degrees of freedom that exist in any space-time region we need to define a measure over an infinite dimensional space, once again leading to divergences in the ultraviolet region. In the perturbative treatment of QFT this issue was dodged rather than solved. In essence, the perturbative method allows us to set up expansions for Z , and thus for scattering cross sections, without really telling us what it is. The one strategy that might suggest itself for extracting a definition of the partition function from the perturbative formalism is to identify it with the sum of the series, but it turns out that, even after the divergences in the coefficients have been removed, realistic QFT perturbation series do not converge. In the absence of a clear characterization of the physical content of QFT, it is hardly surprising that

we struggle to find justifications for the moves made in perturbative calculations. But this situation is also clearly bad news for the epistemic achievement component of realism. How can the empirical success of the standard model possibly support beliefs about unobservable aspects of reality if we can't even specify what the theory is saying about the world?

There is another strand of the QFT programme the realist might turn to here which does directly address the question of where QFT models live in the mathematical universe, namely the axiomatic approach to the theory. In the 1950s and 60s mathematical physicists dissatisfied with the limitations of the perturbative approach tried to put the theory on a firm non-perturbative footing by writing down sets of axioms that any relativistic quantum field could reasonably be expected to satisfy. In fact, two mathematical treatments of QFT came out of this tradition: one based on the machinery of operator-valued distributions (Streater and Wightman, 1965) and the other on von Neumann algebras (Haag, 1996; Halvorson, 2007). Mathematically rigorous work in these frameworks succeeded in showing that one can precisely define path integrals for field configurations that take arbitrarily large momenta, at least in the case of some toy models in space-times with a reduced number of dimensions. Unfortunately though, the axiomatic tradition ultimately doesn't offer much solace to the aspiring realist about high energy physics. The crucial rub is that, as of yet, neither the standard model nor any interacting QFT in four-dimensional Minkowski space-time, has been constructed as a model of these axiomatic systems. While axiomatic treatments of QFT give us a clear set of theoretical principles, and therefore at least potential targets for realist commitment the connection with empirical predictions that the realist needs is missing.

In sum then, both the explanatory and epistemic achievement components of scientific realism seem to run into serious trouble in the QFT context. I should stress that I don't take this to be an existential threat to a realist view of science as a whole. An obvious response to this situation is to weaken the strength of one's epistemic commitments in this context, perhaps citing the theoretical immaturity of the QFT programme. One could admit, for instance, that we cannot yet explain the success of QFT's empirical predictions, or make precise claims about the nature of the unobservable world on their basis while remaining optimistic that future scientific progress will eventually come up with the goods.⁵ Furthermore, some philosophers of science have recently been arguing that realism needs to be weakened anyway in the face of anti-realist critique. Saatsi (2015, 2016) advances a position he calls minimal realism which abandons the epistemic achievement component entirely and focuses instead on a stripped back version of the explanatory thesis. The minimal realism is committed to their being an explanation of the success of current theories in terms of the way they latch onto the world but admits that we may not be able to say how this story goes in our current epistemic

⁵Doreen Fraser (2009) can be read as advocating a position like this. According to her, the axiomatic approach should be viewed as a work in progress that we have good reason to hope will eventually solve the puzzle surrounding the physical content of the standard model.

position. This weakened explanatory thesis is clearly compatible with the state of play in high energy physics as I have just characterized it.

I will be suggesting in what follows, however, that there is scope for rescuing the traditional explanatory and epistemic achievement theses in the QFT context. The preceding discussion may be a fair assessment of the situation as it stood before the 1970s, but developments in renormalization theory open up the possibility of a more full-blooded realist reading of the theory.

3 The Renormalization Group

In the decades following the invention of the perturbative QFT formalism the notion of renormalization broadened in scope and ultimately underwent fundamental conceptual changes. Fruitful exchange of ideas between condensed matter and high energy physics culminated in the emergence of the so-called renormalization group in the 1970s. Nowadays, renormalization methods enjoy applications in many areas of physics and beyond, taking on different forms in different theoretical contexts. I focus here on the incarnation of the renormalization group which is most significant for the philosophy of QFT: the momentum space approach developed by Kenneth Wilson.⁶

The core idea underlying the renormalization group, in all of its guises, is the study of coarse-graining transformations: operations which take us from an initial system of interest to a new one that lacks some of the degrees of freedom associated with high energies and small length scales but shares its large-scale properties. One reason why renormalization group methods are so diverse is that there are many ways of implementing a transformation of this kind. One approach employed in the study of lattice spin systems in statistical physics, for instance, is to replace groups of neighbouring spins with a single ‘block spin’ degree of freedom and tune the dynamics of the new system so as to reproduce the same (or, in practice, similar) macroscopic behaviour (figure 1). In some cases, it may be possible to invert the transformation, forming a group structure—hence the name. But this

⁶Wilson and Kogut (1974) being the classic review of this approach.




Figure 1: A ‘majority rule’ blocking procedure on a two valued spin system.

is not always possible; blocking transformations of the kind just described are typically not invertible. In any case, group theory seldom plays an important role in renormalization group methods, so the terminology is always misleading in one way or another. The real significance of these transformations is that they can be understood as inducing a ‘flow’ on a space of possible theories. Studying this flow turns out to be a powerful source of information about the behaviour of physical systems at different scales.

Rather than working in real space, as in the blocking approach, Wilson pioneered the idea of implementing a coarse-graining transformation in momentum space. To see how this works we need to return to the path integral expression for the partition function, Z . As we already discussed, defining integrals over field configurations is a project fraught with difficulty. Wilson’s key insight was that, rather than considering all configurations at once, we can evaluate the integral in sections. To get started with this approach we need to explicitly remove from the path integral the configurations associated with momenta above some very high value Λ :

$$Z_\Lambda = \int_{|p| \leq \Lambda} \mathcal{D}\phi e^{iS}. \quad (3)$$

We’ll see how this can be rationalised in detail shortly, but for now we can think of Λ as the energy scale at which new fields, or physics beyond the scope of the QFT framework entirely comes into play (in the most extreme case this would be the so-called Planck scale at which quantum gravity effects are expected to become important). This is somewhat analogous to the cutoff imposed in perturbative renormalization, but we are now operating in a completely non-perturbative context. What results from this cutoff procedure then is a well defined physical system that lacks degrees of freedom associated with arbitrarily large momentum, and arbitrarily small length scales. In contrast to the models sought in the axiomatic approach to QFT, these systems have a finite number of degrees of freedom in any finite space-time region. I’ll call these structures cutoff QFT models.

The Wilsonian renormalization group sets up a coarse-graining transformation on these cutoff QFTs as follows. We isolate the contribution to the integral due to the highest remaining field configurations, whose Fourier transforms have support above some value μ ; this part of the integral is then computed separately and absorbed into a shift in the action. In symbols:

$$\int_{|p| \leq \mu} \mathcal{D}\phi \int_{\mu \leq |p| \leq \Lambda} \mathcal{D}\phi e^{iS} = \int_{|p| \leq \mu} \mathcal{D}\phi e^{i(S+\delta S)}. \quad (4)$$

This defines a transformation that takes us from an initial cutoff QFT model to a new one, which has a lower cutoff and a modified dynamics but behaves like the original.

This transformation will not only alter the values of coupling parameters in the initial action but also give rise to new interaction terms. In general, we need to consider all possible terms which are not ruled out by initially demanded symmetries and constraints, including non-renormalizable interactions. For scalar fields,

this means going beyond the ϕ^4 theory to a broader class of actions including interactions like: $\lambda_6\phi^6$, $\lambda_8\phi^8$, and so on. The renormalization group transformation can then be seen as inducing a flow in a space of possible theories, spanned by the couplings $(m, \lambda_4, \lambda_6, \dots)$. The way that these parameters change as μ is lowered, and more and more high momentum degrees of freedom are ‘integrated out’, is described by the so-called beta functions:

$$\beta_i = \mu \frac{\partial}{\partial \mu} \lambda_i. \quad (5)$$

In most interesting contexts the equations governing the renormalization group flow are highly non-linear and cannot be exactly solved. Probing them via various approximation methods, however, has proven to be a powerful tool for answering many important questions about QFTs.⁷

What is most significant for the scientific realism debate is how the renormalization group flow behaves in the low energy regime, far below the cutoff Λ . It turns out to be possible to say something surprisingly general about this. As μ decreases the renormalization group flow is thought to be attracted to a finite dimensional surface spanned by the parameters associated with renormalizable terms in the action. Polchinski (1984) gives a rigorous demonstration of this behaviour in the case of scalar field theories by linearising the renormalization group equations in the region of small couplings. If we do this it is possible to show that variations in the non-renormalizable couplings at initial cutoff scale Λ are almost completely absorbed into variations in renormalizable couplings as we coarse-grain to a much lower scale $\mu \ll \Lambda$, up to powers of μ/Λ (figure 2). While this is a somewhat special result, and rests on the assumption that the couplings are small, this sort of behaviour is believed to hold generally and apply to realistic theories like QED and the standard model.

What this means is that QFT models that differ dramatically in their high energy dynamics manifest very similar physics at lower energies; a phenomenon known as universality in the physics literature. Basically any scalar field theory will look like the familiar ϕ^4 model at sufficiently low energies, for instance. This leads to a new perspective on QFT, and on the cutoff imposed at the beginning of the renormalization group analysis in particular. The renormalization group results just discussed demonstrate that removing the high energy degrees of freedom of a field theory leaves its low energy behaviour more or less unaffected. Since a model’s low energy physics can be almost completely parameterised by renormalizable coupling, all varying the value of Λ (and the details of how it is imposed) can do is move it around this finite dimensional surface. Fixing the values of the renormalizable couplings via a finite number of experimental measurements absorbs almost all of the cutoff dependence of a QFT model’s physical quantities. Cross sections for scattering events at an energy scale E , for instance, will only depend on the cutoff through powers of E/Λ . The cutoff, then, allows us to bracket

⁷The renormalization group also turns out to be relevant to the project of constructing QFT models without cutoffs which satisfy the sets of axioms devised in the axiomatic approach to QFT, for instance. See Hancox-Li (2015).

the question of what the world is like at the fundamental level while accurately modelling its low energy properties.

In the post renormalization group era, QED and the standard model have come to be regarded as ‘effective field theories’: models that are valid in some limited energy regime but should not be trusted beyond them. It is this shift in the outlook and methodology of the QFT programme which opens new paths for the scientific realist, as I will now argue.

4 Brighter Prospects for Realism

I claim that these advances in renormalization theory have improved the situation for both of the components of the realist doctrine in the QFT context. Firstly, the renormalization group framework provides a physical understanding of the original perturbative renormalization procedure. It thus opens up the possibility of achieving the sort of explanation of the empirical success of perturbative QFT the realist seeks. And secondly, it leads to a new perspective on the problem of characterizing QFT’s physical content and ultimately suggests a way of articulating claims about the world that are supported by the success of the standard model.

How does the renormalization group help us make sense of perturbative QFT? One important upshot that has been emphasised in the recent philosophical literature is that it transforms our understanding of the notion of renormalizability (Butterfield and Bouatta, 2014). As I mentioned in section 2, renormalizability was traditionally viewed as a property that any viable QFT model must possess. There was always something mysterious about this way of thinking, however. Why should the world be structured in such a way as to be amenable to perturbative approximation? Without the renormalization group results just described, the fact that the actions of empirically successful theories are renormalizable would seem to be a lucky coincidence. We can now see, however, that limiting one’s attention to renormalizable interaction terms is a very reasonable thing to do. If Λ is taken to be a very high energy scale at which new physics comes into play, we should expect physics at currently accessible energies to be very well described by a renormalizable action, as the effects of non-renormalizable interaction terms will be heavily suppressed by inverse powers of the cutoff.

This does not quite get to the heart of how the renormalization group illuminates the perturbative approach, however. After all, the fundamental puzzle about renormalized perturbation theory was why it produces accurate approximations at all, even granting the renormalizability of the interactions under consideration. I suggested above that the perturbative approach is incapable of answering this question on its own because it lacks a cogent characterization of the systems and quantities it is supposed to be approximating. The renormalization group framework arguably fills this lacuna, providing a non-perturbative framework that

can justify the various steps that go into the perturbative renormalization procedure.

The first step of the perturbative renormalization procedure, you will recall, was to replace the divergent momentum integrals in the perturbative coefficients with finite cutoff expressions. Since the cutoff is removed at the end of the calculation this was typically understood in purely formal terms in the original perturbative treatment of QFT. The renormalization group framework provides a physical interpretation of what is going on here, however. We have seen that it is possible to simulate the effects of high momentum degrees of freedom not included in a cutoff model by tuning the system's dynamics. We can thus legitimize the perturbative cutoff on momentum space integrals in the same terms, on the understanding that the effects of physics beyond the cutoff can be absorbed into an 'effective' action.

The second step of the procedure was to redefine the expansion parameter so as to remove the diverging dependence on the cutoff in the perturbative coefficients. In the original incarnation of the renormalization method, this was viewed as a necessary step to extract sensible predictions from the perturbative formalism. The renormalization group analysis of the low energy regime provides a physical justification for removing the hypersensitivity to the cutoff, however. We saw that low energy physical quantities, and in particular, the scattering cross sections that are typically the targets of perturbative calculations, actually are weakly dependent on the cutoff at low energies. This amounts to a non-perturbative demonstration that the logarithms and powers of Λ that appear in naive perturbative expansion are artefacts of an inappropriate choice of expansion parameter. Once the expansion parameter has been renormalized in the manner described in the previous section the only dependence on the cutoff that remains in perturbative approximations takes the correct form of powers of E/Λ . The perturbative renormalization procedure, on this reinterpretation, is fundamentally about ensuring that our approximations have the right scaling behaviour, not about ensuring that they are mathematically well behaved as the cutoff is removed. Choosing the expansion parameter so as to minimize the dependence on the cutoff can simply be understood as a matter of ensuring that truncations of the series mimic the behaviour of the physical quantity that they are supposed to approximate.

The final step of taking the cutoff to infinity also finds a natural justification in the renormalization group setting. Assuming that the cutoff scale is much higher than the energy scale we are trying to describe, the E/Λ cutoff dependence of renormalized perturbative approximations, and the actual physical quantities they are supposed to approximate, will be very small. In many contexts, they will be much smaller than expected experimental error and can consequently be justifiably ignored. What we are doing when we take the cutoff to infinity in perturbative calculations is neglecting the residual dependence on the cutoff. Since removing the cutoff in the perturbative context has significant computational benefits, and the renormalization group gives us a handle on the kind of errors that result from doing so, it is pragmatically justified.

This is only a sketch of how the renormalization group illuminates the original perturbative QFT formalism and more work is clearly needed to develop this story in detail.⁸ Still, the pieces seem to be in place to explain the success of perturbative predictions in terms which should satisfy the realist. We also find ourselves in a better position regarding the epistemic achievement component of realism in the post renormalization group context. The shift towards an effective field theory perspective on QFT points to a way of extracting knowledge of the unobservable from empirically successful QFT models.

The central problem here was the lack of a clear answer to the question of what empirically successful QFTs are—both mathematically and physically. We saw in the previous section, however, that it is possible to precisely define the path integral for the partition function and explicitly construct realistic QFTs as mathematical models if the degrees of freedom of the field associated with arbitrarily large energies and momenta are removed via the cutoff. This provides a non-perturbative characterization of QFT which has a crucial advantage over the axiomatic systems discussed in section 2; we can explicitly write down cutoff formulations of empirically successful QFTs, and the standard model in particular. Furthermore, we have seen good reasons to regard these cutoff models as conceptually respectable, and empirically successful, theories in their own right. As Wallace (2006, 2011) has argued, this suggests that we should be looking to cutoff QFT models when it comes to the question of what we ought to believe about the world given the empirical successes of high energy physics.⁹

What would it mean exactly to be realist about a cutoff QFT? Following Williams (2017), we can make sense of this along the lines of the selective realist programme introduced in section 2. The basic idea behind this approach was that the realist ought to take a differentiated attitude towards the content of a successful theory, saving their optimism for those theory constituents that underwrite its predictive success. The renormalization group comes into its own again here, providing a powerful tool for developing this sort of selective realist reading of a QFT.

On the one hand, it allows us to identify features of cutoff QFT models that we should not take representationally seriously. Much of the empirical success of the standard model takes the form of predictions of cross sections for scattering events

⁸One potential worry here is that, since perturbation theory itself is often used to analyse the renormalization group flow—and the Polchinski results mentioned in the previous section turns on a small coupling assumption—there is a danger of circularity in appealing to the renormalization group to rationalize the perturbative formalism. I think this objection can be rebutted, but discussing this point in detail is beyond the scope of this chapter. Thanks to Laura Ruetsche for raising this issue.

⁹Taking cutoff models to provide an adequate characterization of the mathematical and physical content of QFT is controversial. Doreen Fraser (2011), in particular, has argued against this move and defended the superiority of axiomatic formulations of the theory. Note, however, that taking cutoff QFTs to be appropriate objects of realist commitment does not imply that the axiomatic approach has nothing to offer philosophically. There are arguably other philosophical issues raised by the QFT programme that are most naturally addressed in the context of the axiomatic tradition. For more on this pluralist approach to the dispute surrounding the formulation of QFT see James Fraser (2016).

produced in particle colliders, with the current energy limit being at the order of 10^{13} electron volts. The renormalization group results just discussed reveal that many features of current QFT models do not really make a difference to these empirical successes, in the sense that they can be varied without much affecting scattering cross sections at relevant energy scales. For one thing, it establishes that these quantities are highly insensitive to the imposition of a cutoff at some much higher energy scale, as well as the details of how this is done. We can also vary the dynamics of a model at the cutoff scale without affecting its predictions; adding small non-renormalizable interactions to the standard model action, for instance, does not undermine its empirical adequacy. What this tells us is that many of the claims QFT models make about the world at the fundamental level do not contribute to, and consequently are not supported by, the predictive successes of modern particle physics.

On the other hand, the renormalization group helps us articulate positive theoretical commitments that are supported by the success of the standard model. As well as sharing empirical content, the classes of QFT models that flow to the same surface under the action of the renormalization group transformation arguably make common claims about relatively large scale, non-fundamental, aspects of the world. They agree, for instance, about low energy correlation functions—expectation values of products of field operators associated with well-separated space-time regions. These quantities are preserved by the renormalization group coarse-graining transformation and encode the long distance structure of a QFT model. They are also directly connected to its successful predictions—unlike the theoretical features mention above, you cannot vary the long distance correlation functions of a theory without drastically affecting its low energy scattering cross sections. Furthermore, the renormalization group tells us that these quantities are extremely insensitive to the details of physics at very high, currently inaccessible, energy scales. It thus demonstrates that they are, at least in one sense of the term, robust, another quality which selective realists often take to motivate belief in a theoretical claim (Wimsatt 2007, Williams 2017).

The idea then is that the information the renormalization group provides about the dependencies that hold between theoretical claims at different scales allows us to sort those we should take representationally seriously from those we should not. The picture that emerges from this analysis is that QFTs enjoy a kind of coarse-grained representational success, capturing some (relatively) long distance, low energy, features of the world without limning its fundamental structure. This accords with the effective field theory methodology, but, crucially, the claim that the standard model is an effective theory is not read as meaning that it is purely phenomenological. It furnishes genuine extra-empirical knowledge on this view—just not about the fundamental.

5 Challenges

Drawing on renormalization group resources in the way I have just outlined is, I think, our best hope for assimilating QFT into a traditional realist view of science. Many questions remain about how this proposal should be fleshed out and defended, however. In the interest of sharpening the position, or, more accurately, identifying those areas where it needs sharpening, I want to conclude by addressing some of Laura Ruetsche's recent objections (Ruetsche, 2018). I will touch, in particular, on two worries raised by her discussion that point to the need for further work in an epistemological and metaphysical/semantic direction respectively.¹⁰

The first objection targets my claim at the end of the last section that results about the low energy behaviour of the renormalization group flow give us grounds to be confident in some coarse-grained properties of QFT models. The worry is that this move falls foul of familiar anti-realist arguments concerning historical theory change. Anti-realists have long pointed to the plethora of predictively successful yet false scientific theories in the historical record as a challenge for realism. A recent incarnation of this sort of argument, which is particularly pertinent in the present context, is Kyle Stanford's (2006) problem of unconceived alternatives. When we look at the history of science, according to Stanford, we find that scientists have repeatedly failed to conceive of rival theories that were just as well supported by the available empirical evidence as the theories they accepted. Inductively then, we should expect present scientists, and scientific theories, to be in the same boat, undermining our confidence in current extra-empirical scientific claims.

The renormalization group results discussed above might seem to offer respite from this problem. The standard model will likely someday be replaced by a new theory that describes the physics of currently inaccessible energy regimes. Since the low energy features of the standard model are highly insensitive to what is going on at these higher energy scales, however, one might think that they are very likely to be retained by these more fundamental theories, whatever they end up looking like. As Ruetsche points out, however, this line of response rests on the assumption that future theories can be situated in the space on which the renormalization group transformation acts and this is not obvious a priori. Rigorous renormalization group results, such that there are, deal with rather circumscribed spaces of theories. Furthermore, the QFT framework itself is expected to break down and be replaced by a radically new quantum gravity theory as the Planck scale is approached. Who is to say this theory can be treated within the renormalization group framework? In fact, if we accept the moral of Stanford's induction, it seems that we should actively expect future theory change to outstrip the reach

¹⁰Another line of objection, which will not be dealt with here, is Doreen Fraser's (2011) claim that renormalization group results reveal widespread underdetermination in QFT and therefore push against realism, rather than coming to its aid. I address this underdetermination worry in James Fraser (2018).

of renormalization group considerations.

How to respond? We should concede, I think, that renormalization group results do not defeat historically generated skepticism on their own.¹¹ Indeed, it is difficult to see how local scientific arguments could ever sway a thoroughgoing anti-realist who is already convinced of their unreliability. It is important to distinguish two tasks facing the scientific realist here, however: spelling out the content of their epistemology and defending it. This chapter has focused on the former; I presupposed a realist perspective and advanced a strategy for articulating two central realist theses in the QFT context. When it comes to the project of arguing for a realist view of science more general epistemological issues that have not been touched on here need to be considered, the implications of historical theory change being prominent among them. A flat-footed response to this sort of attack then is to admit that the plausibility of the programme set out in this chapter is conditional on a successful rebuttal of historical arguments for anti-realism.¹²

Having said that, it is possible that the renormalization group might play some role in the broader project of defending realism. Here is a different way it might feature in a response to Stanford's unconceived alternatives problem, for instance. One supposed advantage of Stanford's argument is that, whereas traditional incarnations of the pessimistic induction generalise over scientific theories, and are therefore vulnerable to the response that false historical theories were unlike currently accepted ones in relevant respects, his induction generalises over scientific theorists. Stanford suggests that it is more difficult to drive a wedge between the inferential capacities of past and present scientists—surely current scientists are no better than luminaries like Darwin and Maxwell at conceiving of relevant alternatives to their theories (Stanford, 2006, 43)? One might reject this assumption, however, claiming that theoretical and methodological progress in the sciences has improved our ability to probe the space of possible theories. The renormalization

¹¹One might try to resist this conclusion by marshalling further theoretical considerations which suggests that the renormalization group apparatus can encompass the class of relevant alternatives to the standard model. We might point, for instance, to Weinberg's (2004) 'folk theorem', which says that any theory that satisfies some general principles, like cluster decomposition and (approximate) Lorentz invariance, will have to recover the basic theoretical structure of QFT at low energy scales (see also Duncan, 2012). This suggests that, no matter how radically novel a future theory of quantum gravity is, its low energy limit should be a QFT that is amenable to the kind of renormalization group analysis described in section 3. Again though, this argument rests on a posteriori scientific premises, so unrestrained historically generated skepticism will also throw it into doubt.

¹²Godfrey-Smith (2008), Chakravarty (2008), Ruhmkorff (2011) and Devitt (2011) are some responses to Stanford's unconceived alternatives problem which the aspiring realist about QFT might draw on. The strategy mentioned below, of urging that science is improving in its ability to rule out alternative theories, is suggested in a number of these papers. I should note that, if it turns out that the epistemic achievement component of realism has to be abandoned in light of arguments like Stanford's the renormalization group might still have something to offer a more minimal realism of the kind advocated by Saatsi (2015). Saatsi suggests that the most promising strategy for articulating the realist position is to point to exemplars, from contemporary science and historical record, of a how a theory's latching onto the world could explain its empirical success. We could thus treat the story told in section 4 as an exemplar of this sort without committing ourselves to its truth.

group is arguably an instance of this phenomenon. Even if it is not exhaustive in its scope, the information the renormalization group provides about how modifications of a theories dynamics at one scale affect its behaviour at others seems to be a novel epistemological resource that was not available to previous generations of physicists. While Ruetsche is right to flag the danger of exaggerating the power and generality of the renormalization group approach then I do not think her discussion gives us reason to abandon the realist programme advanced in this chapter. It does point to a need for a more careful examination of the epistemological mileage that can really be extracted from renormalization group results, however.

The second worry I want to consider concerns the nature of the representational success allegedly enjoyed by effective field theories. Even granting that some of the low energy claims of current QFTs are supported by renormalization group considerations, Ruetsche (2018) suggests, it is not obvious that they really concern the unobservable. This line of attack exploits the fact that renormalization group enthusiasts are often rather vague about the features of the world they take effective field theories to be latching onto. Furthermore, when they do try to make more explicit commitments they seem to be vulnerable to reinterpretation in empiricist friendly terms. In the previous section, I identified low energy correlation functions as plausible examples of the sort of theoretical quantities we ought to be realist about. But correlation functions are intimately related to scattering cross sections measured in collider experiments (via the LSZ reduction formula). If correlation functions can be interpreted as merely encoding information about the empirical signatures of scattering events measured at collider experiments then the view of QFT I have been developing threatens to collapse into a form of empiricism.

Again, I do not think this line of objection is fatal, but it does add urgency to a central challenge facing our realist programme. In order to substantiate the claim that effective field theories are capturing unobservable aspects of the world, in accordance with the epistemic achievement thesis, we need a precise characterization of the non-fundamental descriptive content of QFT models. This is a difficult problem for, at least, two reasons.

First of all, the status of non-fundamental entities and properties is a controversial issue in its own right. As Williams (2017) points out, pervasive methodological assumptions about the nature of the interpretive project have led philosophers of physics to focus on the fundamental ontology posited by physical theories. As a result, non-fundamental physical ontology is undertheorized and some may even doubt that there is such a thing, adding fuel to the suspicions about the realist credentials of my programme. In philosophy more generally, there is no agreed-upon framework for regimenting claims about the non-fundamental, with different approaches gaining currency in different sub-disciplines. Philosophers of science in the structural realist tradition (Ladyman and Ross, 2007; Wallace, 2012) have appealed to the Dennettian notion of ‘real patterns’, for instance, while analytical metaphysics has seen an explosion of interest in notions of metaphysical dependence and grounding (Bliss and Trogon, 2014). There are certainly frameworks

on the table that the effective field theory realist might turn to in order to clarify their position then, but doing so will only be as reassuring as these approaches are well motivated.¹³

A second obstacle here is the perennial controversy surrounding the physical content of quantum theories. I have skirted around the measurement problem thus far but the would-be realist about QFT clearly cannot put this issue off indefinitely. In order to rebut suggestions that low energy correlation functions have no extra-empirical content we seem to need a physical interpretation of quantum operators and states, and different approaches to the measurement problem give different answers to the question of what the quantum world is basically like.¹⁴ Furthermore, in addition to the puzzles inherited from non-relativistic quantum mechanics, QFT raises interpretive problems of its own. As Ruetsche (2011) argues detail, the existence of unitarily inequivalent Hilbert space representations of a quantum theory with infinitely many degrees of freedom arguably poses a novel interpretive challenge.¹⁵

All this suggests that the realist strategies advanced in this chapter will have to be re-examined alongside traditional debates about the interpretation of quantum theory and groundwork on the notions of the fundamental and derivative before the prospects of realism in high energy physics can be fully assessed.

References

Butterfield, Jeremy. and Bouatta, Nazim. (2015), “Renormalization for Philosophers.” in Bigaj, Tomasz and Wütrich, Christian (eds) *Metaphysics in Contemporary Physics*, Brill.

Bliss, Ricki, and Trogon, Kelly. (2014) ”Metaphysical grounding.” in Stanford Encyclopedia of Philosophy, <https://plato.stanford.edu/entries/grounding/>

Chakravartty, A. (2008), “What You Don’t Know Can’t Hurt You: Realism and the Unconceived”, *Philosophical Studies*, 137(1), 149-158.

Devitt, M. (2011), “Are Unconceived Alternatives a Problem for Scientific Realism?”, *Journal for General Philosophy of Science*, 42(2), 285-293.

¹³One problem here is that many discussions of the fundamental derivative distinction in the metaphysics literature are sufficiently far removed from scientific practice that it is difficult to see how to apply them to actual physics. Discussions like McKenzie’s (2014) application of Fine’s notion of grounding to quantum theory show that this gap can be bridged, however.

¹⁴Wallace’s realism about QFT is wedded to his Evertianism. It is worth noting that Everettian interpretations of quantum theory are perhaps the most natural fit with the approach to realism about QFT described in this chapter since other approaches to the measurement problem which modify the basic structure of quantum mechanics, such as Bohmian and dynamical collapse models, cannot currently be extended to empirically successful QFTs.

¹⁵Indeed, she claims that unitarily inequivalent representations ultimately pose a problem for a realist view of QFT. See James Fraser (2016), chapter 6, for a response to this challenge.

- Duncan, Antony. (2012), *The Conceptual Framework of Quantum Field Theory*. Oxford University Press, Oxford.
- Haag, R. (1996), *Local Quantum Physics*, Springer-Verlag.
- Fraser, Doreen. (2008), “The fate of ‘particles’ in quantum field theories with interactions”, *Studies in History and Philosophy of Modern Physics*, 39(4), 841-859.
- (2009), “Quantum Field Theory: Underdetermination, Inconsistency, and Idealization”, *Philosophy of Science*, 76, 536-567
- (2011), “How to Take Particle Physics Seriously: A Further Defence of Axiomatic Quantum Field Theory”, *Studies in History and Philosophy of Modern Physics* 42, 126–135.
- Fraser, James D. (2016) *What is Quantum Field Theory? Idealisation, Explanation and Realism in High Energy Physics*, doctoral thesis, University of Leeds
- (2017) “The Real Problem with Perturbative Quantum Field Theory” *British Journal for Philosophy of Science*, <https://doi.org/10.1093/bjps/axx042>
- (2018) “Why Underdetermination Isn’t Always Bad” unpublished manuscript.
- Godfrey-Smith, Peter. (2008), “Recurrent Transient Underdetermination and the Glass Half Full”, *Philosophical Studies*, 137(1), 141-148.
- Halvorson, Hans. (2007), “Algebraic Quantum Field Theory” In J. Butterfield, and J. Earman (eds), *Handbook of the philosophy of science: Philosophy of physics. Part A*, Elsevier.
- Hancox-Li, Leif. (2015), “Coarse-Graining as a Route to Microscopic Physics: The Renormalization Group in Quantum Field Theory”, *Philosophy of Science* 82 (5), 1211-1223.
- Kitcher, Philip. (2001) “Real Realism: The Galilean Strategy”, *The Philosophical Review*, 110(2): 151-197.
- Ladyman, James, and Ross, Don (2007) *Every thing must go: Metaphysics naturalized*. Oxford University Press.
- Kerry McKenzie (2014) “Priority and Particle Physics: Ontic Structural Realism as a Fundamentality Thesis”, *British Journal for the Philosophy of Science*, 65(2), 353–380.
- Polchinski, Joseph. (1984), “Renormalization and Effective Lagrangians”, *Nuclear Physics B*, 231(2), 269-295.
- Psillos, Stathis. (1999) *Scientific Realism: How Science Tracks Truth*, Routledge, London.
- Ruetsche, Laura (2011) *Interpreting Quantum Theories* Oxford University Press: Oxford.

- (2017) “Renormalization Group Realism: the Assent of Pessimism” forthcoming in *Philosophy of Science*.
- Ruhmkorff, S. (2011), “Some Difficulties for the Problem of Unconceived Alternatives” *Philosophy of Science*, 78(5), 875-886.
- Saatsi, Juha. (2016), “Replacing Recipe Realism”, *Synthese*, 1-12.
- (2016b), “What is Theoretical Progress of Science?”, Forthcoming in *Synthese*.
- Stanford, P. Kyle. (2006), *Exceeding our Grasp: Science, History, and the Problem of Unconceived Alternatives*, Oxford University Press, Oxford.
- Streater, Raymond F. and Wightman, Arthur S. (1965), *PCT, Spin and Statistics, and All That*, New York and Amsterdam: Benjamin Inc.
- Van Fraassen, Bas. (1980), *The Scientific Image*, Clarendon Press.
- Wallace, David. (2006), “In Defence of Naiveté: The Conceptual Status of Lagrangian Quantum Field Theory”, *Synthese*, 151, 33-80
- (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
- (2012), *The Emergent Multiverse: Quantum Theory According to the Everett Interpretation*, Oxford University Press.
- Weinberg, Steven. (2004) “What is Quantum Field Theory, and What Did we Think it is?” in Tian Yu Cao (ed) *The Conceptual Foundations of Quantum Field Theory*, Cambridge University Press, Cambridge.
- Williams, Porter. (2017), “Scientific Realism Made Effective”, forthcoming in *British Journal for the Philosophy of Science*.
- Wilson, Kenneth. G., and Kogut, J. (1974), “The Renormalization Group and the ϵ Expansion”, *Physics Reports*, 12(2), 75-199.
- Wimsatt, William C. (2007), *Re-engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*, Harvard University Press.